Security and Privacy of Overlay-based ICN/IP Coexistence Approaches

Cesar van der Poel Chhagan Lal

Mauro Conti

June 27, 2021

Abstract

Information-Centric Networking (ICN) is a networking paradigm proposed to replace the current IP network. It uses in-network caching to enhance availability. However, as a clean slate approach is unlikely to work, an architecture that allows for the two paradigms to coexist needs to me used to facilitate the transition. Several such architectures have already been designed, but further research is needed to make sure a transition into such an architecture does not pose any risks to the privacy or security of users of the internet. The aim of this paper is to identify several important privacy and security requirements and apply these to several coexistence architectures. The focus of this paper is on underlay-based coexistence architectures, which map IP packets to ICN packets to "tunnel" IP over ICN. The architectures were compared based on the requirements they satisfy and their modes of deployment. The investigated architectures largely supported the same set of requirements to approximately the same extent. However, many important requirements were not or only partially supported by one or more of the architectures. Further research is needed to compare these architectures to overlay-based and hybrid architectures.

1 Introduction

In the early days of the internet, an architecture that was convenient for the use at the time was implemented. However, this host-centric network is not properly suited for the challenges the internet faces today, and does not scale very gracefully [3]. In order to overcome these problems, a new form of networking, called Information-Centric Networking (ICN) has been designed. It uses in-network caching to allow for easier content delivery and less strain on content servers as well as important connections. This form of networking, however, provides its own set of challenges and problems [1]. Furthermore, as we cannot just "reset the internet", we need to devise a way to transition from the current IP-architecture to the new ICN-architecture.

Several ways have been proposed to ensure coexistence between the two paradigms [5]. These can be divided into overlay-based approaches, where ICN is performed on top of the existing IP network, hybrid approaches, where the entire network supports both ICN and IP, and underlay approaches, where IP packets are mapped to corresponding ICN packets to

allow IP to be tunneled over ICN. The network should be safe to use; to confirm this is the case, the architecture needs to support certain important privacy and security requirements.

In this paper, the following question will be answered:

"What privacy and security features are or are not supported by underlay-based ICN-IP coexistence architectures?"

To answer this question, I will first identify three underlay-based architectures, and describe their workings. After this, important security and privacy requirements that need to be satisfied by an architecture will be identified. These will then be applied to the architectures to determine which of them are supported. In the conclusion, these results will be summarised and the most important differences and trade-offs between the architectures will be provided. This will allow future researchers to identify the architecture they consider "safest" based on their priorities, or to identify security and privacy features that are important but not supported by these architectures. It will also allow those working with the architectures to identify their weaknesses, so they can keep these into account.

This paper is structured in the following way. Section 2 will provide a general overview of the underlay architectures to be investigated, including their principles and workings. Section 3 will provide an overview of the privacy and security requirements to consider, after which the extent to which the architectures support these features is investigated. Once this will be done, section 4 will reflect on the ethical aspects and reproducibility of this research, after which section 5 will provide some discussion on the subject and propose future research directions. Finally, in section 6, the conclusions of this paper will be summarised, and the benefits and downsides of the different architectures will be compared to each other.

2 Background and Related Works

There are three different approaches for ICN/IP coexistence. The overlay approach uses the current IP network to transport ICN messages between different ICN-capable devices. The hybrid approach requires all network nodes to have both ICN and IP protocol stacks to support both protocols. Thirdly, there is the underlay approach, which uses gateways to convert IP packets into corresponding ICN packets. This conversion is reversible, essentially allowing for an IP tunnel over the ICN network. This third approach is the one investigated in this paper.

There are currently four different underlay-based architectures. These are the CCN solution proposed by CableLabs [23], the NDN-based DOCTOR [6], PURSUIT evolution POINT [13] and finally RIFE [15], which is also based on PURSUIT. RIFE will not be discussed in this paper as its main focus is on interconnecting remote networks in a delay- and disruption tolerant way, rather than improving the currently existent internet architecture. The focus of the other architectures is rather more towards "upgrading" the current internet to be more scalable and secure, which seems a more appropriate goal to research in this paper.

2.1 CableLabs

The goal of CableLabs was to analyze the technical problems involved in using Content Delivery Networks (CDN), and see how these can be overcome by using Content Centric Networking (CCN) [23]. In the paper, an incremental transition approach is considered, in which CDNs are replaced by CCN clusters. An example of such an island can be seen in Figure 1.

The CCN protocol uses Interest packets, sent by clients, to request data, and Content packets, sent by caches or origin servers, to return the requested data. The first contains the name of the resource that is being requested, while the latter contains (a part of) the resource as well as metadata containing a signature and a key to confirm the integrity of said resource. Such a key is certified by a trusted authority to prevent forgery [11] [23].

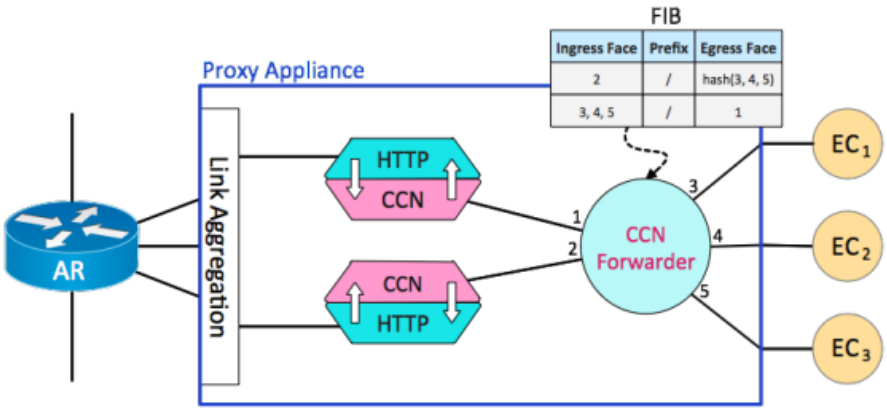

Figure 1: Detail of a CableLabs cluster, where the AR node connects the IP network (left) with the CCN network (right), as described in [23].

