Open Call 1

Self-Certifying Names for Named Data Networking
Deliverable 3: Experiment Results and Final Report

Authors	Nikos Fotiou, Yannis Thomas, Iakovos Pittaras, Vasilios A. Siris, George Xylomenos, George C. Polyzos Mobile Multimedia Laboratory Department of Informatics School of Information Sciences and Technology Athens University of Economics and Business, Greece {fotiou,thomasi,pittaras,vsiris,xgeorge,polyzos}@aueb.gr In collaboration with Prof. Christos Papadopoulos Department of Computer Science The University of Memphis, USA Christos.Papadopoulos@memphis.edu
Due Date	30/6/2021
Submission Date	30/6/2021
Assigned Reviewers	
Keywords	Decentralized Identifiers (DIDs), self-sovereign identities, Public Key Infrastructure (PKI), ICN, NDN, mutable, immutable, content authenticity

Deliverable 3: Part I

Analysis, results, and wider impact

The information contained in sections 1-11 will be used in part to update the NGI Atlantic's public deliverables (including the Experiment Catalogue on the website).

1. Abstract

The SCN4NDN project is verifying and evaluating a novel, self-sovereign, content authentication scheme applied to the Information-Centric Networking (ICN) paradigm. It is extending the Named Data Networking (NDN) ICN architecture, to include a self-managed digital signature scheme, based on the emerging standard of Decentralized Identifiers (DIDs). To this end, it leverages the NDN testbed to perform multiple experiments that measure the resilience of the proposed solution to security attacks, as well as its impact on key ICN functionalities, including caching, multicast, multisource, and multipath.

2. Project Vision

The SCN4NDN project is experimenting 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, the Internet of Things (IoT), and architectures focused on big data and/or cyber security. ICN's main goal is to enable fast and secure content dissemination by leveraging direct and intrinsic information identification; this allows supporting multicast, multipath, and caching, as well as novel trust mechanisms. In this project we are experimenting 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 are experimenting 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 is validating and evaluating 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 is specifically evaluating the use of DIDs as content name prefixes. These DIDs are randomly generated and do 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).

Based on feedback from the NDN community and from conference reviews, we extended our scheme to allow a DID to delegate control of content not only to a public key, but also to another DID. This allows the entity where control is delegated to periodically rotate its signing keys, without invalidating the delegation DID. Such an approach is applicable to Content Distribution Network (CDN) and IoT scenarios, as shown in the experiments section.

3. Details on participants (both EU and US)

This project is a joint effort between the Athens University of Economics and Business (AUEB) through its Research Center, managing research funding, and the Mobile Multimedia Laboratory, base of the researchers, and the University of Memphis (UofM).

The **AUEB team** includes the following members:

George Xylomenos (M), Professor
Role Project Manager
George Xylomenos is a Professor at AUEB and a member of the Mobile Multimedia Laboratory (MMlab). He received his Diploma in Informatics from AUEB and M.Sc. and Ph.D. degrees from the University of California, San Diego. From 2009 to 2018 he was the director of Network Operations at AUEB, and from 2017 to 2020 he was the deputy rector for Finances, Planning and Development at AUEB. His current research interests include future Internet and Information-Centric Networking architectures and improving the performance of digital media applications and protocols over mobile and wireless networks. He has participated in the groundbreaking EU-funded projects PSIRP, PURSUIT, and POINT that developed and implemented an ICN architecture and integrated it as an underlay for the Internet and recently in the SOFIE project on federated IoT with blockchains.

