NEXT
GENERATION
INTERNET

Grant Agreement No.: 871582

Call: H2020-ICT-2018-2020

Topic: ICT-31-2018-2019

Type of action: RIA

NGI ATLANTIC.EU

D3.4 EXPERIMENT CATALOGUE 1

Work package	WP3 “NGI Experiments”
Task	T3.4 “Policy and strategic impact creation”
Submission date	18/12/2020
Deliverable lead	Trust-IT
Version	1.0
Authors	Cristina Mancarella, Patricia Nugent, Sara Pittonet (Trust-IT)
Reviewers	James Clarke (WIT), Halid Hrasnica (Eurescom)

Abstract	1st Version of the Experiment Catalogue, a carousel of the 1st set of NGIatltic.eu EU-US funded projects to be provided via the NGIatlantic.eu platform
Keywords	NGI, NGIatlantic.eu funded projects, EU-US NGI cooperation

Disclaimer

This public deliverable is considered in “DRAFT” format until it is independently reviewed by project reviewers.

Acknowledgements

NGIatlantic.eu has received funding from the European Union’s Horizon 2020 research and innovation programme under the Grant Agreement no. 871582. The content of this document does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of such content.

EXECUTIVE SUMMARY

The NGIatlantic.eu Experiment Catalogue is an online carousel giving a concise, yet comprehensive set of information on each NGI.atlantic.eu funded EU-US experiment, accessible via the NGIatlantic.eu platform at the address https://ngiatlantic.eu/funded-experiments.

This document (D3.4, EXPERIMENT CATALOGUE 1) and its subsequent iterations is an accompaniment in reported format of the online catalogue (D3.3, EXPERIMENT CATALOGUE DATABASE), gathered every six months in a booklet with one final publication at Month 30.

The information contained in the document was gathered from the first contractual deliverable provided by the funded projects at the time of writing structured specifically to this purpose; in addition, information was gleaned from the applications, other documentation and dedicated interviews, which took place as part of the communication and dissemination activities carried out under WP4. This entry will continue to be enriched incrementally over the project duration as the experiments evolve in terms of progress and ultimately impact and results.

Following the introduction, this deliverable report presents the catalogue overview and structure (Section 2) followed by Section 3 presenting the Open Call topics of the first three Open Calls and the insights from the first two closed calls. Section 4 brings the reader into the real catalogue, with the list of the four experiments funded and started at the time of writing, entitled:

- Fairteam;
- Measuring Multi-Carrier Cellular Access International Roaming Performance;
- CacheCash Experiment;
- Self-Certifying Names for Named Data Networking.

The concluding Section 5 introduces the following iterations of the catalogue.

Co-funded by the Horizon 2020
Framework Programme of the European Union

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Co-funded by the Horizon 2020
Framework Programme of the European Union

ABBREVIATIONS

5G-PPP	5G-Public Private Partnership
DG CNECT	Directorate General for Communications Networks, Content & Technology
EAG	External Advisory Group
EC	European Commission
EIT	European Institute of Innovation & Technology
EMT	Experiment Monitoring Team
ICE-T	Internet Core & Edge Technologies
NGI	Next Generation Internet
NSF	National Science Foundation
R&D	Research & Development
RIA	Research Innovation Action
SME	Small and Medium-Sized Enterprise
KPI	Key Performance Indicator

Co-funded by the Horizon 2020
Framework Programme of the European Union

1. INTRODUCTION

The main objective of D3.4, and its subsequent iterations, is to gather a concise, yet comprehensive set of information on each NGIatlantic.eu funded EU-US experiment in a consolidated manner, which will be published incrementally following each Open Call, in an online catalogue accessible via the NGIatlantic.eu platform (https://ngiatlantic.eu/funded-experiments).

D3.4, followed by D3.5, D3.7 and D3.9 yet to come is, therefore, an accompaniment in reported format of the online catalogue, which will also include multimedia material. Every six months, or at the appropriate times of the open calls' projects' implementation, the information gathered will also be printed in a booklet accompanying each deliverable with one final publication at Month 30.

The information contained in the document was gathered from the first contractual deliverable provided by the funded projects at the time of writing structured specifically to this purpose; in addition, material was also derived from the applications, other documentation and dedicated interviews that took place as part of the communication and dissemination activities carried out under WP4.

This document is delivered at M12, as opposed to M6 to include information on the funded projects from Open Call #1, which went through a lengthier selection process than originally envisaged and the reluctance of the open call projects to launch earlier due to the COVID-19 pandemic impacts. However, at the time of writing, it has been possible to capture an initial insight into the results of NGIatlantic.eu's Opens Call #2, as well as the priority topics identified for Open Call #3.

1.1. Relation with other deliverables

This deliverable builds on material made available for D3.3, The Experiment Catalogue Database, whereby uniform information is gathered from each of the funded EU-US NGIatlantic.eu experiments. It is the first of a series of 4 documents to be reiterated incrementally until project conclusion as follows:

- D3.5 Experiment Catalogue 2 (M12);
- D3.7 Experiment Catalogue 3 (M18);
- D3.9 Experiment Catalogue 4 (M30).

Also relevant to this deliverable in that they will draw on the content of the above are:

- D3.6 Interim Experiment Results and Global impact of the EU-US NGI experiments period 1 (M12) - and its reiteration;
- D3.8 Final Experiment Results and Global impact of the EU-US NGI experiments period 2 (M30).

Finally, D3.4, D3.5 and D3.9 will impact the following deliverables due under WP4:

- D4.2 NGI EU-US experimental platform ecosystem status report period 1 (M14) - and its reiteration;
- D4.3 experimental platform ecosystem status report period 1 (M30);
- D4.4 Final report on communication & stakeholder's engagement and promotion of the open call (M30);
- D4.5 NGIatlantic.eu Priority Landscape Report 2030 (M30);
- D4.6 Exploitation report.

Co-funded by the Horizon 2020
Framework Programme of the European Union

1.2. Specific objectives

The main objective of the catalogue is to give **full online visibility** to the experiments funded under NGIatlantic.eu, with succinct, uniform details provided for each enhanced with multimedia materials. The initial information is gathered at kick-off for the projects in the format prescribed by the first experiment deliverable, D1, due at the end of their Month 1, and this will continue to be enriched incrementally over the project duration in terms of progress and ultimately impact and results as the experiments evolve.

The catalogue also addresses some of WP4 specific objectives, as detailed in *D4.1 Plan for communication & stakeholder's engagement and promotion of the open call*, namely:

- Showcase NGIatlantic.eu results, creating awareness of the impact of the experiments carried out with NGIatlantic.eu cascade funding;
- Ensure global recognition and visibility on how the EU and the US are impacting and influencing NGI efforts;
- Deliver exploitation activities for the experiments carried out.

Co-funded by the Horizon 2020 Framework Programme of the European Union

2. CATALOGUE OVERVIEW

D3.4 is a living document that will be updated every six months, or at the project implementation time-line of the experiments coming onstream from each of NGIatlantic.eu's Open Calls. For this reason, all of the funded projects are presented with the same structure, which will be then replicated online in the Experiment Catalogue available on the website.