A router in CCN has several functions. First of all, it routes the request based on the longest prefix in its routing table [23] [11]. The interface through which a request is received is stored. Once a response is received, this response is forwarded over that same interface to route the response to the requester. Another function of CCN routers is interest aggregation; this means that when multiple interest packets for the same resource are received, only one is forwarded, and the response is returned to all requesters [11]. This reduces load on the network between this router and the origin server. A third (optional) function of CCN routers is caching. When new content passes through, the router may decide to store this in its cache (or Content Store) in order to serve subsequent requests [11] [23].

In order to apply the CCN architecture in the current IP network, CableLabs proposes to place numerous CCN clusters in the existing IP network, equipped with gateways that translate a HTTP request to a CCN Interest packet and feed it into the CCN network. If the requested content is not found in the CCN network, another gateway requests the content from the (IP) origin server. In this case, the response is mapped to a CCN content object. After this, or if the content was already cached, it is returned to the original gateway and translated back into an HTTP response [23]. From the IP network's perspective, this island functions as a large caching server. By increasing the amount of clusters and their sizes, this slowly transforms the original IP network into (nearly) a full CCN network.

2.2 DOCTOR

DeplOyment and seCurisaTion of new functiOnalities in virtualized networking enviRonments is a coexistence approach that leverages Network Function Virtualization (NFV) to enable easier deployment and coexistence of new networking architectures. Several of these functions can be run in parallel, as can be seen in Figure 2. In the testbed of the project, a coexistence of IP- and NDN-network stacks was used [6].

The NDN architecture is very similar to CCN; it uses Interest and Content (or Data) packets, which can be forwarded, aggregated in the case of interests, and cached in the case of Data packets [24]. The main difference between these two architectures is the organisations working on it [12].

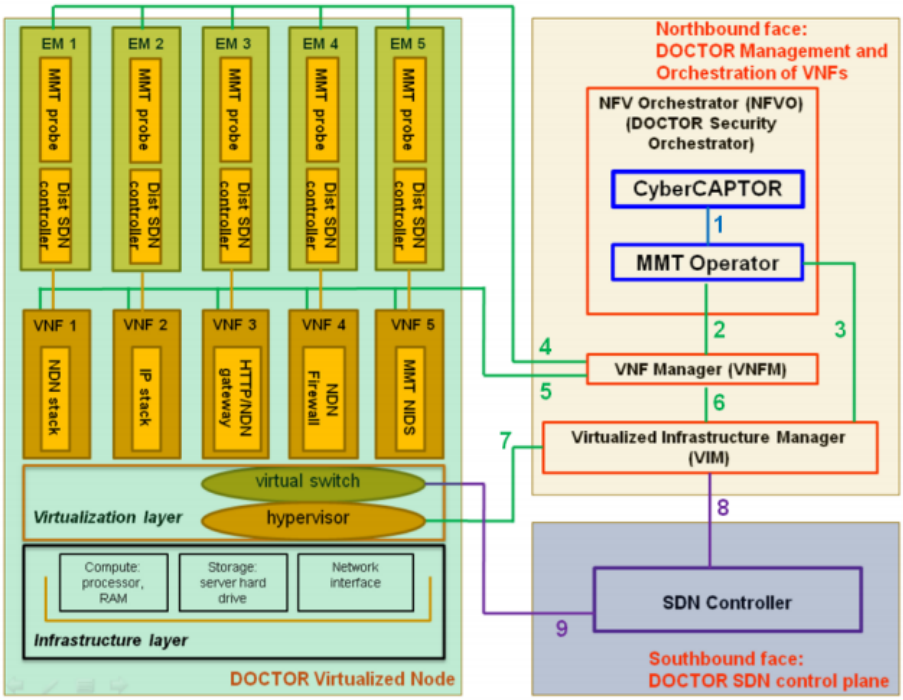

Figure 2: The architecture of the DOCTOR node, as described in [21].

DOCTOR allows all network nodes to have multiple (virtual) network functions; a stack for each protocol it supports, and gateways between them [21]. The support for multiple stacks means that, if needed, NDN parts of the network could still support IP without tunneling. No specific structures dividing the network into NDN and IP were described.

2.3 POINT

The POINT architecture is based on the BlackAdder project [18] and mainly applies some changes to the protocol stack [19]. It uses a publish/subscribe paradigm for its ICN section.

BlackAdder is an ICN implementation that is part of the PURSUIT project [18]. It defines the following functions: publication, subscription, rendezvous, topology management and forwarding. Data is requested via the subscription function; this can be done before publication. Once the content is available, the rendezvous function locates the content by recursively checking caching nodes until the content is found or the origin server is reached. The topology management function identifies a route from content to requester, and forwarding nodes transport the content along this route [2] [18]. A schematic of the node with its various functions can be seen in Figure 3.

