An Orchestrated NDN Virtual Infrastructure transporting Web Traffic: Design, Implementation and First Experiments with Real End-Users

Guillaume Doyen, Thibault Cholez, Wissam Mallouli, Bertrand Mathieu, Hoang-Long Mai, Xavier Marchal, Daishi Kondo, Messaoud Aouadj, Alain Ploix, Edgardo Montes-de-Oca, Olivier Festor.

Abstract—After more than one decade of research efforts, ICN technologies seem mature enough to move from a design and implementation phase to early deployment trials. However, if one wants an ICN stack to be operated in a production context, some major locks, related to (1) the lack of adequate deployment infrastructure, (2) the migration of relevant services and (3) the capability to accurately monitor an overall ICN domain, must be addressed. In this paper we present a feedback experience of the deployment of a Named Data Networking island. The latter considers HTTP over NDN as a primary service deployed through dedicated gateways and NFV as a substrate for deploying and orchestrating ICN components. From the performance assessment of individual components up to the opening of the testbed to end-users in university campus, we propose an analysis that can guide further research efforts in the ICN area.

Index Terms—Named Data Networking, Service Migration, Network Function Virtualization, *In vivo* experiments

I. INTRODUCTION

After more than one decade of research and development, the Information Centric Networking (ICN) [1] paradigm has now reached a level of maturity which makes it a promising candidate to replace or complete standard IP stacks. Among the numerous insights that attest of this maturity, one can observe that (1) several ICN architectures proposed to date have stable and fully-operational implementations (e.g. CCNX, NFD, Pursuit's Blackadder, Netinf, (2) some worldwide testbeds allow the deployment of real services for research and experimentation purposes (e.g. the NDN testbed¹ or the ICN testbed federation²), and finally (3) the ICN Research Group (ICNRG), a dedicated working-group at the IRTF, actively works on the standardization of several key points of this paradigm such as deployment considerations [2].

Guillaume Doyen and Alain Ploix are with trans-disciplinary cyber-security team, ICD, UMR 6281 CNRS, Troyes University of Technology, Troyes, France.

Hoang-Long Mai, Wissam Mallouli and Edgardo Montes-De-Oca are with Montimage Research Labs, 39 rue Bobillot, 75013, Paris, France.

Thibault Cholez, Xavier Marchal, Daishi Kondo and Olivier Festor are with LORIA-CNRS, 54506 Vandoeuvre-les-Nancy, France.

Bertrand Mathieu is with Orange Labs, Lannion, France.

This work is partially co-funded by (1) the French National Research Agency (ANR), DOCTOR project, <ANR-14-CE28-0001>, started in 01/12/2014 and supported by the French Systematic cluster and (2) the CRCA Excellence grant <A2101-03>, and the CRCA-FEDER CyberSec Platform <201304601>.

¹<https://named-data.net/ndn-testbed/>, Accessed on 04/02/19

²<http://www.icn2020.org/2018/01/15/testbed-federation/>, Accessed on 04/02/19

Consequently, ICN is now moving from a pure research stage toward early trials of deployments, which exhibit novel challenges this paradigm has to face. The first one is related to deployment consideration, requiring novel infrastructure means able to host all ICN network functions. The second relies in the identification and migration of relevant services that can benefit from a deployment over an ICN. Finally, novel monitoring and security facilities have to be designed to enable stakeholders to deploy ICN domains in a safe and manageable way.

Focusing on the Named Data Networking (NDN) solution [3], the most federating ICN architecture to date, and after having summarized our contributions about security in [4], we present in this paper our achievements about (1) service migration with the presentation of a HTTP/NDN gateway able to carry web traffic [5] (2) infrastructure means which leverages Network Function Virtualization (NFV) to propose content-oriented orchestration [6]. Original contributions depicted in this paper especially focus on NDN caching of web content and first results of experiments performed with real users.

