Maintaining Efficiency and Privacy in Mobile Networks through Information-Centric Networking

Dirk Kutscher *, Giovanna Carofiglio †, Luca Muscariello ‡, Paul Polakos †
*NEC †Cisco Systems ‡Orange
dirk.kutscher@neclab.eu, gcarofig@cisco.com, luca.muscariello@orange.com, ppolakos@cisco.com

***Abstract*—In this paper we present a solution to overcome the impasse of deploying confidentiality at the cost of breaking most of current network traffic engineering in mobile networks. Our proposition is based on Information-Centric Networking (ICN) which is a data-centric network architecture that gracefully incorporates security and traffic optimization.**

I. Introduction

The dramatic growth of mobile communication services has been driven by three main factors: the availability of smart end devices, the success of innovative web services and the investments made to support traffic growth into the radio mobile data networks. Privacy has always been a major concern in radio mobile networks and the RAN has employed encryption by default from the second generation standards. LTE commonly employs IPSEC in the backhaul network between the eNB and the PDN gateway. The importance of guaranteeing privacy on all Internet services goes beyond the specific radio mobile data network or the all IP core; it includes the communication service in all its components: over-the-air transmission, end-to-end IP, HTTP employing e.g. `Set-Cookie` and `User-Agent` headers, and the web service that might employ user trackers or store private content in clear text in the server end. The current solution guarantees authentication, integrity and confidentiality by delegation to the transport layer on an end-to-end basis using TLS, which may serve a valuable short-term purpose, but significantly impairs the long-term flexibility, reliability, and manageability of the Internet. The usage of TLS implies significant cost that has been quantified in [7], and is a by-default solution in the HTTP/2 standard [1], provoking several criticisms about the technical solutions to the privacy problem [4].

Today, roughly half of the costs of operating a mobile data networks comes from the RAN, including spectrum license fees, infrastructure deployment and management, network operations. Most of the remaining cost comes from the backhaul network which is a large infrastructure interconnecting a huge number of towers to the mobile core, which, in the end, constitutes a negligible relative cost. Mobile traffic engineering is today based on bearer management, that constitute virtual channels from the UE to the mobile core traversing the backhaul using GTP tunnels (using IPsec possibly). This means that bandwidth optimization in mobile networks is today rather sub-optimal as the backhaul is supposed to ideally connect transparently the RAN to the core, whereas congestion management and overload control cannot be implemented within it. A number of attempts have been made in the last five years to implement congestion management by installing middle boxes in the tower and in the PDN gateway to realize traffic offloading by caching data in the tower. Such techniques have been swept away by the growth of TLS traffic from 2013.

In Europe, but increasingly also in the US, the continued ARPU decrease combined with rising investment demands is placing operator markets on an unsustainable trajectory inevitably leading to market consolidations and reduced investments. Effective resource management would thus be more important than ever. Approximately 50% of today's HTTP traffic is carried through TLS connections. We know that web traffic is highly cacheable as shown in [3] and also [12] [9]. Despite the usage of TLS we know that a significant portion of encrypted web traffic is delivered from local appliances, like Google caches, which gives an estimate of the opportunities of caching data close to the user. The need for innovative solutions that can provide network resource optimization as well as confidentiality is driven by economic incentives that the industry urges to deploy with standardized and inter-operable solutions ([8]).

There are different approaches to implement web traffic encryption, integrity and peer authentication. HTTPS is a convenient solution as confidentiality is provided by a transport layer that guarantees HTTP to remain unchanged. Other proposed solutions like S-HTTP, that unsuccessfully competed with SSL and HTTPS [10], embedded security features within HTTP without relying on a secured transport middle layer but instead by signing content hashes. The idea has also been updated more recently in [11] to modern HTTP (e.g. supporting the `Transfer-Encoded` header set to `Chunked`).

Content based security instead of connection based is the foundation of the Information-Centric Networking (ICN) architecture. In ICN, we provide a network service that directly implements the desired information-access abstraction. The network forwards requests for named data and corresponding responses containing the data. The name can be cryptographically bound to the data for ascertaining authenticity. This enables the network to replicate data objects in arbitrary locations, thus enabling ubiquitous caching. Object data can also be encrypted for user privacy, leaving other network-relevant information such as the name intact – thus maintaining options for traffic management, policing etc. The performance gains of having ICN in the mobile backhaul have been evaluated experimentally in [2]. ICN incorporates these ideas into a novel network layer providing all of the mentioned objectives without using mad-in-the-middle like solutions [5].

II. Information-Centric Networking

ICN makes use of data-centric security instead of connection security. We summarize the model here and refer to [6][13] for a detailed description of the model. ICN secures data itself by requiring producers to cryptographically sign every data packet: the signature constitutes the integrity meta-data. The data is uniquely identified by a name that is bound to the data via the signature. The producer's public key to implement signature verification can be obtained by using the `KeyLocator` field which can be the name of the data containing the key of the producer ([13]). Authentication is implemented via the producer's key that makes use of a trust model, e.g. PKI, Web-of-Trust that can be extended using key chaining to delegate trust to different sub-namespaces (for hierarchical naming). Confidentiality is obtained by encryption of the data payload using the producer's key. Notice that authenticity, integrity and confidentiality are independent features.