George C. Polyzos (M), Professor
Role Technical Manager
George C. Polyzos is leading the Mobile Multimedia Laboratory (MMlab) at AUEB, where he is Professor since 1999 and Director of the Graduate Program in Computer Science. Previously, he was Professor of Computer Science and Engineering at the University of California, San Diego (1988-1999), where he was co-director of the Computer Systems Laboratory, member of the Steering Committee of the Center for Wireless Communications

and Senior Fellow of the San Diego Supercomputer Center. He received his Diploma in Electrical Engineering from the National Technical University in Athens, Greece and his MASc in Electrical Engineering and PhD in Computer Science from the University of Toronto, Canada. Under his leadership the MMlab has participated in a series of research projects funded by the European Commission, the European Space Agency, and Greece that co-developed Publish-Subscribe Internetworking, an Information-Centric Networking architecture, with project PURSUIT receiving the Future Internet Award. His current research interests focus on the Internet-of-Things, security and privacy, Internet architecture and protocols, and wireless mobile multimedia networking. He is currently leading MMlab's participation in H2020 project INTERCONNECT and until very recently SOFIE, investigating Distributed Ledger and Interledger Technologies, including smart contracts, for IoT systems federation and smart homes and energy grid evolution, focusing on openness, security, privacy, and business incentives. He is on the editorial board of the *IEEE Transactions on Mobile Computing* and the *Journal of Reliable Intelligent Environments*. He was the founding chair of the Steering Committee of the ACM SIGCOMM conference on Information-Centric Networking and is now on the Steering Committee of the IFIP Wireless and Mobile Networking Conference. He has also been reviewer and panelist for the US NSF and the EC and reviewer for the California MICRO program, the Swiss National Science Foundation, the European Coordinated Research on Long-term Challenges, and the Greek General Secretariat of Research and Technology.

Vasilis A. Siris (M), Professor
Role He will contribute to DID-related activities and the integration with NDN
Vasilios Siris is Professor at AUEB and a member of the Mobile Multimedia Laboratory (MMlab). Previously he was an Assistant Professor at the Department of Computer Science of the University of Crete (2002-2008) and a Research Associate at the Institute of Computer Science, Foundation for Research and Technology - Hellas / FORTH (1993-2015). His current research interests include resource management and traffic control in wired and wireless networks, exploitation of Distributed Ledger Technologies (DLTs) and blockchains for trusted communication with constrained and mobile devices in the Internet of Things, and architecture of future mobile terrestrial/satellite and pervasive communication systems. He participated in the H2020 project SOFIE (Secure Open Federation for Internet Everywhere), where he was leading the architecture and framework evaluation.

Nikos Fotiou (M), PhD, Researcher
Role He will lead the DID-related tasks and security-related experiments design
Nikos Fotiou is a senior researcher with the MMLab/AUEB, where he completed his PhD in 2014 on "Information-Centric Networking: Security Requirements and Solutions." He received his MSc in Internetworking (2007) from the Royal Institute of Technology (KTH), ICT Dept., and the Diploma in ICSD Engineering (2005) from the University of the Aegean,

Dept. of Information and Communications Systems Engin. He was a key MMLab researcher on projects PSIRP, PURSUIT, POINT, ϕSAT, I-CAN and others, and until recently led the contributions of the MMLab in H2020 SOFIE on security solutions based on decentralized identifiers and verifiable credentials. His current research interests include user privacy, access control mechanisms, and content integrity and provenance verification. He has co-authored 18 journal publications and more than 44 conference and workshop papers, with more than 2700 citations.

Yiannis Thomas (M), PhD, Researcher

Role He will lead NDN integration and network-related experiment design

Yannis Thomas is a postdoctoral researcher at the Department of Informatics of AUEB and member of the Mobile Multimedia Laboratory (MMlab). He received his Diploma in Informatics (2009), his MS in Information Systems (2012) and his PhD in Computer Science (2018) from the Department of Informatics of AUEB, Greece. He has contributed to a series of research projects funded by the European Commission (EC), the ESA, and Greece, such as the EC-funded projects "Publish Subscribe Internet Technology", "iP Over IcN-the betTer IP" and "Secure Open Federation for Internet Everywhere", and the ESA-funded project "The role of satellite in the future Internet". His research interests include multiflow transport, congestion and flow control in the current Internet and Next Generation Networks, including Mobile Adhoc NETworks (MANETs) and heterogeneous Satellite-Terrestrial networks. He has coauthored more than 20 refereed articles in international journals, books, and conferences, including top-tier venues such as the IEEE INFOCOM and *IEEE/ACM Transactions on Networking*. He has also been reviewer for international journals, such as the *ACM Transactions on Mobile Computing*, Elsevier *Computer Networks* and *Computer Communications*.