The paper is organized as follows. In Section II, we present the related work on ICN deployment and service migration. In Section III, we present two major contributions which are an HTTP/NDN gateway enabling web traffic to cross NDN islands and a content-oriented MANO which enables the automated deployment and dynamic enforcement of NDN policies. In Section IV, we depict the different evaluation results from unitary tests of the components and a measurement campaign involving real end-users. Finally, Section V draws some lessons from this first deployment experience that can help guiding any further contributions in this area.

II. RELATED WORK

A. SDN and NFV for ICN Deployment

It is commonly admitted that NDN will not replace IP in a one-shot phase, but that the deployment will rather be progressive and Software Defined Networking (SDN) and NFV are two key technologies that allow it.

Many research efforts argue for the cohabitation of IP and ICN on a common Layer-2 network by leveraging the network programmability offered by SDN. In [7], the authors propose to use a dedicated UDP or TCP port to identify ICN protocol and to extend the SDN controller with an ICN module. Salsano et al. [8] propose a framework for deploying

ICN functionalities over SDN using the IP option header as a name field for ICN. Meanwhile, the authors of [9] propose and implement an ICN module in the SDN controller to process the forwarding path computation for NDN flows, separately from IP flows. Nguyen et al. [10] implement an intermediate layer between a CCN node and an OpenFlow switch called *Wrapper*, and the combination of three elements acts as an ICN router.

NFV rather argue for the separation of IP and ICN protocols by leveraging the isolation property of virtualization. It also enables the deployment of ICN without requiring any change in the current network infrastructure thus acting as key enabler. Sardara et al. [11] follow this direction by proposing vICN (Virtualized ICN). The authors provide a flexible unified framework for ICN, which includes several functions such as monitoring. The H2020 FLAME project aims at fully integrating ICN into an overall media function platform using the concepts of Service Function Chaining (SFC). The 5G next generation system architecture defined by 3GPP leverages NFV and SDN technologies to provide the flexibility to deploy ICN-as-a-Slice. Finally, in [12], the authors benefit from NFV to provide contextualized edge services relying on ICN protocol stacks.

To date, several solutions which aim at providing practical deployment solutions of ICN with SDN and NFV, have been proposed. However, to the best of our knowledge, our architecture is the first approach which pushes the content-oriented paradigm of ICN up to high-level orchestration templates while keeping interoperability with existing standards for network virtualization such as the ETSI MANO reference architecture, TOSCA as high-level specification language, and Docker as a virtualization substrate.

B. Service Migration

The prime protocol to deliver content nowadays is HTTP, and it is natural that a few initiatives already tried to make the bridge between HTTP and the NDN world. A first approach was proposed by Wang et al [13] that translates HTTP requests to CCN *Interests* packets but it exhibits weaknesses such as the inability to communicate with a native CCN producer/consumer and the misuse of the "metadata" field in the CCN *Interest* packet to carry HTTP request's header. The authors of [14] explain how CDN could benefit from ICN and identify CCN/HTTP translation and CCN/IP tunneling as key technologies of their architecture, but they only provide a high level description of the gateway. Another more generic approach, also relying on a gateway, was proposed by Moiseenko et al [15]. They succeeded to carry TCP over NDN, and consequently all the upper protocols that use TCP, which includes HTTP. While their work is a significant step toward the adoption and deployment of NDN, their generic solution also misses an important incentive because it does not use one of the main features of NDN, the cache, when carrying web contents because it is not possible to have an efficient content-level caching when only TCP-level information is considered.

To conclude, to the best of our knowledge, our work is the sole to address the most relevant service migration case

while preserving all the features of both the HTTP protocol and the NDN one. It is also the sole to have been extensively evaluated in a realistic deployment situation with end-users accessing real web sites.