Although at first glance the approach of this architecture seems to be vastly different

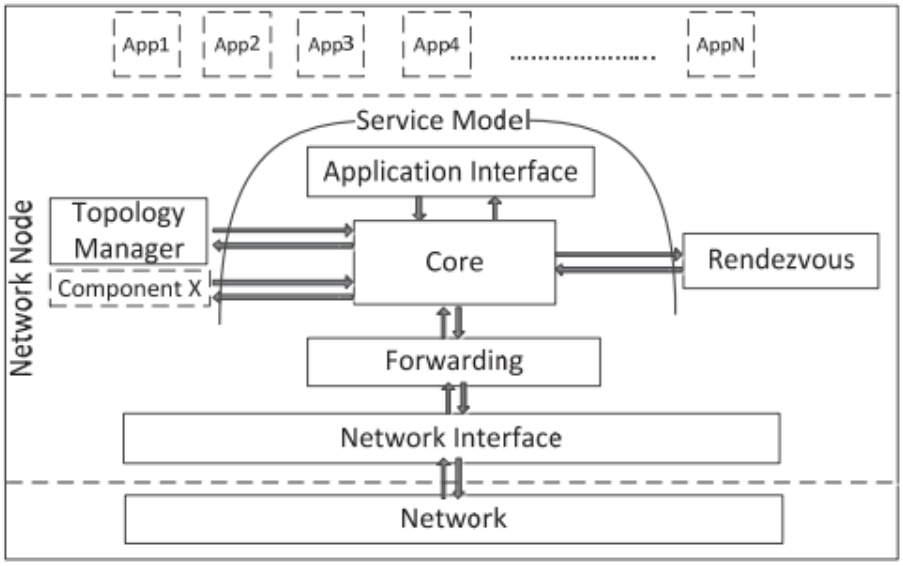

Figure 3: The architecture of the BlackAdder node used in POINT, with an overview of the interactions between the different functions of the node. Further described in [21].

from CCN and NDN, a thorough understanding of it will show that they are still rather similar. With normal usage, the main differences are the fact that in this approach, the rendezvous and topology management functions locate and route the content instead of the routers in the network. Secondly, this approach allows for content to be requested before it is available; this simply means they do not receive it right away [18].

3 Privacy and Security Evaluation

In this paper, four privacy aspects will be considered: anonymity, meaning different nodes are indistinguishable; request secrecy, meaning there is no definitive record of previously performed requests; confidentiality, meaning secret content cannot be seen by other parties; and unlinkability, meaning independent packets cannot be shown to be linked to each other. Furthermore, four security aspects will be considered: availability, meaning requested content can be served within reasonable time constraints; integrity, meaning nodes can verify the correctness of content; access control, meaning some can access certain content while others cannot; and non-repudiation, meaning a publisher cannot unpublish publishments.

In the following subsections, these aspects will first be provided with a more detailed definition. After this, the ICN and IP paradigms, as well as the coexistence architectures, are investigated to see what effects their designs have on the satisfaction of these requirements.

3.1 Privacy

As stated in the previous paragraph, four privacy requirements will be used to evaluate the architectures. Before seeing how much these architectures support the mentioned requirements, I will elaborate on what exactly these four entail:

- *Anonimity* - What is meant by anonymity is that what specific device sends out a request should not affect the response to it. If a device sends out a certain request to a network in a certain state, it should receive the exact same response any other device would receive to said request in said network.
- *Request secrecy* - Request secrecy entails that any device requesting certain content can plausibly deny having requested that content. This means that no evidence of a request must be left behind.
- *Confidentiality* - Confidentiality means that data can only be read by the intended receiver(s), and not third parties. If a device sends out secret information intended for some other device, only that device should receive this information.
- *Unlinkability* - Unlinkability in this setting is the inability to link multiple requests (or other data points) to each other. Even if the above requirements make sure we can't link these requests to one specific device or user, a certain pattern in requests might allow for indirect identification.

3.1.1 Internet Protocol

As IP is currently used, and has been used since the internet's early days, all lacking privacy requirements can be directly exploited in a heavily used public network. This part of the paper will investigate how private this has made IP.

Anonymity The IP network routes packets based on IP addresses specified on them [14]. This has the unfortunate side effect that all IP traffic contains information specific to its sender (and receiver). Transmitted over a public network, anyone can view this information, and thus link these requests to the device with the given IP address. **In short, the fact that IP addresses are on each IP packets violates anonymity in IP networking.**

Request secrecy Once an IP packet is sent into the network, routers forward it based on its destination IP address [8]. Although routers could decide to record what packets move through them, they do not normally do this, although a malicious router could. **As internet traffic is generally not recorded by IP routers, we can conclude request secrecy is mostly satisfied.**

Confidentiality In IP, connection-based encryption is used to provide confidentiality [23]. This encryption ensures that only the two parties using the connection can access the decrypted data. As a result, any unrelated third parties cannot retrieve this information from the network. **As the encrypted connection makes sure the data is only available to the sender and (intended) receiver, confidentiality can be considered satisfied in IP networks.**

Unlinkability As mentioned previously, IP packets contain the source and destination IP addresses [14]. Packets sent to or from the same device contain the same IP address, meaning they can be linked based on this information. **As IP addresses allow for linking multiple requests to each other, we cannot consider unlinkability satisfied for IP networks.**

3.1.2 Information-Centric Networking

Although pure ICN will not be applied to the internet, it is still interesting to see what privacy and security requirements apply to it to which extent. This will help to investigate what effect a certain coexistence architecture has on these requirements.