Iakovos Pittaras (M), PhD student

Role He will be responsible for the execution of experiments

Iakovos Pittaras is a PhD candidate at the Department of Informatics of the Athens University of Economics and Business (AUEB) and member of the Mobile Multimedia Laboratory (MMlab). The title of his PhD thesis is "Secure Interoperability for Internet of Things Data and Actuation". Previously, he received his Diploma in Informatics from AUEB (2017), where he was admitted eleventh in a class of 180 students. He obtained his M.Sc. in Computer Science (GPA: 9.25/10, first) from AUEB (2019). During his postgraduate studies, he received a scholarship, funded by Omilia Ltd., for being first in the class. The title of his M.Sc. thesis was "Interacting with the Web of Things using Blockchains", which received a grade of 10/10. Until recently he participated as a researcher in the H2020-SOFIE project.

The **UofM team** is composed of the following member:

Christos Papadopoulos (M), Professor
Role He will assist in designing experiments for the NDN testbed and the UofM servers
Christos Papadopoulos received his PhD in 1999 from Washington University in St. Louis, MO. Before joining the Computer Science Department at the University of Memphis he was an assistant professor at the University of Southern California and a professor at Colorado State University. His research interests include network and cyber-physical systems security, global Internet measurements, information-centric networks and smart and autonomous systems. From 2018-20 he was a program manager at the Department of Homeland Security (DHS) Science and Technology directorate (S&T), where he managed projects on cyber-physical systems security, including industrial control systems, law enforcement and automotive cybersecurity. Dr. Papadopoulos was involved in research projects in excess of $10M and served as a principal investigator on projects funded by DHS, DARPA and the NSF. Recent projects address Internet-wide security problems such as Distributed Denial of Service (DDoS) attacks and detecting evasive Internet bots, analysis of global Internet routing, and applications of future Internet architectures in scientific domains such as climate and high-energy Physics through Named Data Networking (NDN). He is a senior member of the IEEE and has served on numerous ACM and IEEE conference program committees.

4. Results

In this section we describe the experiments performed in the NDN testbed as part of the project. The baseline and first two scenarios were presented in D2, the remainder is first reported in this deliverable. We also describe the modifications to our original design based on feedback from the NDN community and the additional options (delegation to DIDs) implemented during the (authorized) extension of the project.

Our first set of experiments aimed to validate our concept. As a baseline scenario we performed the following experiment. We implemented a Producer script installed in an endpoint VM attached to the testbed node at Colorado State University (CSU), and a Consumer script installed in an endpoint VM at a MMlab testbed node. The Producer script advertises a content item name prefix. Initially, the Consumer script sends an interest message for that prefix (see Figure 1 for the sequence of events). After 0.5 seconds a packet arrives that includes the content item's metadata. The metadata indicate that the content item includes 10 chunks. Then, the Consumer script requests chunks one by one, i.e., it requests the second chunk after the first chunk has arrived, and so forth. After 4.39 seconds all chunks have been received.