III. A VIRTUALIZED NDN INFRASTRUCTURE FOR WEB TRAFFIC

A. Leveraging NFV as a Substrate for NDN Deployment

The overall architecture of our content-oriented MANO is illustrated in Figure 1. It strictly follows the ETSI reference architecture specification³. It leverages Docker as the core technology for the *Network Functions Virtualization Infrastructure (NFVI)* and VXLAN as an encapsulation strategy for the NDN data-plane traffic, thus making the *NFVI* agnostic to the carried traffic nature. Consequently, the *Virtual Infrastructure Manager (VIM)*, does not need to be extended to support NDN traffic too, and we have selected Docker Swarm as a ready-to-use technology (arrow 9). As such, the methodology we followed has consisted in solely extending or redesigning the MANO components which need NDN awareness, without diverting them from their initial purpose. These are the *VNF Manager (VNFM)* and the *NFV Orchestrator (NFVO)*.

To provide network administrators with a high level and intuitive way to specify virtualized NDN services, we have designed a novel TOSCA profile for NDN which also integrates novel policies to dynamically react against NDN specific performance and security incidents. The specifications for *Virtual Deployment Units*, *Virtual Links* and *Connection Points* have been kept unchanged since they only relate to the infrastructure layer while those for *VNF*, *Forwarding Path* and *policies* have been extended to encompass NDN features. For instance, the *VNF* specification includes configuration parameters that represent the set of NDN prefixes to be announced as well as the status of a signature verification module in an NDN router, while the *Forwarding Path* specification captures the list of VNFs that a particular set of NDN packets will follow defined on either the forwarder or the NDN prefix. Finally, our NDN TOSCA extension enables the specification policies modelled with Event-Condition-Action (ECA) rules that apply dynamically during service runtime for dynamic re-configuration operations (e.g. enforcement of signature verification to be applied on *Data* packets) and scale-out.

We have designed and implemented a dedicated NFVO which includes two main blocks: a TOSCA parser and an Orchestration engine. Given our TOSCA profile for NDN (arrow 1), the TOSCA parser reads TOSCA templates (arrow 2) and creates an in-memory graph of TOSCA nodes and their relationships (arrow 3). The graph is then passed to the NFVO Core component (arrow 4) which delegates any NDN-specific operation to the NDN Engine (arrow 5) which extracts the forwarding information defined in the *Forwarding Path* and it translates it into FIB configuration for each VNF which includes the NDN prefix, the address of the next hop and the port number of the NDN process (e.g. NFD, NDN firewall

³<https://www.etsi.org/technologies-clusters/technologies/nfv>, Accessed on 04/02/19

or signature verification module). In case of a deployment automation, the NFVO Core orders the VIM to deploy a management network (arrow 7) to ensure the communication between VNFs and the VNF Manager, followed by the deployment of VNF Manager as a container connecting to this network (arrow 6).

B. Migrating HTTP through Dedicated Gateways

An NDN island using our HTTP over NDN architecture must instantiate two kinds of gateways: (1) an ingress gateway (iGW), that converts HTTP user requests into NDN messages and returning NDN messages into HTTP responses; and (2)

The gateways follow a naming pattern based on a naming proposition to convert URL to ICN names⁴. Since an NDN *Interest* packet cannot carry data while an HTTP request’s header does, iGW (or a native NDN web-client) must exchange different messages in 3 steps, defined in Table I, to retrieve the web content. First, iGW sends an *Interest* whose name components contain, as illustrated in Table I.a:

- This *Interest* packet is sent in the NDN network to ask someone to handle the request (Figure 2 steps 1-2). Consequently, eGW (or a native NDN web-server) knows upon reception that an HTTP request must be satisfied, but also the network name to contact the client to get the request details. Please note that the SHA1 is used to be sure that we respect the parameters of the HTTP request to match users' properties. However, choosing carefully a subset of fields to be considered

⁴<http://www.icn-names.net/>, Accessed on 04/02/19

Table I
NAMING PATTERN OF THE MAPPING PROTOCOL

a.	/http/reverse_splitted_domain_name/path/sender_route/sha1
b.	/sender_route/sha1(/segment)
c.	/http/reverse_splitted_domain_name/path/sha1(/version/segment)

to compute the hash can vastly improve NDN caching while giving consistent results to users.