2.1. Structure of the catalogue

For each experiment, the Catalogue presents the **overall participants details**, including the name of the project, the EU Coordinator name and country of origin, as well as the organisation type, and, finally, the US participants information. The project topic and the TRL Level of the solution proposed are included too.

The **experiment description** follows, providing an abstract of the experiment, as well as the implementation plan and expected impacts.

Where available, multimedia material is included too, such as interviews carried out with the project's participants.

Finally, each project presents the **future plans**, risk and contingency measures and the needs for Communication, Dissemination and Exploitation.

Co-funded by the Horizon 2020
Framework Programme of the European Union

3. NGIATLANTIC.EU FUNDED EXPERIMENTS AT A GLANCE

3.1. Topics of the research

From the outset of the Project, the topics of the NGIAtlantic.eu Open Calls were selected to reflect the emerging results from other NGI Research and Innovation Actions (RIA). Therefore, after mapping results from the previous open calls, Open Call #1 launched with three NGI topics for the experimentation of results between EU and US applicants in areas related to i) **Privacy and Trust enhancing technologies**, via the NGI Zero Trust and NGI Trust projects; ii) **Decentralised data governance**, via LEDGER, and iii) **Discovery and identification technologies** (better search and discovery) via NGI Zero DISCOVERY open call projects.

Open Call #2 maintained the topics already identified to ensure continued coverage of priority topics of the NGI initiative in a mature state with results ready for experimentation on EU – US experimental platforms, while opening a new stream focussed on the **interconnection of EU-US infrastructures for NGI experimentation**. The intention here was to place more focus on the bridging of already matured EU-US experimental platforms, bringing these to the forefront of the project and helping set up their availability for the future calls.

At the time of writing, Open Call #3 has just launched on 1 December 2020 and will run until 26 February 2021. With input and support provided by the External Advisory Group (EAG), two new, emerging topics have been introduced to the stream of NGI initiatives identified in i) to iii) above, with iv) **Strengthening Internet Trustworthiness with electronic identities**, embracing self-sovereign identity (SSI), federated and/or decentralised technologies for supporting internet-wide e-identities with various levels of identification, reputation and trust, to serve as a basis for new business models for verifying and valuating personal data. The second new topic for EU-US experimentation is v) **Greening the Internet: a Sustainable and Climate-friendly NGI,** which focuses on solutions to contribute to the fight against climate change, awareness of the environmental impact of the Internet, and promotion of technologies that help reduce energy consumption and carbon emission.

It should also be noted that all three NGIatlantic.eu Open Calls clearly emphasised in the call text and supporting documentation that applications from new EU innovators not already involved in other NGI projects with their own results to experiment on EU-US platforms were also welcome to submit applications.

3.2. Open call results 1 and 2 insights

For **Open Call #1,** a total of **21 EU – US proposals** were submitted via the [NGIatlantic.eu Trust Grant platform](). The projects applying for funding were all in the Long Term (LT) category, of six-months in duration and with an average request of just over 130,000 Euro in funding. The proposals provided **good geographical coverage both in the EU, with coordinators from twelve different Member States, as in the US, with 21 organisations from 14 different states**, as shown in the figure below.

Co-funded by the Horizon 2020 Framework Programme of the European Union

FIGURE 1 *INSIGHTS FROM EU-US PROPOSALS RECEIVED IN OPEN CALL 1*

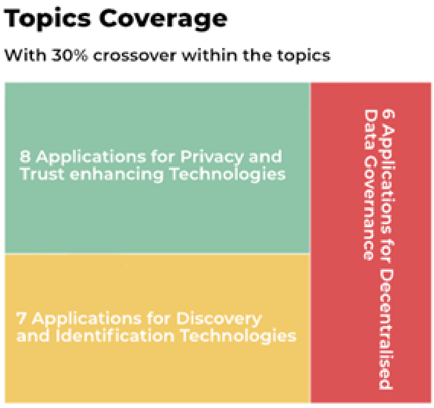

FIGURE 2 TOPICS OF COVERAGE OF OPEN CALL 1

A good balance in terms of the three NGI topics prioritised under this call was also provided with 8 applications received on Privacy and Trust Enhancing Technologies, 6 on Decentralised Data Governance and 7 on Discovery and Identification Technologies. As to be expected, an intersection of about 30% of the proposals between topics could also be observed.

As to the type of organisations applying for funds, on the EU side 15 applications were from industry (mainly SMEs), 9 from Research/Academia, and 2 from non-profit organisations. As to the US side, most of the organisations involved were from Research & Academia, 10, with 8 were from large companies & SMEs, and 3 from non-profit organisations.

At the time of writing, the improved evaluation process adopted for Open Call #2 is just coming to a close.

In terms of applications submitted this time a total of **23 EU-US proposals** was reached, with 19 falling within the Long-Term category, of six-months duration and 4 Short-Term projects. The average funding requested was lower than that of the first call at 130,000 Euro, due to the Short-Term applications which brought the overall average down. Once more, a **good geographical coverage** could be observed with **12 EU countries** represented as coordinators and **18 US states**, surpassing that of the 1st Call, also possible due to intensified efforts to promote the NGIatlantic.eu **[Twinning Lab](#)** that provides a dedicated online space for researchers, innovators and start-ups that want to establish complementary partnerships in addressing NGI challenges across a number of topics.

FIGURE 3 INSIGHTS OF *EU-US PROPOSALS RECEIVED IN OPEN CALL 2*

As explained in Section 4.1, Open Call #2 introduced a new stream for proposals focussed on the interconnection of EU-US infrastructures for NGI experimentation, resulting in 9 applications, 5 of which intersecting with the other NGI topics maintained from Open Call #1. Of the other topics, as also in Open Call #2, a number of proposals addressed more than one topic with 11 applications involving Privacy and Trust Enhancing Technologies, 12 involving Decentralised Data Governance and 9 within Discovery and Identification Technologies.

As to the type of organisations applying for funds, once more, on the EU side the majority coordinating were from industry with 11 applications, all of which are SMEs, followed by Research and Academia with 8 and finally, 4 from non-profit organisations. On the US side, similar to Open Call #1, the majority of US Partners were from Research & Academia sector with a total of 16, 7 from large companies & SMEs, and 4 from non-profit organisations.