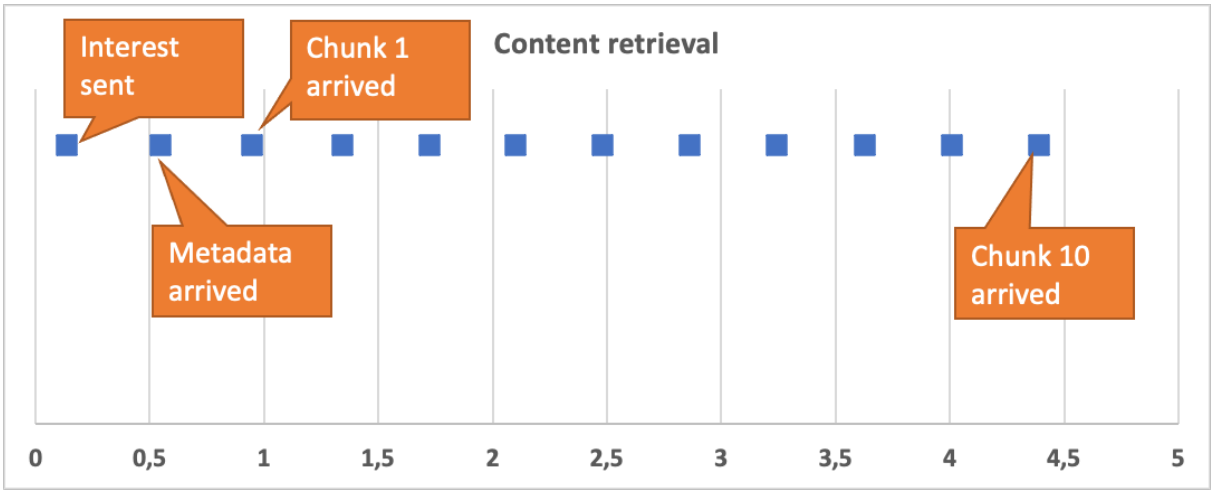

Figure 1: Content retrieval events in baseline scenario. The x axis shows time in seconds.

Scenario 1: Accelerated content retrieval using multisource transfer

With this experiment we validated our approach of using multiple sources (multisource). In particular, we validated that it is possible to receive a content item simultaneously from two different sources (with chunks coming from both sources). For this experiment we extended our baseline scenario to include another producer, located at an endpoint VM attached to the testbed node of the University of Memphis (UofM). This producer advertises the same content item using a different name prefix; however, the two content items are linked through their metadata.

As in the baseline scenario, the Consumer script sends an interest message for the prefix advertised from CSU (see Figure 2 for the sequence of events). After 0.5 seconds a packet arrives that includes the content item's metadata. The metadata indicate that the content item includes 10 chunks, as well as that the item has an alternative name. Then, the Consumer script requests half of the chunks using the original name, and at the same time it requests the rest of the chunks using the alternative name. After 2.43 seconds all chunks have been received.

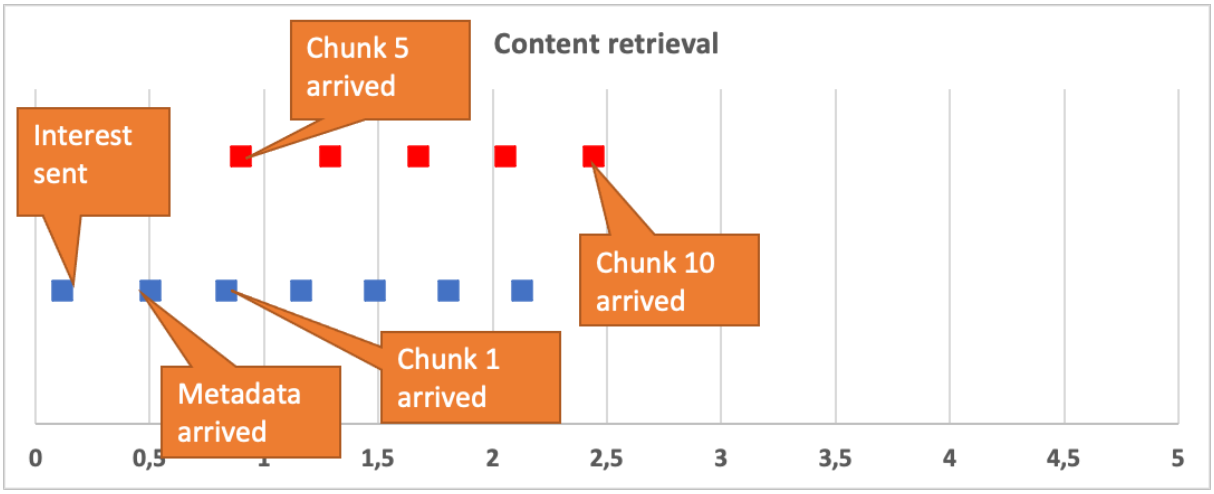

Figure 2: Content retrieval events with multisource. The red chunks are received from a second source. The x axis shows time in seconds.