Indeed, our preliminary experiments showed that users cannot benefit from the cached responses initiated by another web browser. This is mainly due to the relative uniqueness of HTTP requests sent by a browser. As such, HTTP responses' packets can only come from the NDN cache if users ask the same content in the very same way regarding all options in HTTP requests' header. By default, a user reloading a page gets 75% of packets from the NDN cache, while a subsequent user with the same browser only gets 38%, and another subsequent user with a different browser 0%. This can endanger the ability of HTTP traffic to take advantage of the NDN island. A solution is to ignore or replace some common but "useless" fields of the header when computing the hash in order to maximize the re-usability of concurrent requests to a given content. Our tests show that the cachability is thus vastly improved without affecting the accuracy of delivered content by modifying *accept-encoding* and ignoring the *user-agent*, *accept-language*, *accept* and *cookie* fields for static contents. With this optimization that preserves all significant HTTP features, a second user with the same browser now achieves 78% of cache hit and a subsequent user with a different browser, 61%.

In the second step, the eGW extracts information from the first *Interest* sent by iGW, more precisely the two last components: the sender route and the hash, in order to retrieve the full HTTP request (Figure 2 steps 3-4 and Table I.b). The sender route is coded as a single binary field name component to be extracted easily. This preliminary exchange including the sender route is mandatory because *Interest* packets cannot transport *Data* in NDN, so that the HTTP request must be retrieved specifically. Once the full HTTP request is received by the eGW, it can now ask the HTTP server in the IP domain for the actual web content. After the reception of the HTTP response from a regular HTTP server, the eGW prepares the NDN *Data* packets. Following NDN principle, it is up to the NDN client to send *Interest* packets to retrieve each chunk of the HTTP response (Figure 2 steps 5-6 and Table I.c). Please notice that the six steps of the mapping protocol are actually done in 2 RTT, because some steps are done in parallel (3-4 and 5-6). In the case of a native NDN web server, the mapping protocol is the same but content is directly sent by the server without issuing traffic on the IP network.

IV. EVALUATION

We followed a specific experimental approach to evaluate each of our contributions: unitary tests for the HTTP/NDN gateway, test scenarios with emulated users for the content-oriented orchestration infrastructure and finally *in vivo* tests with real end-users.

Figure 2. Sequence diagram of the communication protocol between gateways

Figure 3. Distribution of the web sites based on percentage retrieval

A. Evaluation of Service Gateways

We evaluated our gateway both in terms of performance and reliability and in terms of interest and efficiency for end-users. For the performance tests, we used one scrapper we internally developed. For the interest and quality for end-users, we used Webview⁵, a tool developed by Orange, allowing to perform automation of tests for the web browsing service and measuring quality metrics, such as the Page Load Time (PLT), as defined by the W3C consortium.

For the performance and accuracy of our gateway, we evaluated the success of retrieving the requested objects. We plot in Figure 3 the frequency distribution of the page content, given by the number of HTTP objects retrieved from the set of all objects that the infrastructure successfully requests and retrieves. The result (plot in semi-log scale) shows that most of the top-1000 HTTP websites can be retrieved entirely, those with bad results being mainly remote websites (Chinese, Korean and Russian).

We evaluated the additional delay resulting from the usage of the gateways. This experiment shows that the additional delay resulting by the usage of the gateways is nearly constant and in our case equal to 29+/-3ms. We can get a constant time because eGW does not wait the completion of the HTTP

⁵<https://webview.orange.com/>, Accessed on 04/02/19

responses and starts to generate *Data* packets as soon as possible.