Anonymity According to [2], ICN packets can be routed in various ways. Some of these (reverse request path) have no need for identifying information, fully supporting anonymity. Others (IP connection) do carry such information, thus violating this property. However, as there are some ICN architectures that are fully anonymous, it is fair to say ICN technically does support anonymity. **Although some designs violate anonymity, several other designs have proven ICN can fully support anonymity.**

Request secrecy A key element of ICN is in-network caching to serve content more easily [3]. However, this has the negative side effect that the router nearest to a device will cache content requested by said device, and possibly only said device. Such an "edge router" would contain a definitive record of the request history of this device. **In ICN, request secrecy is violated as edge routers near a device contain a record of the activity of said device.**

Confidentiality As stated previously, confidentiality can be achieved through encryption [16]. In ICN, this is achieved by encrypting content objects [3]. This way, only devices with the correct decryption keys have access to the decrypted data. As long as there are secure ways to exchange keys between the intended parties ([7]), this excludes any external parties potentially interested in the data. **Content-based encryption allows for confidentiality in ICN networks.**

Unlinkability Depending on the extent to which a specific ICN design supports anonymity, the lack of it could be used to link requests. However, under the assumption anonymity is satisfied, this is not an option. A different potential approach could be to investigate a router through which a specific device requests content and conclude links based on the cached content. However, this can only be done when one has access to said router, which will not be the case in most (realistic) situations. **As there is no feasible way of linking multiple packets reliably, unlinkability is mostly satisfied by ICN networks in the general case.**

3.1.3 CableLabs

As previously described, CableLabs uses CCN networks as "caching servers" in an IP network; IP devices can connect to "islands" of CCN nodes via the gateways. Although its incremental approach will eventually allow devices to use CCN-native operations, this means that in earlier versions both servers and clients will be surrounded by the IP network. In this part of the paper, I will see what effects this has on the mentioned privacy features.

Anonimity In the CableLabs architecture, nodes in the IP network communicate with gateways [23]. This communication is not anonymous, as it carries IP addresses, but as this communication is with (trustworthy) CableLabs gateways we can assume it does not alter the response based on this information. As this gateway "unpacks" the request, the

communication from this point forward is in no way connected to the original requester. **Anonymity is mostly satisfied by CableLabs, as the identifying data (the IP address) is removed at the gateways of the CCN cluster.**

Request secrecy As stated in section 3.1.2, request secrecy is violated in ICN due to edge routers containing cached data of the nearest device. However, as such edge routers only exist inside the CCN clusters, one would still be shared by a relatively large amount of devices, thus not providing a lot of information on what device is responsible for what content. This is combined with the fact that the IP network outside the clusters mostly supports request secrecy. **As the "edge routers" of the CCN clusters serve many devices, request secrecy can be concluded to be largely satisfied by CableLabs.**

Confidentiality As described in [23], CableLabs maps IP to ICN packets and back. However, IP encryption is connection-based, while ICN encryption is content-based. [23] does not explicitly consider any way of overcoming this difference. A possible solution could be to decrypt content at any gateway, but this would have the side effect of allowing any third party with access to said gateways to receive decrypted confidential data. **As CableLabs does not consider ways of converting between ICN and IP encryptions, we cannot consider confidentiality satisfied for this architecture.**

Unlinkability CableLabs uses IP and CCN packets [23]. As described in [11], CCN packets do not contain any data that could be used for linking requests, satisfying unlinkability. In the IP sections of the network, packets could still be linked based on the IP addresses contained in them. However, this section of the network would become increasingly smaller as CableLabs reached later stages of its development. **As packets can practically only be linked using the IP addresses they contain in the shrinking IP sections of the network, unlinkability is satisfied increasingly more overtime by the CableLabs architecture.**

3.1.4 DOCTOR

The DOCTOR architecture, unlike CableLabs, does not explicitly state how the network can be divided into NDN and IP, probably because they do not believe it matters. As a matter of fact, since all their nodes support both network stacks and gateway functionalities, any node can be a gateway, meaning the boundaries between IP and NDN are as flexible as they need to be [21]. We will now have a look at what this means for the privacy features.

Anonimity As there is no predefined structure dividing the network into NDN and IP sections, different cases need to be considered individually. Packets in NDN sections, like CCN, do not contain identifying data, making any communication through such a section anonymous; assuming once again gateways are trusted, even requests starting in IP sections are. There could, however, still be IP-only communication, which is not anonymous. **As IP addresses are stripped from packets at gateways, but DOCTOR sometimes still uses IP-only communication, anonymity is largely, though not entirely, supported.**

Request secrecy As all DOCTOR nodes support gateway- and router functionalities, a gateway is presumably also a router. As such, caching nodes are very near to devices, meaning it is likely that these handle requests (almost) exclusively from one device. This violates request secrecy in sections using mainly NDN. In parts of the network with more IP (only) routers, responses are not cached, thus respecting request secrecy. Even if such requests pass a gateway, the original device is hard to determine. **Considering the fact that gateways act as "edge routers" in NDN sections, request secrecy is only partially satisfied by the DOCTOR architecture.**

Confidentiality In DOCTOR, all nodes support not only NDN and gateway functionalities, but also IP functionalities [21]. This means that a DOCTOR network would be able to provide a connection based encryption for confidential communication between two IP-compatible nodes by utilizing these IP network stacks, thus avoiding the transfer to content-based encryption. However, the authors of the relevant sources did not explicitly consider this (nor any way of transfer). Furthermore, this defeats the purpose of having an underlay-based architecture in the first place, as it prevents tunneling over ICN. **As IP stacks could be exploited to provide confidentiality, but this actively evades the use of NDN and gateways, confidentiality can only be considered partially satisfied.**

Unlinkability As described in sections 3.1.1, the IP sections of DOCTOR allow for linkability based on IP addresses. The NDN sections do not; these packets contain no identifying information [24]. Gateways use their own IP addresses, meaning linking requests of devices through them is unfeasible. As there is no particular structure for the DOCTOR architecture ([21]), this also does not affect the support for this property. **As IP-only communication violates unlinkability but communication through NDN sections supports it, unlinkability can be considered increasingly satisfied for DOCTOR.**

3.1.5 POINT

The POINT architecture does not specify any specific structure considering the boundaries of the IP and ICN networks. This means that nodes can be either in ICN sections of the network, or in IP sections of the network, and can communicate with servers in either.