4. NGIATLANTIC.EU FUNDED EXPERIMENTS - DECEMBER 2020

4.1. Fairteam

Project Coordinator (EU): **Fairkom Gesellschaft**
Country of the EU Coordinator: **Austria**
Organisation Type: **SME**
State of US partner: **Delaware**
NGI related Topic: **Privacy and Trust Enhancing Technologies**
Call Reference**: 1**
Starting date: **1/10/2020**

PROJECT PARTICIPANTS

EU	US
fairkom	**Rocket.Chat**
Roland Alton: CEO, IT professional with focus on management of open source projects, member of the board, also teaching at Vorarlberg University of Applied Sciences. Role in the project: supervisor. Armin Felder: CTO, IT professional with focus on programming, large scale deployments, identity management systems. Role in the project: technical consulting, mapping of groups in ID management systems. Dimitri Hofer: support technician and key account manager: coordination of experiments. Christian Fischer: UX specialist Florian Koller: user experience evaluation	Sing Li: Chief Opportunities Officer. Role in the project: supervisor. Milton Rucks, product manager of Rocket.Chat's team collaboration vertical Rodrigo Nascimento, core developer Renato Becker, core developer Gabriel Engel: Rocket.Chat Founder. Role in the project: approve

TABLE 1 PARTICIPANTS IN THE FAIRTEAM EXPERIMENT

The EU – US FairTeam project is addressing the NGI topics of Privacy and Trust enhancing technologies and Decentralised Data Governance, in an open source and cloud-based workspace environment to improve the way that teams can be mapped together, based on the concept of "fairness". In addition, the project is also addressing the integration of identification technologies and vertical integration between the communication services.

Making a service available for independent teams is usually happening at a mature level of a product stage. **Rocket.Chat** supports federation of instances, which perfectly covers the use case of connecting chat groups among instances. Rocket.Chat also supports "discussions",

Co-funded by the Horizon 2020 Framework Programme of the European Union

which inherit team members of a chat group. However, within an organisational perspective, we want to be able to assign selected chat groups to a team as well.

Several solutions have been discussed in issue 658 for various use cases, but none has been really satisfying yet for a wider deployment. In this project, the team will be focussing on the mapping of group members with a nextcloud circles via API or with an identity managements system like keycloak to Rocket.Chat teams. As a user base the team can rely on **fairchat**, which is a large RocketChat based platform provided by fairkom to private users, NGOs and several companies. **fairchat.net** can be used with any browser and also has its own apps in the Android and Apple stores.

Two experiments will be carried out, working closely with the Rocket.Chat development team.

Experiment 1

On the fairchat staging server the team will be testing:

- 1a Synchronizing teams via IDM / LDAP;
- 1b Synchronizing teams via API.

Experiment 1 will be conducted with 3 teams with 3 groups each, with a total of 20 users. The following groups active in fairchat will be approached to participate: **fairkom-team** (fairkom team members, closed group located in Austria), **fairmove-it** (network of IT professionals for sustainability projects, located in German speaking countries, open channel) and **transition-austria** (network of Transition Towns, open channel).

Experiment 2

On the fairchat production server the team will be testing:

- 1a Synchronizing teams via IDM / LDAP;
- 1b Synchronizing teams via API.

Experiment 2 will be conducted with 5 teams with 3 groups each, with a total of 50 users. The following groups active in fairchat will be approached to participate: **fairmove-it** (network of IT professionals for sustainability projects, located in German speaking countries, open channel), **transition-austria** (network of Transition Towns, open channel) and **fairmeeting** (users who purchased a fairmeeting or fairteaching video conferencing package, located in Europe).

The EU partner, Fairkom Gesellschaft, based in Austria, is a competence network for open source based solutions, with a focus on the use of online services for the fair use of shared resources in a number of fields related to cloud services, software development, commons, and social design. Their vision is based upon an orientation towards the common good that does not reward shareholders, but people who use open source tools to set meaningful activities, which is very much in line with the human-centric vision of the NGI initiative of the EU.

The US partner, Rocket.Chat Technologies Corp., whose headquarters is located in Delaware, USA, with a global presence, has access to larger deployments and know-how of organisational structures when it comes to defining team memberships; they were selected for this NGIatlantic.eu team to offer test beds to run the EU – US NGI experiments, as they have significant experience in know-how for small and large deployments, including tuning, and adaptation of such systems over diverse network and hardware topologies from the smallest to largest clusters of distributed servers servicing millions of global users.

Co-funded by the Horizon 2020
Framework Programme of the European Union

IMPLEMENTATION PLAN

In the preparation phase of this experiment, the Rocket.Chat API is being extended with functions for managing teams. This allows us to connect various group management systems.

The experiments will be performed with selected Rocket.Chat instances and group management systems. This requires the team to set up either staging copies of this instances or work on the production platforms to experiment with, depending on the willingness of the invited leads to take a little risk in the experiment, that it may fail. Recovering procedures will be prepared, in case something goes badly wrong.

As group management systems, Nextcloud and Keycloak will be used. In the experiment, the team will try to align the APIs so that creating or deleting a team or adding and removing team members in Rocket.Chat is being synchronised with the groups in the IDMs.

The experiment will be starting with a laboratory set-up to be classified as TRL 4 and it will be moving quickly on to a TRL 5 simulation of a real-world setting. Experiment 1 is to be classified as a demonstration of a relevant environment (TRL 6) whilst experiment 2 is a system prototype demonstration in operational environment (TRL 7).

After the successful completion of this project, the experiment ideally can enter TRL 8 and qualify the results soon as system complete to be applied also in customer projects.

EXPECTED IMPACT

- Enhanced EU – US cooperation in Next Generation Internet, including policy cooperation;
- Reinforced collaboration and increased synergies between the Next Generation Internet and the Tomorrow's Internet programmes;
- Developing interoperable solutions and joint demonstrators, contributions to standards.

EXPECTED RESULTS

Provision of feedback to the Rocket.Chat team regarding the API extension for team management; a report on how to integrate keycloak as an IDM with the RocketChat team concept and how to mirror existing department structures. Lastly, a report on how to mirror nextcloud circles with RocketChat teams.

FUTURE PLAN

Nextcloud + Rocket.Chat + Keycloak is a top selling open source stack combination and Fairteam can lead to the vertical integration of market leading open source framework to manage and synchronize teams. At the same time, Rocket.Chat intends to integrate the team feature in the EE version while fairkom will offer the team feature as an extension to its portfolio of SaaS and consulting services.

Expected TRL at experiment completition : 5 (Experiment 1) and 7 (Experiment 2)

Co-funded by the Horizon 2020 Framework Programme of the European Union

4.2. Measuring Multi-Carrier Cellular Access International Roaming Performance

Project Coordinator (EU): **Universidad Carlos III de Madrid**
Country of the EU Coordinator: **Spain**
Organisation Type: **Academia**
State of US partner: **Illinois**
NGI related Topic: **Discovery and Identification Technology**
Call Reference**: 1**
Starting date: **1/11/2020**

PROJECT PARTICIPANTS

EU	US
Universidad Carlos III de Madrid Marcelo Bagnulo, associate professor UC3M. He is the coordinator of the project and will be PI for UC3M. **TelefonicaResearch** Andra Lutu is an Associate Researcher with TelefonicaResearch. Andra will work on the measurements of mobile access performance while roaming in the EU.	**Northwestern University** **Fabián E. Bustamante** is a professor at Northwestern University. Fabian will work on the measurements of mobile access performance while roaming in the EU

TABLE 2 PARTICIPANTS IN THE MCCA INTERNATIONAL ROAMING PERFORMANCE EXPERIMENT

Multi-Carrier Cellular Access (MCCA) providers such as GoogleFi erupted recently in the mobile Internet access market claiming to offer enhanced performance at reduced prices. MCCA providers dynamically attach to different carriers based on real time measurements of performance. The goal of the project is to do an experimental comparative study of the performance of these emerging MCCA providers and the incumbent cellular providers with a special focus on the performance while roaming. The proposed methodology is to acquire MCCA services and traditional mobile services in the US and experimentally compare their performance both at home (US) and while roaming (EU).