B. Evaluation of Content Oriented Orchestration Components

The first experiments we have performed to assess the performance of our content-oriented MANO framework have consisted in evaluating the capability of our solution to automatically deploy an NDN virtual topology. Different topologies exhibiting various numbers of VNFs and different connectivity degrees have been considered (e.g. star, ring and mesh). The collected results have shown that (1) our content-oriented MANO is able to spawn and configure VNFs in a reasonable time (from 10 to 30 seconds on average per VNF), (2) the NDN configuration enforcement is reasonably time consuming as compared to standard container spawning and (3) the total deployment time grows linearly according to the number of VNFs. The bottleneck we identified for the deployment automation resides in the connection degree of VNFs which induces orchestration operations of both infrastructure layer endpoints as well as the convergence of forwarding routes in the NDN virtual network, thus leading to an potential exponential grow of the deployment time in case of a full-mesh topology.

Then, in order to evaluate the dynamic part of our orchestration framework, we considered the case of a Content Poisoning Attack (CPA) which is one of the current threats in NDN. Our solution could be extended to other security threats (e.g. Interest Flooding Attack, cache probing, route poisoning) but such a study is left for future work. CPA detection is achieved through a dedicated detector presented in [4] and from an orchestration perspective, we consider three mitigation policies that bring a virtual NDN network under attack back to normal. First, on a CPA detection by a security probe, our orchestrator dynamically enforces the signature verification module, which checks the integrity of each *Data* packet, of NDN routers located at the edge of the NDN virtual network. Then, for each corrupted prefix, the orchestrator dynamically reconfigures the black list of the upstream virtual NDN firewall to prevent such traffic from entering the network. Finally, to deal with any potential overload of NDN routers enforcing the signature verification which may induce a too high computing cost, the orchestrator enforces a scale-out policy which dynamically spawns and configure replicas of NDN routers.

Figure 4 exhibits time series of *Good Data* (blue lines), *Bad Data* (red lines) and *Lost Data* (black lines), expressed in terms of number of packets per second, received by a good user under a CPA. The green dotted line shows the time at which the network detects the attack and triggers the firewall and signature verification policies to mitigate the attack. We can observe that, after this moment, the number of *Bad Data* abruptly drops to zero. However, the number of *Good Data* also decreases slightly, and the number of *Lost Data* increases remarkably. The reason lies in the overload the routers performing the signature verification suffer from, which prevents them from reliably achieving their basic forwarding operations. Then, the two black dotted lines show the moment when the scale-out configuration operations start and

Figure 4. Time series of the dynamic orchestration of CPA mitigation ({red, blue, black} lines respectively stand for {good, bad, lost} data packets)

terminate. After that time, the network returns to a normal state: *Good Data* gets back to its initial rate and *Lost Data* drops to a negligible value. We can also observe that during the period of scale-out, there is no downtime of the network. Overall, these results show the effectiveness of the content-oriented orchestration in terms of security and performance.

C. Real User Experiments

To assess the performance of our overall deployment solution for NDN, we conducted a real-scale experiment involving real users. During five days, the HTTP traffic issued by a dozen of volunteers students from the Troyes University of Technology campus has been routed to a basic version of our testbed, integrating our gateways and a single NDN router but without dynamic reconfiguration.

In total, 38 reports of users' perception of the HTTP service operated over our ICN island were provided. They were questioned on (1) any additional delay they perceived as compared to their everyday-life navigation, (2) the incomplete retrieval of web contents in the pages they accessed and (3) their will to adopt this technology in a real life.