Anonimity The packets used for the PURSUIT parts of the POINT network do not contain much identifying information, as with most ICN architectures. However, as seen in 3.1.1, the IP parts do not share this property. With no specific network structure to consider, the anonymity is a combination of these two properties. **As anonymity is satisfied by PURSUIT but not by IP, the POINT architecture supports anonymity for a large part, but not entirely.**

Request secrecy From section 3.1.1 we know that request secrecy is mostly satisfied for the IP sections of the POINT network, while section 3.1.2 discusses how this is not so much the case for the ICN sections of the network; PURSUIT is no exception to the information provided there. **As request history is not recorded in IP sections while it is in the PURSUIT sections, request secrecy can only be considered to be partially satisfied by the POINT architecture.**

Confidentiality [18] and [19] do not go deep into the matter of encryption; it feels safe to say they did not consider it. One source ([20]) on PSIRP, which PURSUIT is based on, did consider the matter. Their considerations amount to delegating it to the "home" rendezvous functions, which encrypts the key needed to access the confidential content specifically for the subscriber [20]. However, there is no proper way to transfer this encryption to IP networks, or the other way around. Besides this, giving a gateway access to content would allow all subscribers using this gateway to access the content. This problem is not considered by the authors of [19]. **As both types of sections of the POINT architecture support encryption but [18] does not consider any way of transfer, confidentiality cannot be considered to be supported by POINT.**

Unlinkability In PURSUIT, the route content packets take is not decided on the go by routers, but determined up front by the topology management function [18]. This function calculates a Bloom filter that tells forwarding nodes what to do. This filter stays attached to the packet until its destination is reached. Based on similarity between these filters, a (remote) attacker could link several requests with a certain probability (though they would never be certain). Besides this fact, IP address can be used to link packets in IP sections of the network as previously seen. **As IP addresses allow for linkability in IP sections and Bloom filters for partial linkability in PURSUIT sections, unlinkability is increasingly, though at best partially, supported by POINT.**

3.2 Security

Besides the four privacy requirements stated in the previous subsection, there are also four security requirements we want the architectures to satisfy. Again, before the evaluation, a more comprehensive definition of each aspect:

- *Availability* - Availability entails that any existing content that is requested is actually delivered; the content must, at all times, be available. This is, however, under the assumption that the requester has access to the requested content. Of course, in the case that the content is not meant for the entity requesting it, there is no need to deliver it, but this concept falls under access control.
- *Integrity* - Integrity is the ability of a requester to verify that the received content is correct and supplied by the correct source. Given a piece of content, we can verify that it has not been tampered with. The subgoal of verifying author is generally referred to as authenticity.
- *Access control* - Not all content should be available to all network users. Access control is the ability to provide some with said content while prohibiting others from seeing it.
- *Non-repudiation* - Non-repudiation means that a publisher cannot (plausibly) deny having published a certain piece of content it has, in fact, published. As long as the content is publicly available, anyone receiving the content will be able to see that is has been published by this entity.

3.2.1 Internet Protocol

Most cyber security knowledge available these days focusses on IP networking. From this, one would expect IP to be rather secure. In this section, the extent to which this is true will be investigated.

Availability In an IP network, content is available from the server serving the content. If this server can, for some reason, not provide this service, the content is not available. This is usually achieved through a DoS- or DDoS-attack, meaning the server is overloaded due to a larger amount of requests than it can handle [10]. Until the attack is over, the content remains unavailable. **As content is generally available, but the network is prone to DoS-attacks, availability is only partially satisfied.**

Integrity [16] describes how integrity in IP networks can be achieved through hashing. By sending both the data and the hash of said data as a response to a request, the integrity can be easily verified through a recalculation of the hash. The authenticity of the server is verified through certificates provided by trusted authorities. **As hashing provides a working integrity framework for IP networks, and certificates allow for authentication, integrity can be considered satisfied.**

Access control As described in [4], access control can be provided by IP servers. Before providing the requested content, the server can perform several checks to see if the content should or should not be served. This way, parties that should not receive certain data can be denied access. **As servers handle requests individually and can perform checks beforehand, access control is satisfied in IP networks.**

Non-repudiation As previously stated, content is only available at the server providing it. This means that as soon as a server stops providing said content, there is no longer a real record of it; certain third parties might keep track of such information, but there is no general reliable way of proving that content was once available at a server. **As there is no reliable way of proving a server once provided certain content, non-repudiation is not satisfied in IP networks.**

3.2.2 Information-Centric Networking

Although cyber security does not particularly focus on ICN, this paradigm was designed by computer science experts, presumably including people with a background in security. In this section, the extent to which these influences affected the satisfaction of our requirements will be discussed.

Availability Content in ICN starts at a certain publisher. Once it is requested, it is also available in caches along the request route [3]. These caches are able to provide the content when requested, even if the publisher is unavailable. This also means that if a DoS-attack is performed for one or a small set of content objects, the caches near the publisher will be able to serve these requests, effectively protecting the publisher from the attack. Even in the event of a successful attack, nearby caches might have stored enough content to still serve (a portion of) the regular requests. **As a DoS-attack is significantly harder to execute while also having less effect, availability can be considered mostly satisfied.**

Integrity In ICN, the integrity of data can be verified with metadata bound to the object [2]. This can, be done by signing it with the private key of the publisher. Such a signature is relatively easy to verify, while being rather difficult to forge. As the signature is a hash encrypted by the publisher, both the integrity of the data itself (by comparing hashes) and the author (by the key needed to decrypt the hash) can be verified. **As signatures provide a reliable way of verifying the integrity of content objects, this requirement can be considered satisfied for ICN networks.**

Access control Though not inherently part of the design, [7] provides a way of performing access control in ICN networks. By encrypting data with a key that all intended users have access to, the access to the content itself is controlled by the publisher. **As encryption can be used to provide a safe and reliable access control system, ICN supports access control.**

Non-repudiation As mentioned in the paragraph on integrity, ICN content is signed by the publisher for authenticity. This signature binds the content object to its publisher, at least as long as a publisher does not change its key pair, making it non-repudiable. A second possible way of repudiation is to have the content disappear entirely. This not only requires the publisher to stop providing the content, but also requires it to be removed from caches that store it; a matter not in the hands of the publisher. **As the only ways of repudiation are to change the key pair of a publisher, which can be considered unlikely, or to have all instances of the content object removed from the network, which can be considered infeasible, non-repudiation can be considered almost entirely satisfied by ICN networks.**

3.2.3 CableLabs

The specifics of the CableLabs architecture of course do not only affect the privacy features of the networking strategies it uses, but also the security aspects. Exactly how it affects them is described in the following paragraphs.