Scenario 2: Recovery from network failure using multisource

With this experiment we validated our approach of using multisource to recover from network failures at the application layer. The setup of this experiment is the same as in scenario 1. However, in this experiment the Consumer script uses the alternative name as a backup solution. The Producer script attached to CSU is configured to stop responding to new interests after transmitting the 5th chunk. At this point, the corresponding interest "times out" and the Consumer requests the rest of the chunks using the alternative name. After 5.49 seconds all chunks have been received, as shown in Figure 3.

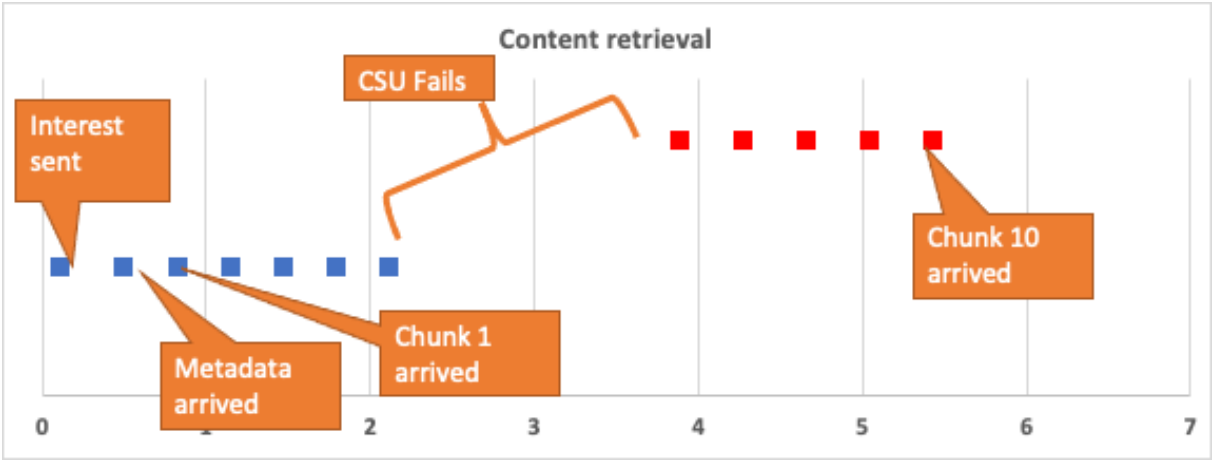

Figure 3: Use of multisource to recover from failures. The red chunks are received from a second source. The x axis shows time in seconds.

Scenario 3: Application layer multipath

This scenario experiments with the use of multiple paths (multipath) which are exploited simultaneously for receiving an item. For this experiment, and similarly to Scenario 1, two alternative prefixes are used for the same item. Nevertheless, in this scenario both prefixes are advertised by the same producer. Therefore, the producer script advertises the two

prefixes (prefix1 and prefix2) from a testbed node located in UMemphis. The Consumer script creates two faces: face1 is used for connecting to a MMlab testbed node and face2 is used for connecting to a CSU testbed node. Then, the Consumer script registers a route towards prefix1 through face1, and a route towards prefix2 through face2. The Consumer sends an interest message for prefix1. The first packet that arrives includes the content item's metadata. These metadata indicate that the content item includes 10 chunks, as well as that the item has an alternative name (i.e., the prefix2-based name). Then, the Consumer script requests half of the chunks using the original name, and at the same time it requests the rest of the chunks using the alternative name: requests for the alternative name are routed through face2, therefore they reach the publisher through an alternative path. With this scenario we obtain the same results as in the case of multisource (Figure 2), using however a single source.