The collected results are synthesized in Table II. They are globally promising since for all questions, results are satisfying. To the question related to their perceived additional delay, more than half of the users did not notice any change while two almost identical subsets notice a degradation (23,68%) or improvement (21.06%) making us conclude that the island does not significantly alter the perceived Quality of Experience while it acts as an additional domain to cross before reaching the Internet. Similarly, to the question about the capability of our solution to retrieve any web content, more than half of the users have been able to retrieve all web objects they accessed to, and only a very small portion notice an important lack in the content retrieval (5,26%) due to a bug occurring with non-occidental encoding. Finally, to the last question related to the ICN adoption in the future, a large part of these testers

Table II
QUALITY OF EXPERIENCE PERCEPTION OF REAL USERS OF OUR
HTTP/NDN ISLAND

Proposal	% of answers
<i>Delay to retrieve web content</i>	
Really worse than usual	5,26%
Slightly worse than usual	18,42%
No difference	55,26%
Slightly better than usual	10,53%
Much better than usual	10,53%
<i>Missing elements in the retrieved content</i>	
Many elements are missing	5,26%
Some elements are missing	36,84%
All items are fully loaded	57,89%
<i>Adoption for a real life usage</i>	
Yes	89,47%
No	10,53%

indicated they are in favor of adopting this technology given their quality of experience during this experiment.

These results make us conclude that an ICN island is a promising candidate for a telco to provide web content to its users, but the diversity and richness of web content must be well caught by ICN components and especially gateways to enable a fully reliable content retrieval.

V. LESSONS LEARNED AND FUTURE RESEARCH DIRECTIONS

In this paper, we presented two contributions which aim at pushing forward the deployment of ICN, and NDN in particular. The first one is a content-oriented orchestrator for NDN networks which brings the content paradigm up to service specification. By extending the TOSCA standard profile, the latter allows to dynamically manage the lifecycle of an NDN network deployed in a virtualized infrastructure and to dynamically react to events or anomalies by reconfiguring a virtualized NDN network. As a use-case of service that can benefit from a migration to NDN, we designed a protocol mapping from HTTP to NDN and implemented an operational gateway which allows end-users using current web browsers to reach public web servers, via an NDN network. Our solution is compliant with current standards of network virtualization, transparent for end-users, adaptive to allow the communication between the NDN and IP worlds, efficient and reliable. The components have been exhaustively evaluated in the context of a security scenario, emulated web traffic and also a real deployment campaign where real end-users from a university campus, via our NDN testbed, accessed public websites using their own browsers. In this context, we also discovered that caching is not as efficient as we could expect with real end-users, due to the variety of accessed web sites, the various web browser configurations and the personalization of content. As a guideline for future work, novel NDN caching features moving away from a generic usage should be refined for the particular web browsing service to improve its general performance and consequently the quality of experience of end-users.

Future direction of work in our research area is as follows. Our short-term perspective consists in analyzing the set of metrics we have collected during the experiments with real

end-users. Our purpose is to provide recommendations and models for subsequent deployments of NDN. Then, to promote our gateway, we will join the federated NDN testbed, and hope to cause an increase of its inner-traffic once the web is made available and with the iGW constituting a cheap entry point for users. Finally, the genericity of our solution which is currently restricted to NDN is an open-question. One can envisage an adaptation of our work for similar ICN solutions such as CCN, which would ease the mapping with HTTP because CCN Interest packets can directly transport the HTTP request, but the design and implementation of content-oriented MANO components is still dependent on the actual ICN technology they manage and would require some efforts.