Availability The CableLabs architecture aims to increase availability of content through caching [23]. This means that this cluster might be able to return previously requested content, even if the origin server is unresponsive. The storage capacity of a cluster is not unlimited, but as it increases in size and more content is requested through it, the cluster should be able to serve increasingly more otherwise unavailable content. A potential DoS-attack through the cluster would not be feasible, as the caches serve the content without consulting the origin server, but in the case where some path through IP-only routers is used a server can still be overlaoded. **The CableLabs architecture will satisfy availability increasingly more as it is used more because it can provide previously requested content if the origin server is unavailable.**

Integrity Judging by the structure described in [23], content is generated in IP networks and mapped to CCN packets. CCN content objects use encrypted hashes called signatures for integrity and authenticity [11]. IP uses a hash, transmitted over an encrypted connection. These hashes are interchangable, meaning this integrity measure can be transferred between the protocols, but the private key needed for the signature is not available, thus not allowing for authentication. **As the integrity of the data itself can be verified through hashes,**

but the content cannot be authenticated, integrity can be considered partially satisfied by CableLabs.

Access control The IP way of access control is not applicable to the CableLabs architecture; a connection running through the CCN cluster would, as previously described, need to be decoded, possibly altering the request. Furthermore, anyone using the same gateway could exploit this access. As [23] does not consider any particular way to deal with this problem, it must be considered unsolved. **Since the CableLabs architecture does not provide a proper way of transferring IP access control to CCN access control, this requirement is not supported by the architecture.**

Non-repudiation As discussed in the part on integrity, there is no proper way for publisher signatures to be available inside the CCN cluster used by CableLabs. This means that this cluster has no effective way of providing non-repudiation. The IP network does not support non-repudiation either. **As ICN non-repudiation cannot be exploited by the CableLabs architecture and IP fails to provide it at all, non-repudiation is not supported by CableLabs.**

3.2.4 DOCTOR

The DOCTOR architecture uses networking principles very similar to those of the CableLabs architecture, but a vastly different implementation. Of course, this also affects the security of the overall network in certain ways.

Availability DOCTOR nodes can lie anywhere in the network and intercept and handle any incidental traffic that passes through them. This content can be cached, and subsequently served from cache. This spread can make sure that a lot of content passes through the caches. In the case of DoS-attacks, these caches might be able to handle a portion of the requests to protect the servers, or serve content from affected servers if the attack is successful. **As DOCTOR uses caching to provide higher availability and reduce the risk and effects of a successful DoS-attack, this requirement can be considered mostly satisfied.**

Integrity The DOCTOR architecture explicitly considers verification via signatures [22]. However, it does not mention any way of transferring this between IP and NDN sections of the network; it only uses it inside the NDN network to ensure cached content has not been tampered with. This means that, in NDN sections, the origin of content that originated in the IP network cannot be verified, but that of potential publications made directly in NDN can. Any content transferred to IP again keeps the hash, but loses the potential signature. **As the transfer of hashes allows for the integrity of content to be verified but full signatures or certificates are only occasionally available, integrity can be considered largely satisfied in the DOCTOR architecture.**

Access control Access control is only shortly discussed in the DOCTOR documents [9] and is performed using a firewall with a blacklist. This firewall restricts access to certain content for devices behind it. However, this means all devices behind a firewall share their permissions. In the case of a gateway to the IP network, this is very inconvenient. Second of all, if some malicious actor would be able to circumvent its firewall, it would have

unrestricted access, violating access control. It should be possible for the NDN network to use the previously mentioned encryption based access control ([7]), but this is not considered in the DOCTOR documents. **As the black-/whitelist approach provided by the DOCTOR authors is unsuited to provide proper access control, and ICN- or IP-only methods would fail in many cases, this requirement is considered to not be satisfied by the architecture.**

Non-repudiation As mentioned in the integrity paragraph, content published directly into the NDN network can be signed correctly, meaning that this content is largely non-repudiable. As signatures cannot be transferred from or into the IP network, content from servers inside this portion of the network is, in fact, repudiable. **As only NDN-based publishers have non-repudiable content inside the NDN network, non-repudiation can only be considered partially satisfied for the DOCTOR architecture.**

3.2.5 POINT

Despite the general idea of POINT being rather similar to that of the other two architectures, the networking principles are somewhat different. We have seen that this has some effect on the privacy features it supports; we will now also look at the security features this affects.