Once data is published by the producer it can be stored in any location without affecting the security properties of the data which are location independent. Inter-networking of encrypted data is included by design in ICN and in-network caching is always possible with or without confidentiality. Authenticity might not be necessary in many cases so the authentication of the identity of the producer is optional. It is not mandatory either to verify the integrity of the data by verification of the signature. It is important to remark that ICN disantangles authenticity, privacy and integrity so that they can be handled in different ways and without the interaction of end-hosts.

TLS provides web security by encrypting a layer 4 connection between two hosts. Authenticity is provided by the web of trust (certification authorities and a public key infrastructure) to authenticate the web server and symmetric cypher on the two end points based on a negotiated key. In presence of TLS many networking operations become unfeasible: filtering, caching, acceleration, trans-coding.

ICN takes a radically different approach to guarantee confidentiality, authenticity and integrity by embedding them into a redefined network layer. Indeed, ICN builds on the abstraction of data requested, accessed, cached and forwarded by name: the network forwards requests coming from the consumer for named data and routes back data packets on the identical reverse path (symmetric routing).

Users retrieve Data using a pull flow control protocol based on subsequent packet queries, triggering Data packets delivery. Name-based routing and forwarding guarantee that queries are properly routed towards a repository, where a permanent copy of the content is stored, following one or multiple paths. Network nodes maintain three major data structures: Content Store (CS), Pending Interest Table (PIT), and Forwarding Information Base (FIB). The CS caches Data packets received, which can be potentially useful to satisfy future Interest packets. The PIT stores Interests that have been forwarded and waiting for matching Data packets to return. The FIB is similar to IP routing table and is maintained by a name-based routing protocol. A strategy module defines the policy for output interface(s) selection at each FIB entry. For each arriving data packet, a router finds the entry in the PIT that matches the data name and forwards the data to all downstream interfaces listed in the PIT entry. It then removes that PIT entry, and caches the Data in the CS. Indeed, Data may come from the repository, or from any intermediate cache along the path with a temporary copy of the Data packet. Packets of the same content can therefore be retrieved in a multi-path fashion.

The ICN communication model allows network nodes between a web server and a web client to operate as forwarding *and* storage functions to implement various inter-networking functionalities like caching or load balancing without relaxing any security feature. As a fully fledged data-centric network architecture, ICN incorporates mobility, storage, security and multi-point communications by design.

III. Conclusions

We believe that ICN has the potential to overcome the impasse brought by the massive deployment of TLS to guarantee transport level confidentiality. ICN provides inter-networking of named-data embedding data-centric security features like authenticity, integrity and confidentiality that are location independent. ICN hence enables all relevant in-network optimizations that are of fundamental importance for mobile traffic engineering in today and next generation RAN, backhaul and core networks: caching, traffic management and bandwidth optimization.

References

[1] M. Belshe, R. Peon, and M. Thomson. Hypertext Transfer Protocol Version 2 (HTTP/2). RFC 7540, Aug. 2015.

[2] G. Carofiglio, M. Gallo, L. Muscariello, and D. Perino. Scalable mobile backhauling via information-centric networking. In *Proc. of IEEE LANMAN*, 2015.

[3] C. Imbrenda, L. Muscariello, and D. Rossi. Analyzing Cacheable Traffic in ISP Access Networks for Micro CDN Applications via Content-centric Networking. In *ACM SIGCOMM ICN*, 2014.

[4] P.-H. Kamp. HTTP/2.0 - The IETF is Phoning It In. *ACM Queue*, 13(1):10:10–10:12, Dec. 2014.

[5] J. Mattsson, R. Skog, H. Spaak, G. Gus, D. Druta, and M. Hafeez. Explicit trusted proxy in http/2.0. Draft, Internet Engineering Task Force, Aug. 2014.

[6] M. Mosko and I. Solis. CCNx Semantics. Internet Draft. Status Experimental. https://www.ietf.org/id/draft-irtf-icnrg-ccnxsemantics-00.txt, 2015.

[7] D. Naylor, A. Finamore, I. Leontiadis, Y. Grunenberger, M. Mellia, M. Munafò, K. Papagiannaki, and P. Steenkiste. The Cost of the "S" in HTTPS. In *Proc. of ACM CoNEXT*, 2014.

[8] A. Pai. Statement of FCC Commissioner Ajit Pai On Netflix's conduct with respect to open video standards, January 16 2015.

[9] B. Ramanan, L. Drabeck, M. Haner, N. Nithi, T. Klein, and C. Sawkar. Cacheability analysis of HTTP traffic in an operational LTE network. In *In Proc. of WTS*, 2013.

[10] E. Rescorla and A. Schiffman. The secure HyperText transfer protocol (S-HTTP). RFC 2660, Internet Engineering Task Force, Aug. 1999.

[11] K. Singh, H. J. Wang, A. Moshchuk, C. Jackson, and W. Lee. Practical end-to-end web content integrity. In *Proc. ACM WWW*, pages 659–668, 2012.

[12] S. Woo, E. Jeong, S. Park, J. Lee, S. Ihm, and K. Park. Comparison of caching strategies in modern cellular backhaul networks. In *Proc. of ACM MobiSys*, 2013.

[13] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang. Named data networking. *SIGCOMM Comput. Commun. Rev.*, 44(3):66–73, July 2014.