Multi-Carrier Cellular Access (MCCA) allows cellular users to dynamically connect to different cellular networks without switching their SIM card. There are a number of advantages that can be realized through MCCA. MCCA users can improve coverage and performance, combining measurement-based dynamic carrier selection and predefined preferences on radio access technology (RAT) use (3G, 4G, 5G, wifi), irrespectively of which carrier is offering it a particular location. Google Fiis one provider offering commercial MCCA services. Google Fi carrier selection policy relies on static preferences regarding available carriers and RATs, while also leveraging local or crowdsourced monitors to track various metrics to inform carrier selection [Yuan18]. While MCCA services are becoming increasingly popular, little is known about their technical intricacies beyond what is described in their commercial offers. To date, there is no quantitative-based analysis of the relative performance benefits of MCCA services compared to traditional services. In particular, there is no information about MCCA performance while

Co-funded by the Horizon 2020 Framework Programme of the European Union

roaming, which is one of the key selling points of several of these offers. Measuring international roaming performance is challenging, as it requires international cooperation and logistics are daunting. In a previous study, we analysed the European roaming ecosystem for traditional carriers and its performance [Mandalari18]. We observed that international roaming imposes performance penalties because traditional carriers systematically use Home Routing (HR) while roaming. HR implies that traffic from and to the roaming device is routed back to the home carrier before reaching the Internet, instead of exiting directly through the visited-network Internet access. This extra leg adds latency, penalizing roaming performance. While it would be technically feasible to provide local Internet access to roamers (called Local Break-out), this is avoided by traditional carriers in order to enable visibility of the roaming-client traffic by the home operator. Visibility of roaming-client traffic is deemed necessary to avoid disputes and abuses involving the computation of the traffic carried by the visited network on behalf of the roaming client. Seizing the performance gains resulting from local breakout is yet another potential benefit that MCCA can achieve for their customers in the international roaming context.

The objective of the project is characterise the performance of MCCA services while roaming. We will collect quantitative information to help us understand the benefits and downsides of MCCA service when used abroad. We will focus our analysis to clients from MCCA operators in the US that are roaming in the EU and investigate how their performance compares to both traditional US carriers while roaming and local EU carriers while providing local services in the visited network.

FIGURE 4 VIDEO INTERVIEW WITH MARCELO BAGNULO, AVAILABLE AT HTTPS://YOUTU.BE/4ZD4VX2CUUS

EXPERIMENT DESCRIPTION

As there is little knowledge about both the mechanisms used by MCCA providers and the resulting performance, we propose to design and execute a first set of experiments focussed on functionality and a second set of experiments focussed on performance. Experiments targeted to discover functionality. The team will design a set of experiments to infer how MCCA providers work in different scenarios. In particular, the team pretend to learn about the carrier selection mechanism, by observing how different parameters affect the carrier selection. Among the parameters we will look into we include, available carriers, available radio access technology (including the different cellular technology generations, wifi), signal strength and other radio parameters. In addition, The team would like to understand how does the MCCA provider clients exit to the Internet And how session continuity is provided. Experiments targeted to discover performance. The team will design experiments that allow us to measure

Co-funded by the Horizon 2020 Framework Programme of the European Union

and compare the performance of the communication service provided by MCCA providers and the one offered by traditional providers. The team will use experiments to assess general performance metrics, such as latency and throughput and will also perform experiments to assess the performance of specific applications such as The web browsing, streaming, interactive applications. The team propose to do a first phase of local experiments to learn how MCCA services available locally behave and perform. The team compare their performance with the underlying carriers that the MCCA provider is relying on. In a second phase, The team will measure the behaviour and performance of MCCA commercial services while roaming. To do this, The team will deploy eSIMs of MCCA services hired in the US in the EUand perform the tests. The tests will allow us to learn how MCCA service work abroad. In particular, The team will learn what type of breakout is used (Local Breakout, Home routing or else). The team will also execute the performance tests. The team will compare the resulting performance with the one of the native users of the landing carrier as well as with the roaming service offered by traditional US carriers.

The project addresses the Discovery and identification technologies topic ofNGIatlantic.Throughout the project, The team propose to experimentally quantify the impact in quality of experience that results from the use of MCCA technology to access different types of data, in particular multimedia content and interactive services. Dueto the international nature of the NGIAtlantic, The team will focus on users that are roaming from the US to the EU. The key enabling technology that will be investigated in MCCA. MCCA technology has the potential to radically change the mobile communication landscape. MCCA is the approach chosen by large content providers (Google, Apple) and by device manufacturers (Samsung, Huawei, Apple, and Google to an extent) to enter in the mobile Internet service provider market. These players are by themselves so large, thatthe potential impact in the mobile Internet market is huge.However, little is known about the real advantages that MCCA technology provides in terms of accessibility, performance and quality of experience. Moreover, the actual benefits realized heavily depends on the specifics of the MCCA mechanism implemented. The goal of the project is to obtain ground truth information about the real improvements in the quality of experience that commercially available MCCA services provide. The team propose doing so through experimentation in the wild.Through the experiments, The team pretend to shed light on different aspects of the MCCA services available including: carrier selection policy in terms of access technology, coverage, and other metrics, performance and quality of experience impact, use of Local breakout or Home routing while roaming, roaming agreements, use of wifi and its impact. The experimental results of the proposed investigations can significantly impact the adoption and evolution of MCCA technologies both by the incumbent operators and by new players. Also in the context of roaming, depending on the results of the experiments, traditional carriers may have to revisit their current approach regarding roaming agreements and extensive use of Home-routing. Indeed, if MCCA is proven to provide significant performance benefits while roaming, traditional carriers may need to explore alternative approaches to the currently used to remain competitive.

IMPLEMENTATION PLAN

The project will be structured in the following tasks.

Task 1: Preparation - Duration: M1 to M4

In this task, the team will do the preparation work for the experiments. This includes selection of the MCCA providers to be tested and buy their eSIMs, selection of hardware used for the testing and purchasing it, identify the MNOs the MCCA users will be landing and obtain SIMs from them (to be used to benchmark the roaming service of the MCCA provider against a native visited network user), design and implement the experiments. perform local experiments to validate the implemented tests.

Task 2: Testing - Duration: M3 to M5

Co-funded by the Horizon 2020 Framework Programme of the European Union

In this task the team will execute the measurement campaign planned in the Preparation task. This will include several rounds of measurement and coordination of measurements across multiple countries.

Task 3: Analysis of results - Duration: M5 and M6

In this task, the team will analyse the results obtained in the measurement campaigns. They will start processing as soon as The team have the earliest results so to use the analysis to improve the subsequent tests.

Task 4: Management, coordination, dissemination and communication - Duration: M1 to M6.

This task will take care of managing the project, including the coordination between the EU and US partners, as well as with the other entities hosting probes. It will also take care of the dissemination and communication actions, including the publication of a repository including data sets and other information regarding the experiments.