Availability POINT does not provide any fixed network structure to use, but the papers mainly consider a case with two (relatively large) sections of different protocols with a gateway node in between. Due to caching, DoS-attacks are harder to execute and less effective in the PURSUIT section of such a structure. They would still be feasible in IP sections, though the PURSUIT sections might reduce its effects. **As caching allows for some protection against DoS-attacks and some availability after a successful one, availability can be considered mostly satisfied by POINT.**

Integrity The sources on POINT ([19]) and PURSUIT ([18]) do not mention verification of content, but [20] describes this in the context of PSIRP, which PURSUIT is based on. This uses signatures of the publisher and the scope, to allow for easy verification using short-lived keys backed by more complex verification with long-lived keys [20]. However, as previously mentioned, there is no way of converting ICN signatures to IP verification, or the other way around, meaning that this integrity only applies to isolated PURSUIT clusters. Of course, IP-only integrity measures can be used in IP sections. **As POINT has integrity measures for both its protocols, but has no way of transferring such integrity measures between the protocols, integrity can only be considered partially satisfied by POINT.**

Access control Access control is not mentioned in the sources on POINT ([19]) and PURSUIT ([18]), but it is mentioned in the source on PSIRP. Here it is said that "The rendezvous systems are expected to perform a level of access control (...)" [20]. This is done in a way similar to the way confidential information is handled. As a result we have the same benefits, meaning that inside the PURSUIT network everything is handled correctly, but also the same downsides, meaning there is no way to transfer this to an IP section of the network without giving all nodes using a specific gateway access (see 3.1.5). **As access control in PURSUIT is achieved through encryption, and there is no way to**

transfer this encryption to IP sections of the network, nor IP access control to PURSUIT, access control is not properly supported by POINT.

Non-repudiation [20] describes that content is signed using a temporary publisher key authenticated using a more permanent scope key, meaning a publisher could be able to deny having published content if its scope is removed. This situation can occur in a PURSUIT network, as scopes can be removed once there are no external references to the elements inside the scope [18]. **Non-repudiation is not satisfied in the POINT architecture, as the keys binding content to an author are short-lived and only valid through the key of the scope, which can be destroyed.**

	IP	ICN	CableLabs	DOCTOR	POINT
Anonymity	Not	Fully	Mostly	Largely	Largely
Request secrecy	Mostly	Not	Largely	Partially	Partially
Confidentiality	Fully	Fully	Not	Partially	Not
Unlinkability	Not	Mostly	Largely	Largely	Partially
Availability	Partially	Mostly	Mostly	Mostly	Mostly
Integrity	Fully	Fully	Partially	Largely	Partially
Access Control	Fully	Fully	Not	Not	Not
Non-repudiation	Not	Mostly	Not	Partially	Not

Table 1: Compact overview of what features are supported by which architectures.

4 Responsible Research

Of course, there are various ethical aspects to this research, which will be discussed in this section. The first subsection will focus on the effects and impact this study has. The second subsection will focus on the reproducibility, and how the sources used in this paper were found.

4.1 Implications

The goal of this research is to facilitate a safe transition between the current IP network and an ICN network infrastructure. As this is the case, many parties involved with the current internet could be (indirectly) affected by this research. Most cyber security experts will not be able to apply their current knowledge anymore, and almost all current websites will have to adjust their workings. However, the effects for these parties are not necessarily negative. Cyber security experts might need some additional education, but can still apply their knowledge to ensure any potential security issues are resolved. The transition would at first cost website administrators a lot of effort, but still be beneficial in the long run because of increased scalability.

4.2 Reproducibility

As this is a literature study, the results can be reproduced by looking at the sources used to write this paper in the first place. These sources were found in different ways. First of all, some sources were provided by the supervisor and responsible professor of this research.

As they have done previous work in this field, they have the relevant expertise to identify useful and correct sources. A second type of sources used original documents describing investigated principles: for example, sources reviewed for the architectures are the documents in which they are first introduced, or in which further details are specified. As these sources are directly linked to the projects, they can be considered useful and correct as well. These sources were found by searching for the projects online, or via the references of a different paper. A final category of sources are those describing what is currently used in modern day internet. A significant part of this information could be considered general knowledge, making it easy to verify whether the information in these sources is correct. These sources were found by searching for (related) terms online.

5 Discussion and Future Work

This section will go further into what exactly this research tells us, what questions are still unanswered and what new open questions we have discovered.

Compare to other approaches This research (purposefully) only focuses on underlay-based architectures. It might be more convenient to use a different architecture, such as an overlay or hybrid one. In order to draw definitive conclusions on which type is safest, additional research needs to be done comparing the three paradigms to identify which ones do or do not support the important privacy and security requirements.

Satisfying all requirements This paper has shown that none of the investigated architectures support all requirements. This means that, if an underlay-based coexistence scheme is to be used, these architectures need to somehow be improved, or an entirely new architecture supporting more aspects needs to be designed. Thirdly, certain ICN principles could already be applied to the current IP network. If, for example, content based encryption would be used, this encryption could easily be transferred to ICN. Additionally, this could save the device serving the content encryption time, as it only needs to encrypt the content once.

Host-centric applications Some applications might be better off as a connection oriented service. Take, for example, dynamically created content, which cannot be cached but is created on the fly [17]. Though both are possible, transport over IP would be easier for this content; every ICN packet needs to be signed individually to verify the source, while an IP packet is transmitted over a connection that is verified once. Similar examples could arise for SSH tunnelling, remote desktop protocol and similar applications. Another case where IP might be more beneficial than ICN is advanced access control. Where ICN requires each individual case to be a separate content object, IP allows a single web page to be changed or expanded based on permissions, thus allowing more flexible access control.

Transition participation A point that could endanger the transition to ICN is the risk of parties refusing to use the new paradigm, as they consider IP to be good enough. This might mean that rather than supporting ICN, some parties might prefer to provide their servers with more resources instead. If this is only the case for smaller web applications, it should not pose a problem, but if larger applications draw the same conclusions there is a serious risk of incompatibility with the new network.

Deployment approach Where or how to deploy the architecture is a further subject to be thoroughly investigated. The exact approach used during deployment might have large effects on usability and acceptance of the system. Only once these effects have been investigated, the most effective approach can be determined.