Scenario 4: Security attack

In this scenario we consider the setup of Figure 4. In particular, there are three Producer scripts, one attached to a CSU testbed node (producer *P*) and two attached to MMlab testbed nodes (Producers *A* and *B*). All producers are part of a service identified by a did:self DID. Producers *A* and *B* provide "IoT measurements", published using the service DID as a prefix, whereas producer *P* publishes a "list of producers" also using the service DID as a prefix. Initially, the Consumer script, which is attached to a testbed node of UMemphis, sends an interest for an item advertised by *P* and receives the list of producers, i.e., it learns about producers *A* and *B*. Then, it periodically sends two interest packets, one for the measurements of producer *A*, and another for the measurements of producer *B*. We emulate an attack where an attacker gains access to the private key that producer *B* uses for advertisements (but not the private key that corresponds to the DID and used for signing Data) as follows: we initiate a new Producer script attached to the same node as the Consumer script; this Producer script (*Attacker*) is configured with the private key of B and sends an advertisement to the network. Since the advertisement is signed using a valid key, it is accepted. Moreover, since the attacker is attached to the same node as the Consumer script, the next interest sent by the Consumer script is satisfied by the *Attacker*. Meanwhile, we assume that the breach has been detected and producer *B* is renamed to producer *C,* and is configured with a new private key (used for signing NDN advertisements). Upon receiving the data packet from the attacker, the Consumer script fails to validate the signature (since this signature should be generated using the private key that corresponds to the DID, and the attacker does not have access to it). As a result, the Consumer script sends a new interest for the item published by publisher P, and it receives the updated list of producers.

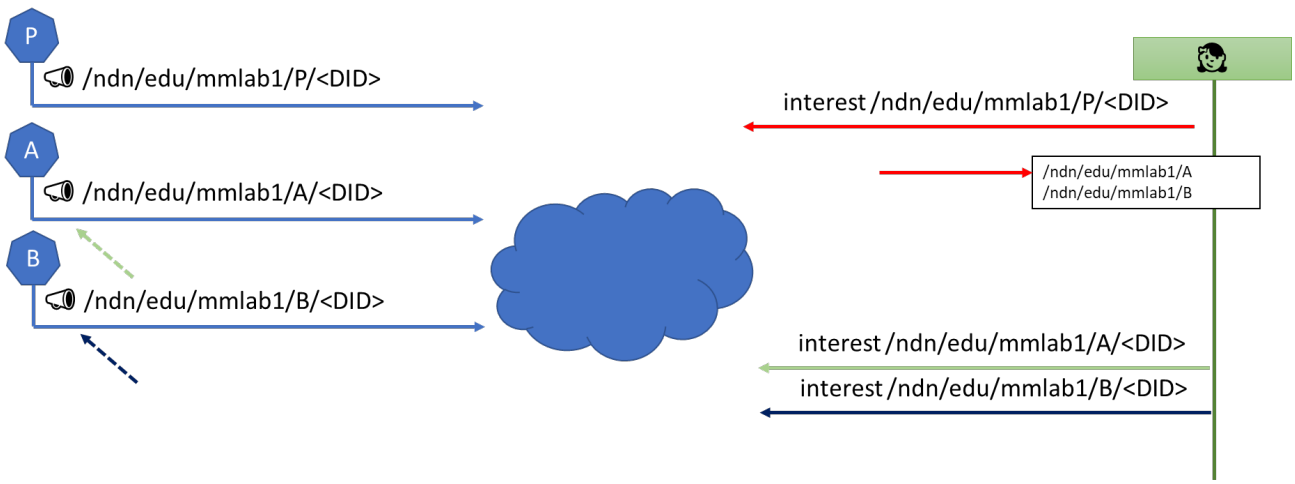

Figure 4: Scenario 4 topology.

Update of the design based on the feedback of the NDN community

In our initial design we were using did:self DIDs as content names. Nevertheless, after interacting with the NDN community it became clear that we can increase the compatibility of our approach with NDN semantics, without losing any of its desired properties, if we use did:self DIDs as content name *prefixes*.

In our updated design, we consider content producers that are responsible for advertising and publishing content items, and content consumers that express interest in content items. Content item names are hierarchical, and we consider two items with the same name as the same item. We also consider "protected" namespaces rooted in a DID, that is, the DID is the prefix of the namespace. These namespaces are owned by the corresponding DID owner.

With this update, a new property is added to the metadata of a content item, which indicates the item's name. Our experiments are not affected, as it is only the interpretation of the DIDs that changes, and not the procedure for asking for and receiving content.