EXPECTED RESULTS

MCCA has received significant interest from the industry, perceived as a potential threat by incumbent Mobile Network Operators (MNOs) and as an opportunity by potential new players to this market (e.g. Google, Apple, Samsung, Huawei). The irruption of such large players has the potential to drastically change the mobile provider market altogether. MCCA enables the provision of mobile communication services without the need of having a radio access network infrastructure (much like a virtual MNO), but unlike virtual MNOs which are usually perceived as a lower quality offer, the MCCA technology has the potential to offer an even better performance than traditional MNOs. This is particularly true in the roaming context, as MCCA can improve performance by supporting Local Breakout. However, there is little information about the performance benefits that can be achieved through MCCA in general and no information at all in the case or roaming. To date, The team are only aware of the work in [Yuan18] which analyses and models conflicts between inter and intra provider policies for selection of radio access in Google Fi. They perform a “small-scale user study” [sic] focussed on measuring the impact of policy conflicts in performance. In [Mandalar18], The team performed an analysis of performance of mobile communications while roaming for traditional MNOs. The work proposed in this project aims to inform the future evolution and adoption of MCCA commercial services. It expects to provide quantitative information about the performance benefits achieved by MCCA, focusing in US-EU roaming scenarios. This information will inform both the MCCA technological evolution as well as the response from traditional MNOs, in particular whether currently used home routing strategies can be sustained in the future in the light of their performance implications.

More specifically, the results of the project will include:

- Definition and implementation of a set of tests to functionally characterize MCCA services, including carrier selection and breakout strategy while roaming;
- Definition and implementation of a set of tests that allow the comparative analysis of the performance of MCCA services and traditional mobile services;
- A number of experiments characterising selected MCCA services using the aforementioned tests;
- Results of the experiments and their analysis. The resulting dataset will be made public.

Co-funded by the Horizon 2020
Framework Programme of the European Union

EXPECTED IMPACT

The Project will address the impacts in relation to the NGI initiative:

- Enhanced EU –US cooperation in Next Generation Internet, including policy cooperation.: The project will serve to nurture the existent relationships between the US and the EU partners of the project consortium. Because the experiments will also involve additional institutions in the EU to host probes to measure roaming in different countries, the project will also serve to create new links between the US institutions and these other EU entities involved in the experiments.
- Reinforced collaboration and increased synergies between the Next Generation Internet and the Tomorrow's Internet programmes.
- An EU -US ecosystem of top researchers, hi-tech start-ups / SMEs and Internet-related communities collaborating on the evolution of the Internet: During the lifetime of the project and beyond, the project will enable a number of joint activities between the US and the EU partners of the project and potentially include additional entities of both sides contacted throughout the project (e.g. host of probes in different countries). The joint activities expected include joint publications in top conferences/journals, exchange of researchers and internships of master and PhD students of the different entities, joint data repositories to publish measurement results obtained throughout the project (all measurement results will be publicly available after proper anonymization).

Co-funded by the Horizon 2020
Framework Programme of the European Union

4.3. CacheCash Experiment

Project Coordinator (EU): **Sorbonne University**
Country of the EU Coordinator: **France**
Organisation Type: **Academia**
State of US partner: **New York, Washington**
NGI related Topic: **Discovery and Identification of Technologies**
Call Reference: **1**
Starting date: **1/10/2020**

PROJECT PARTICIPANTS

EU	US
Sorbonne University The Principal Interface Timur Friedman is a faculty member at Sorbonne Université, and a member of the LIP6 and LINCS computer science laboratories. His research expertise is in the area of internet measurements. The principal technical work on the project is carried out by Berat Senel, currently a research engineer at Sorbonne Université,	**New York University's Tandon School of Engineering and university of California (UC) Berkeley** Principal Investigators are Justin Cappos, an associate professor in the Computer Science and Engineering department at New York University's Tandon School of Engineering, and Rick McGeer holds a Ph.D. from the University of California (UC) Berkeley where he is principal developer of EdgeNet.

TABLE 3 PARTICIPANTS IN THE CACHECASH EXPERIMENT

Delivery of large chunks of content, such as video, accounts for a substantial percentage of all internet traffic. Such content, when distributed by large corporations, benefits from CDNs, with nodes hosted by internet service providers (ISPs) providing local service to end users in those ISPs. CacheCash (https://ssl.engineering.nyu.edu/projects#cachecash), the CDN technology under testing by the project, has the potential to change the nature of CDNs by involving the end users themselves directly in serving content through machines that are under their control. Putting users in charge can lead to a wider range of content benefitting from CDNs. CacheCash provides a service in which interested users run caches, and they are incentivised to participate by receiving a crypto-currency (Cachecoin) in exchange for serving content to other users. Both cryptographic and economic analyses demonstrate that the incentives lead CacheCash users to honestly serve content. Analysis has also shown that CacheCash can scale to meet the workload of even the most popular services used today.

CacheCash, the CDN technology, will be tested on EdgeNet to evaluate its performance at a distributed edge network. The measurements would present performance results to see if any additional requirements to be handled before running CacheCash at scale by involving the end-users.

The experiment is based on two cutting-edge technologies:

- **CacheCash, the CDN technology** that being tested, has the potential to change the nature of CDNs by involving the end-users themselves directly in serving content through machines that are under their control. It provides a service where interested users run

Co-funded by the Horizon 2020 Framework Programme of the European Union

caches and are incentivized to participate by receiving a cryptocurrency in exchange for serving content.

- **EdgeNet** brings service deployment to the edge cloud, which is the key enabling technology, with customized extensions for Kubernetes. It has been successfully piloted in Fed4FIRE+ open call experiments and is now ready for larger-scale experiments.

The CacheCash experiment addresses the "Discovery and identification technologies" topic of the call. In testing a new blockchain-based CDN technology, it addresses the topic's focus on "new methods of search, discovery, and access of large heterogeneous data sources" and "multi-media content". CacheCash's focus on end-users addresses the "personalised information retrieval" aspect of the topic, and the evaluation of the efficacy of CacheCash will be in terms of "improved quality of experience".

The key enabling technology is the edge cloud technology provided by PlanetLab Europe through EdgeNet. Edge cloud nodes are an integral part of the 5G vision. EdgeNet is based on contemporary de facto industry standard technologies Docker and Kubernetes. The EdgeNet cluster consists of approximately 50 nodes located in North America, South America, and Europe. It has been successfully piloted in Fed4FIRE+ open call experiments and is now ready for larger scale experiments engaging real internet users.

FIGURE 5 VIDEO INTERVIEW WITH TIMUR FRIEDMAN, AVAILABLE AT HTTPS://YOUTU.BE/3N2QF8GJYGO

There are several ways to look at the innovative aspects of this work. First, the innovative technology we deploy has never been tried at a scale even remotely approaching this. It is unknown how the cryptographic techniques for distributing content securely in a way that economically incentivises cheating will work given the latency, non-transitive connectivity, and scalability constraints of the internet. Further, our effort to use real world software distributed to real world users will provide a basis for understanding how such technology will work in practice and at scale.