6 Conclusion

The aim of this paper was to analyze different underlay-based ICN-IP coexistence approaches to see what privacy and security requirements they supported. To this end, the principles and inner workings of different instances of these architectures, CableLabs, DOCTOR and POINT, were described. After this, different security and privacy features were identified, described and applied to said architectures, to see to what extent they were supported. In doing so, it turned out that none of the architectures analyzed was able to satisfy all identified requirements. Although some requirements were satisfied by some of the architectures, most failed to provide proper ways of converting certain IP properties to their corresponding ICN features.

Most architectures were rather similar in the privacy and security aspects they support. The only significant difference is the fact that the DOCTOR architecture would be able to support confidentiality by retaining the IP stack to allow for connection-based security. Besides this, the main trade-offs between the architectures are flexibility and ease of deployment. To be more concrete, CableLabs provides a very intuitive way of deploying ICN, by placing "caching servers" in the IP network that slowly but surely extend to cover the entire network. However, as these ICN nodes are clustered together, they are (in the early phases) not distributed throughout the entire network, which means they will be less easily integrated and handle less requests. The converse is true for the other two architectures.

References

[1] Eslam G AbdAllah and Mohammad Zulkernine. "A Survey of Security Attacks in Information-Centric Networking". en. In: (2015), p. 15.

[2] Bengt Ahlgren et al. "A Survey of Information-Centric Networking". en. In: *IEEE Communications Magazine* (2012), p. 11.

[3] Fernando Almeida. "Information Centric Networks â Design Issues, Principles and Approaches". en. In: (). URL: `https://www.academia.edu/19555977/Information_Centric_Networks_Design_Issues_Principles_and_Approaches` (visited on 05/18/2021).

[4] John Barkley et al. "Role Based Access Control for the World Wide Web". In: (May 1997).

[5] Mauro Conti. "The Road Ahead for Networking: A Survey on ICN-IP Coexistence Solutions". en. In: (), p. 28.

[6] *DOCTOR project*. URL: `http://www.doctor-project.org/index.htm` (visited on 04/29/2021).

[7] Jun Kurihara, Ersin Uzun, and Christopher A. Wood. "An Encryption-Based Access Control Framework for Content-Centric Networking". In: (2015). URL: `http://sprout.ics.uci.edu/projects/ndn/papers/ccnac15.pdf`.

[8] Ravi Malhotra. *IP Routing*. en. Google-Books-ID: UyarMUrGKuQC. "O'Reilly Media, Inc.", Jan. 2002. ISBN: 9780596002756.

[9] Bertrand Mathieu et al. "Design and specification of DOCTOR MANagement and Orchestration of security remediations and countermeasures". In: (July 2018). URL: http://www.doctor-project.org/outcome/deliverable/DOCTOR-D3.1.pdf.

[10] J. Mirkovic, G. Prier, and P. Reiher. "Attacking DDoS at the source". In: *10th IEEE International Conference on Network Protocols, 2002. Proceedings.* ISSN: 1092-1648. Nov. 2002, pp. 312–321. DOI: 10.1109/ICNP.2002.1181418.

[11] Marc Mosko, Ignacio Solis, and Christopher Wood. *CCNx Semantics*. en. URL: https://tools.ietf.org/html/draft-irtf-icnrg-ccnxsemantics-01 (visited on 05/03/2021).

[12] *NDN Frequently Asked Questions (FAQ)*. en-US. URL: https://named-data.net/project/faq/ (visited on 05/05/2021).

[13] *POINT*. URL: cordis.europa.eu/project/id/643990 (visited on 04/29/2021).

[14] J. Postel. *hjp: doc: RFC 0791: Internet Protocol*. Sept. 1981. URL: https://www.hjp.at/doc/rfc/rfc791.html (visited on 05/10/2021).

[15] *RIFE*. en-US. URL: https://www.rife-project.eu/ (visited on 04/29/2021).

[16] *ShareTechnote*. URL: http://sharetechnote.com/html/Handbook_IP_Network_Confidentiality_Integrity.html (visited on 05/22/2021).

[17] Reza Tourani et al. "Security, Privacy, and Access Control in Information-Centric Networking: A Survey". en. In: (), p. 36.

[18] Dirk Trossen and George Parisis. "Designing and realizing an information-centric internet". In: *IEEE Communications Magazine* 50.7 (July 2012), pp. 60–67. ISSN: 1558-1896. DOI: 10.1109/MCOM.2012.6231280.

[19] Dirk Trossen et al. "IP over ICN - The better IP?" In: *2015 European Conference on Networks and Communications (EuCNC)*. June 2015, pp. 413–417. DOI: 10.1109/EuCNC.2015.7194109.

[20] Dirk Trossen et al. "PSIRP Publish-Subscribe Internet Routing Paradigm". In: (Sept. 2009). URL: http://www.psirp.org/files/Deliverables/FP7-INFSO-ICT-216173-PSIRP-D2_4_ArchitectureUpdateAndSecurityAnalysis.pdf.

[21] Patrick Truong et al. "Architecture of the DOCTOR Virtualized Node". In: (Dec. 2015). URL: http://www.doctor-project.org/outcome/deliverable/DOCTOR-D1.2.pdf.

[22] Patrick Truong et al. "DOCTOR Management and Orchestration: Security evaluation and Performance enhancement via micro-services". In: (Sept. 2018). URL: http://www.doctor-project.org/outcome/deliverable/DOCTOR-D3.2.pdf.

[23] Greg White and Greg Rutz. "Content Delivery with Content-Centric Networking". In: (Feb. 2016). URL: https://www.cablelabs.com/wp-content/uploads/2016/02/Content-Delivery-with-Content-Centric-Networking-Feb-2016.pdf.

[24] L. Zhang et al. "Named data networking". In: *CCRV* (2014). DOI: 10.1145/2656877.2656887.