REFERENCES

- [1] B. Ahlgren et al., "A survey of information-centric networking," *IEEE Communications Magazine*, vol. 50, no. 7, 2012, pp. 26–36.
- [2] A. Rahman et al., "Deployment considerations for information-centric networking," *ICNRG draft*, available at <https://tools.ietf.org/html/draft-irtf-icnrg-deployment-guidelines-04>, 2018.
- [3] L. Zhang et al., "Named Data Networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, 2014, pp. 66–73.
- [4] T. N. Nguyen et al., "A security monitoring plane for Named Data Networking deployment," *IEEE Communications Magazine – Feature Topic on Information-Centric Networking Security*, vol. 56, no. 11, 2018, pp. 88–94.
- [5] X. Marchal et al., "Leveraging NFV for the deployment of NDN: application to HTTP traffic transport," in *2018 Proceedings of IEEE/IFIP NOMS*, 2018, pp. 1–5.
- [6] H. L. Mai et al., "Toward content-oriented orchestration: SDN and NFV as enabling technologies for NDN," To appear in: *2019 Proceedings of IEEE/IFIP IM*, 2019, pp. 1–5.
- [7] M. Vahlenkamp et al., "Enabling information centric networking in IP networks using SDN," in *2013 Proceedings of IEEE SDN4FNS*, 2013, pp. 1–6.
- [8] S. Salsano et al., "Information centric networking over SDN and openflow: Architectural aspects and experiments on the OFELIA testbed," *Computer Networks*, vol. 57, no. 16, 2013, pp. 3207–3221.
- [9] N. L. M. van Adrichem and al., "Ndnflow: Software-defined Named Data Networking," in *2015 Proceedings of IEEE NetSoft*, 2015, pp. 1–5.
- [10] X. N. Nguyen et al., "Efficient caching in Content-Centric Networks using Openflow," in *2013 Proceedings of IEEE INFOCOM*, 2013, pp. 1–2.
- [11] M. Sardara et al., "Virtualized ICN (vICN): Towards a unified network virtualization framework for ICN experimentation," in *2017 Proceedings of ACM ICN*, 2017, pp. 109–115.
- [12] P. TalebiFard et al., "An Information Centric Networking approach towards contextualized edge service," in *2015 Proceedings of IEEE CCNC*, 2015, pp. 250–255.
- [13] S. Wang et al., "On adapting HTTP protocol to Content Centric Networking," in *2012 Proceedings of CFI*, ACM, 2012, pp. 1–6.
- [14] G. White and al/ "Content Delivery With Content Centric Networking," Cable Television Laboratories, Tech. Rep., 2016, pp. 1–26.
- [15] I. Moiseenko et al., "TCP/ICN: Carrying TCP over Content Centric and Named Data Networks," in *2016 Proceedings of ACM ICN*, 2016, pp. 112–121.

BIOGRAPHIES

Guillaume Doyen has been an associate professor at UTT since 2006. His current research focuses on the design of autonomous management solutions for content distribution and virtualized infrastructures.

Thibault Cholez is an Associate Professor at Université de Lorraine where he works in the Inria RESIST research team on Data Network Monitoring and Analytics.

Wissam Mallouli is currently a research and development project manager at Montimage, France. His topics of interest cover formal methods for monitoring of functional, performance and security aspects of networks and applications.

Bertrand Mathieu joined Orange Labs in 1994. Since 1999, he is working on distributed computing, programmable networks, overlay networks, QoS and QoE and Information-Centric Networking.

Hoang-Long Mai is a Ph.D. student in a Industrial Convention of Formation by Research (CIFRE) contract between Montimage, UTT, and INRIA Lorraine. His Ph.D. topic focuses on the autonomous monitoring and control of virtualized network functions.

Xavier Marchal is a PhD student at LORIA/INRIA, he works on ways to efficiently deploy new networks like NDN thanks to new paradigms like NFV and SDN but also thinks about ways for these networks to cohabit.

Daishi Kondo received his Ph.D. degree in computer science from the University of Lorraine in 2018. His research interests include Information Centric Networking, network security, privacy, and Peer-to-Peer networking.

Messaoud Aouadj is a postdoc researcher at UTT. His research interests focus on works on network virtualization and software-defined networking.

Alain Ploix has been an associate professor at UTT since 2004. He has been at the head of the network and telecommunication department from 2010 to 2018.

Edgardo Montes De Oca founded Montimage, a research oriented SME in 2004. His main interests include monitoring the security and performance of 4G/5G networks.

Olivier Festor is a professor in computer science at the University of Lorraine and director of TELECOM Nancy, the Graduate Engineering School in Computer Science.