To date, most efforts around building distributed data delivery have focused primarily on data storage or on trusted nodes that deliver content. The expected impact of the CacheCash experiment will be to push the boundaries more toward low latency and bandwidth, which are highly valued in the CDN space. The relaxation of trust assumptions in existing platforms also will open up the set of participants. Somewhat counterintuitively, if we are able to demonstrate

Co-funded by the Horizon 2020 Framework Programme of the European Union

that high performance can be provided by untrusted nodes, this may enable even CDNs with trusted nodes to have better security as they may alter their architecture to remove trust assumptions without harming performance. By having our experiment be used in a practical deployment by real internet users, the team will be able to evaluate the strengths of such approaches, the weaknesses, and to identify research challenges that must be tackled to allow these technologies to become widespread.

IMPLEMENTATION PLAN

The CasheCash team will use PlanetLab Europe (https://planet-lab.eu/), which is one of the EU Fed4FIRE+ platforms. PlanetLab Europe offers an edge cloud service that is particularly apt for 5G experimentation. Researchers can deploy their experiments on containers that run on servers distributed around the edge of the internet. These servers are completely open to the internet, meaning that researchers can deploy experimental services to real users, and those services will be subject to real internet conditions (latencies, routing changes, cross traffic, etc.) Currently, about 50 such servers in Europe and North America are available.

PlanetLab Europe is migrating these servers to the new EdgeNet (https://edge-net.org/) technology that Sorbonne University (EU/PI Friedman) is developing in a currently unfunded collaboration with NYU Tandon School (US/PI Cappos) and US Ignite (US/PI McGeer). With EdgeNet, we update PlanetLab Europe to the most recent and widely used containerisation technologies: Docker, and the emerging industry standard for container orchestration, Kubernetes.

EXPECTED IMPACT

Sorbonne University's collaboration with US Ignite and the NYU Tandon School is an outcome of the Global Experimentation for Future Internet (GEFI) workshop, 26-27 October 2017, in Rio de Janeiro, Brazil. It became clear that there was a shared interest in working towards a common successor platform to the existing LXC-based PlanetLab technology, to move to the contemporary de facto industry standards of Docker and Kubernetes. This collaboration has, to this date, not been formalised through joint funding, but it is a very real and enduring collaboration, featuring weekly meetings to advance our joint goals. The US side has in the past benefitted from an NSF EAGER grant to support some of this work, and the EU side has received partial support from French national funding. At times, the development initiative has been taken in the United States and Canada, and at times in France. Our ability to continue together has, however, constantly been in question due to the uncertainties on both sides of the Atlantic regarding ongoing financing. The NGIatlantic.eu funding is the first official recognition of our collaboration, enabling the US partners to seek similar funding, and it sets the entire collaboration on a firmer footing for the future. This experiment also increases the appearance of all partners and allows them to exchange information about the NGI Generation Internet and the Tomorrow's Internet programmes to seek further potential collaborations in the future.

Besides all of these, a successful CacheCash trial on EdgeNet means that CacheCash is capable of involving any Kubernetes-based cluster such as EdgeNet to run caches on its nodes, and EdgeNet can host any CDN solution. This mutual interest will allow both software to be demonstrated to the Inter-Atlantic community. Because both are open source solutions, they can either be adopted by high-tech startups or empower spin-off startups. Moreover, EdgeNet, as a testbed, can attract the EU-US ecosystem of top researchers to run their experiments.

Co-funded by the Horizon 2020
Framework Programme of the European Union

EXPECTED RESULTS

The project will gather a large, real world data set using cutting edge blockchain data distribution technology. Few studies of this sort have been done, and regardless of the success or failure of the technology, solid evidence that shows the reality today would be helpful. Vast sums of money have been invested in such technologies (e.g., over 250M USD on FileCoin) and understanding their likely real-world properties may help to inform future research and development.

Prior edge CDN efforts like Akamai's NetSession or Quilt provided good performance, but largely failed due to a lack of incentives for participants. These steps will help us see if a secure system with good game-theoretic anti-cheating properties and (later) financial incentives can still provide strong performance. By having altruistic use initially, we can experiment and understand future performance issues which may arise.

FUTURE PLAN

The purpose of the experiment will be to trial CacheCash in the real world with real internet users consuming content that is of interest to them. The content will initially be from The Update Framework, or TUF (https://theupdateframework.io/). This technology is used by Microsoft, Google, Amazon, IBM, Docker, RedHat, CloudFlare, and many other companies. It is being integrated right now to distribute Python packages on their community repository. We will slowly start by first evaluating the properties of CacheCash with a small subset of TUF users. As our capacity and stability allows, we will add more users and projects to get a more representative data set. Moreover, another exploitation result may be the adoption of CacheCash by a start-up working on the content delivery field.

Sorbonne University plans to use the EdgeNet software to run PlanetLab Europe. With a successful trial of CacheCash on EdgeNet, one exploitation result will be that PlanetLab Europe's users migrated into EdgeNet, and the migration of PlanetLab Europe's existing infrastructure into the EdgeNet cluster will follow that.

As for technology readiness levels of EdgeNet and CacheCash, the team might evaluate EdgeNet at about TRL 3 Experimental Proof of Concept, as there is a functioning instance of EdgeNet that is supporting experiments. CacheCash is closer to TRL 2, a technology concept that has been formulated. The aim is to move CacheCash to TRL 3, as it will involve an experimental proof of concept of the technology being deployed on EdgeNet. The work will also move EdgeNet closer to TRL 4, as using the platform to support an ambitious experiment such as CacheCash brings it further towards being lab-validated.

4.4. Self-Certifying Names for Named Data Networking

Project Coordinator (EU): **Athens University of Economics and Business - Research Centre (AUEB)**
Country of the EU Coordinator: **Greece**
Organisation Type: **Academia**
State of US partner: **Colorado**
NGI related Topic: **Privacy and trust enhancing technologies**
Call Reference: **1**
Starting date: **5/11/2020**

PROJECT PARTICIPANTS

EU	US
Athens University of Economics and Business (AUEB) Project Manager: George Xylomenos(M), Professor Technical Manager: George C. Polyzos (M), Professor Vasilis A. Siris (M), Professor, DID-related activities and the integration with NDN Nikos Fotiou (M), PhD, Researcher, DID-related tasks and security-related experiments design Yiannis Thomas (M),PhD, Researcher, NDN integrationand network-related experiment design Iakovos Pittaras (M), PhD student, execution of experiments	**University of Memphis (UofM) members:** Christos Papadopoulos (M), Professor, running experiments on the NDN testbed and the UofM servers.

TABLE 4 PARTICIPANTS IN THE SCN4NDN EXPERIMENT

The SCN4NDN project will experiment with the merger of two promising NGI technologies: Information-Centric Networking (ICN)[1] and Decentralized Identifiers (DIDs)[2]. ICN has been on the spotlight of many research efforts for more than a decade. It has been explored as a standalone future Internet architecture, as well as an enabler for other NGI architectures, including 5G, IoT, and architectures focused on big dataand/orcyber security. ICN's goal is to enable fast and secure content dissemination

[1] G. Xylomenos, C.N. Ververidis, V.A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K.V. Katsaros, G.C. Polyzos, "A Survey of Information-Centric Networking Research," IEEE Communications Surveys & Tutorials, vol. 16, no. 2, pp. 1024-1049, 2014.

[2] W3C Credentials Community Group, "A primer for decentralized identifiers," 2019; available athttps://w3c-ccg.github.io/did-primer/

Co-funded by the Horizon 2020
Framework Programme of the European Union

by leveraging direct and intrinsic information identification; this allows supporting multicast, multipath, and caching, as well as novel trust mechanisms.

In this project, the team will experiment with the Named-Data Networking (NDN)[3] ICN architecture. A DID is a new form of identifier under standardization by W3C. A DID system can be regarded as a key-value lookup system, where the key is the identifier, the DID, and the value is a DID document. A DID document contains "properties" including information that can be used for verifying DID ownership, as well as the document's integrity. In this project we will experiment with a self-sovereign DID system, i.e., a system where DID documents are managed by the DID owners themselves (as opposed to systems where DID documents are managed by a trusted "registry"). The project is driven by the goals of improving ICN security, enhancing content-owner's privacy, and enabling decentralized data governance. To this end, the project will validate and evaluate a solution that uses self-sovereign DIDs to protect content authenticity in NDN. The integration of DIDs into NDN is expected to provide robust security against fake content (that is, content that does not correspond to its name) without relying on third parties, as well as efficient spam prevention. The project will evaluate the use of per-content DIDs. These DIDs will be randomly generated and they will not reveal any information related to the (content) owner (as opposed, for example, to a digital certificate bound to an owner-specific identifier): this is expected to provide enhanced privacy and resilience against user censorship attacks, since it will not be possible to track and/or filter content belonging to specific owners. Furthermore, the project anticipates improved decentralized data governance by enabling content owners to specify and integrate into their content items lists of authorized storage nodes, as well as basic access control policies. The project vision includes achieving these goals without degrading key functionalities of ICN (and of NDN in particular), including support for advanced traffic management (such as multicast and multisource).

FIGURE 6 VIDEO INTERVIEW WITH NIKOS FOTIU, AVAILABLE AT HTTPS://YOUTU.BE/CSJIHNKBEPI

The proposed solution builds on the emerging paradigm of Decentralized Identifiers (DIDs), a new form of identification under standardization by W3C which allows self-sovereignty. DIDs

[3] V. Jacobson, D.K. Smetters, J.D. Thornton, M.F. Plass, N.H. Briggs, R.L. Braynard, "Networking Named Content,"Proc. ACM CoNEXT 2009, Rome, Italy, December 2009

are associated with a DID document that contains cryptographic materials that can be used for performing various security-related tasks.

The solution proposes the use of DIDs as content names. A content owner will be able to define the public keys that correspond to various “roles” and/or “rights” in the DID document. For example, an owner will be able to define who can “host” an item, who can “change its contents”, or even who can “change its encoding.”

Additionally, a content owner will be able to digitally sign some “metadata” and include them with the corresponding content items; then, the cryptographic material included in the DID document will be used for verifying the metadata signature. While most DID systems assume a registry for managing DID documents, the proposed solution uses an innovative concept that enables decentralized and autonomous DID document management.

This approach allows DID documents to be integrated in the content items themselves in a secure and verifiable way. With respect to key aspects of NDN (and ICN in general), the solution is expected to have (measurable) impact in the following areas:

- **Security**. The solution allows the detection of fake replicas of a content item without relying on 3rdparties, i.e., it can be used not only for verifying an item’s integrity, but also an item’s authenticity;
- **Multicast and Caching**. The solution favours the so-called “channel” mode of delivery which facilitates multicast, but it creates some challenges to caching, which we address;
- **Multisource**. The solution supports content name “aliases.” Using aliases, it will be possible to implement multisource in a transparent way;
- **Multipath**. Our solution allows the “definition” of authorized in-network nodes allowed to modify “control fields” used for orchestrating multipath transmission.

The SCN4NDN experiments will be conducted on the Named Data Networking (NDN[4]) testbed. The NDN testbed is a global shared resource, with origins, management, and majority presence in the USA. It was created for research purposes and it relies on software routers at several participating institutions, application host nodes, and other devices. Currently, the testbed includes 37 nodes located in the USA, Portugal, Korea, Thailand, Norway, Japan, Brazil, Italy, and France. The testbed is centrally managed and runs a routing protocol that allows communication with every other node. The testbed enables broad areas of research on virtually any type of application and forms a network for the real-life evaluation of NDN technologies, protocols, and applications. The project will have access to (through the UofM) a 24-server lab environment that can emulate dozens of virtual machines connected as various network topologies for controlled experiments. The project will implement two types of experiments: end-to-end experiments and incremental deployment experiments. With end-to-end experiments, our solution will be used at the application layer of the communicating endpoints and it will be adjusted to be compatible with the NDN protocols as deployed in the NDN testbed. For incremental deployment experiments, the project will consider a dual approach: (i) a local testbed with fewer nodes with full access will be deployed at AUEB premises, and (ii) a number of VMs connected to the NDN testbed will be provided by UofM, where modifications to the NDN software stack and incremental deployment of our solution will be explored. The proposed experiments will address key NGI topics, including privacy and trust enhancing technologies, decentralized data governance, and discovery and identification technologies. The project is expected to investigate a multitude of traffic patterns, content types, content delivery modes, as well as security scenarios.

[4] https://named-data.net/ndn-testbed/

Co-funded by the Horizon 2020 Framework Programme of the European Union

IMPLEMENTATION PLAN

The project's DID implementation is based on the specifications of "DID:self", a DID method we have published in Mobile Multimedia Laboratory, "DID:self method specification"[5] A Python3-based software library[6] currently at TRL3, provides DID document (self-)management functionality by implementing the corresponding Create, Update, and Read methods. The Create method is used for creating the initial DID document, the Update method is used for modifying it (including key rotation), and the Read method outputs the final DID document and a proof chain that can be used for verifying the binding between a DID and the corresponding document. An application layer solution will combine our implementation with python-ndn[7] , an NDN client library, to provide the desired functionality as follows. For each content item, the application will generate a public-private key pair from the Curve25519 elliptic curve. The public key will be then treated as a DID and will be used as the name of the content item. The private key that corresponds to a DID/content name, will be used for signing content item metadata. Metadata will include information that can be used for verifying various properties of an item, such as its integrity, version, type, and alternative names. Eventually, the DID document that corresponds to an item, the appropriate proof, and the signed metadata will be included in the payload of a NDN packet. Using python-ndn, the application will interact with the "NDN Forwarding Daemon" (NFD) of a testbed node in order to perform the appropriate (ICN) operations. Our final software will be at least at TRL4, and possibly at TRL5.

The experiments will consider the following content item types: immutable items, mutable items that may have multiple "representations" (e.g., an image file stored under different encodings), and mutable items that may have different "versions." When mutable items are used, we need a mechanism to distinguish among different versions or representations of the same item: this is achieved by including the corresponding information in the metadata (e.g., using a "version" and a "type" field). In addition, for mutable items, the DID document may specify the public keys of the entities that are allowed to generate new versions and/or representations of an item. Similarly, our experiments will consider two content delivery modes: the document mode and the channel mode. In the former mode, a content name will be mapped to a data "bucket," e.g., an image file, whereas in the second case, a content name will be mapped to a data "stream," e.g., a streamed video. During the project we will perform experiments related to the following ICN functionalities:

- **Caching and Multicast Caching and multicast functionalities are provided by the NFD.** Therefore, experimentation with these functionalities simply requires the allocation of a suitable topology within the testbed and the application of an appropriate workload at the endpoints.
- **Multisource**. In order to experiment with multisource, we will implement content name "aliases". In particular, a content item will be provided by multiple sources and each source will use a different content name. All these name "aliases" will be included in a special field of the item's metadata called "alsoknownas". A client application will request simultaneously many of the item's names, making sure that each request concerns different chunks. This will result in the client receiving different chunks of an item from different sources.

[5] https://github.com/mmlab-aueb/did-sel

[6] Mobile Multimedia Laboratory, "DID:self method python library," available at https://github.com/mmlab-aueb/did-self-py

[7] Named Data Networking, "A Named Data Networking client library with AsyncIO support in Python 3," available at https://github.com/named-data/python-ndn

Co-funded by the Horizon 2020 Framework Programme of the European Union

- **Multipath**. In order to experiment with multipath, the team will add a new node in the NDN testbed which will be attached to, at least, two nodes, located in different locations. Furthermore, we will allow the content owner to include in DID documents the public keys of in-network nodes that are authorized to modify "control fields" used for orchestrating multipath transmission
- **Security**. Our security related experiments will focus on detecting fake content. In particular, the team will consider cases where an attacker signs a fake item using a revoked or a breached key. Our scenarios will consider various types of compromised keys, including keys allowed to update the DID document, as well as keys allowed to generate new versions/representations of an item. In our experiments, the team will "inject" fake content from various points of the testbed.

EXPECTED IMPACT

With respect to the NGI initiative, the project is anticipated to have impact in the following areas:

- **Enhanced EU –US cooperation in Next Generation Internet,** including policy cooperation. Beyond ICN, and ICT research in general, we believe that our project can be a starting point for better future EU-US relations in science and technology: both partners, AUEB and UofM, through their active collaboration in organizing international events, such as the ACM SIGCOMM ICN conference, and in participation in international working groups, such as IRTF's ICNRG, have already established a fruitful relationship that guarantees a successful collaboration. The outcomes of the project will become an excellent showcase of cooperation between the two continents and will set up an example for other similar activities.
- **Reinforced collaboration and increased synergies between the Next Generation Internet and the Tomorrow's Internet programmes.** The project combines EU-based and US-based researchers and resources to experiment with networking architecture and components that are of interest to both the Next Generation Internet and the Tomorrow's Internet programmes. For instance, our "DID:self" method is applicable to a number of emerging authentication and authorization standards. Furthermore, the DID-based content authentication mechanism can be applied in other networking and application contexts, such as the emerging Inter-Planetary File System (IPFS), or even HTTP-based services.
- **Developing interoperable solutions and joint demonstrators, contributions to standards.** The project is expected to be a showcase of the merger of two emerging standards, managed by different standardization bodies. On the one hand, DIDs are primarily pursued by the W3C. On the other hand, ICN standards are mainly developed under the umbrella of the IETF. Both standardization efforts involve partners from academia and industry. Beyond the demonstration of the joint standards, the project is anticipated to inspire new activities in the respective standardization bodies. In particular, we expect to ignite discussions related to self-managed DIDs, as well as to novel content authentication mechanisms. An EU -US ecosystem of top researchers, hi-tech start-ups / SMEs and Internet-related communities collaborating on the evolution of the Internet. The team envisions that this project will not be a mere collaboration between two ICN pioneers but will also establish a permanent link between EU-US ICT research based on the Future Internet ICN approach. EU ICN research efforts are more human-centric, focusing mostly on security and trust, self-sovereignty, and distributed data governance. US efforts on the other hand prioritize deployment and real-world exploitation. Research teams on both continents will benefit from this complementary partnership.

EXPECTED RESULTS

For each considered ICN functionality, it is expected to obtain the following results:

- **Caching and Multicast.** For these functionalities, the project will compare this approach to “vanilla” NDN by experimenting with mutable and immutable items, as well as document and channel delivery modes. When it comes to mutable items and channel mode, we expect to see O(N) reduction in network overhead, where N is the number of recipients receiving the stream. On the other hand, when mutable items and document mode are combined, we expect to see an increase in the number of content “advertisements” required to keep caches up to date, since there is a need to distinguish among the various versions. The degree of the increase depends on two factors: (i) whether or not a new version of an item deprecates the old ones, and (ii) whether or not caches are capable of parsing our item metadata format (a case that will be explored using “incremental deployment” scenarios). Finally, it is expected to see similar performance between the project approach and vanilla NDN when mutable items are used (no matter the delivery mode).
- **Multisource and multipath** When it comes to multisource and multipath, the team will perform two types of experiments. Firstly, the team will utilize multiple sources/paths simultaneously and we will measure the improvement in throughput; this improvement is expected to be log(N) compared to vanilla NDN, where N is the number of paths/sources. Secondly, the team will use the additional paths/sources as a “fallback” mechanism and introduce network failures. In that case, the team will measure the amount of time required to re-gain the original throughput; this is expected to be a few milliseconds. For these experiments, the expected results will also be affected by whether or not in-network nodes can parse and verify DID documents.
- **Security**. With respect to the security aspects of our system we will perform two types of experiments. Firstly, the team will measure the time required for users to start receiving the correct version of a file in the presence of attackers that have access to the signing key of the content owner. Secondly, the team will measure the time required for users to start receiving the correct version of a file in the presence of attackers that have been authorized by the content owner to host/modify an item, but after some time they start behaving maliciously hence their privileges have to be revoked. In both cases, we expect to see results in the order of a few seconds.

Expected TRL at experiment completion: 4

Co-funded by the Horizon 2020
Framework Programme of the European Union

5. CONCLUSIONS

Every NGIatlantic.eu funded experiment has an entry in the online catalogue available at https://ngiatlantic.eu/funded-experiments, not only while the experiment is running but beyond, in order to track the impact and exploitation path of the project results. The catalogue contains details of the project vision, participant contact details, project results (expected and actual), exploitation and impact results.

Every six months, or based on the timing of the open calls' implementations, a printable version of the catalogue of the experiments is made publicly available via the NGIatlantic.eu website. This is the first version of this printable catalogue with the four running projects at the time of writing, as a giveaway to showcase NGIatlantic.eu results, creating awareness of the impact of the experiments carried out with NGIatlantic.eu cascade funding, ensuring global recognition and visibility on how the EU and the US are impacting and influencing NGI efforts.

A second iteration of this printable catalogue is foreseen in June 2021.

Co-funded by the Horizon 2020
Framework Programme of the European Union