Exploration of new constructions during the extension

Our project was extended by 1.75 months, and this gave us the opportunity to explore an additional construction that facilitates "content storage delegation." We performed the following experiment. A producer *P* owns the prefix */<DID>* and wishes to allow producer *A* to publish a content item under */<DID>*. Producer A owns a did:self DID, i.e., *<DID A>*, and in the corresponding DID document has defined a public key. Producer *P* enables delegation by a generating a DID document for *<DID>* that includes in the "assertion" property the *identifier* of the public key defined in the DID document of *<DID A>* (see also Figure 5) and by sending this document to producer *A*. The advantage of this approach is that producer *A* can change its public key without having to receive a new DID document from *P.* In a certificate-based system, this is the equivalent of being able to change keys without having to receive a new certificate.

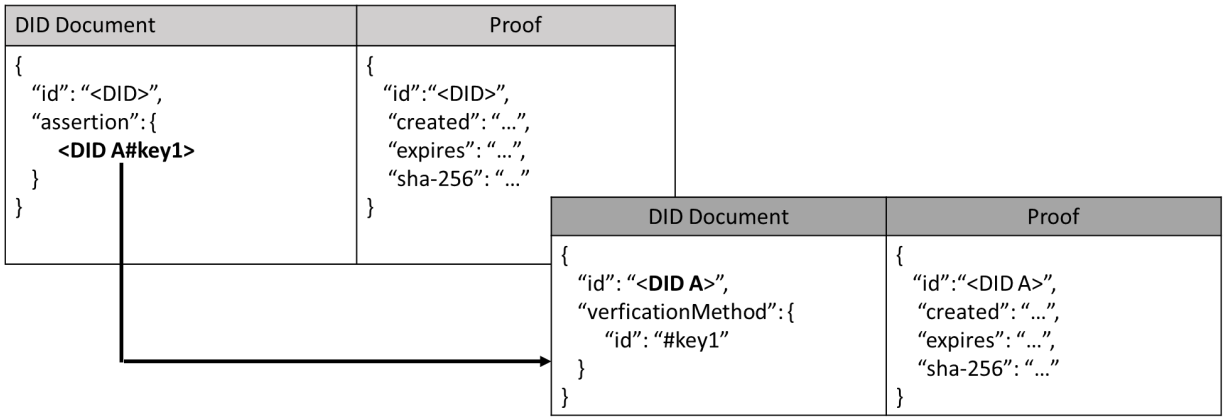

Figure 5: Item metadata based on the new construction.

5. Discussion and Analysis of Results

Our solution supports multiple communication paradigms that improve the performance and the resilience to failures of NDN, as shown in the previous section: data can be downloaded simultaneously from multiple sources or via multiple paths; alternatively, multiple sources and/or paths can be used in failover mode to improve fault tolerance. At the same time, it achieves significant security properties that have the potential to advance NDN security.

Security properties achieved

Assuming that a consumer knows the DID used as the prefix of a content item name, our solution achieves the following security goals in the presence of an *active* adversary:

- The integrity of transmitted items is protected: a digest of the content item's data is recorded in the metadata header. Since metadata also includes the item's name, an attacker cannot replace the transmitted item with another, valid item.
- The authenticity of transmitted items is protected. The authenticity of an item, i.e., the verification of the "binding" between the transmitted item and the item name included in the corresponding interest message, is achieved by including the item's name in the metadata. If the item's name was not included, an attacker could send an interest for a different item that has the same prefix as a legitimate request and then, replace a legitimate response with the response it received to his request.
- Authorized producers can easily rotate keys. A producer can replace its signing key with a new one; as long as the new key has the same key identifier as the replaced one, no more actions are required, e.g., the producer does not have to receive a new DID document for the DID user as the name prefix. In a certificate-based solution on the other hand, an endpoint would have to receive a new certificate from the service provider. This feature is particularly useful in use cases such as IoT systems where the same service is provided by different devices, which are periodically "rotated" (e.g., because their battery is drained).

Improvements to NDN security

In our solution, content name prefixes are used as "trust anchors," therefore there is an explicit binding between a content item name and the corresponding trust anchor. NDN on the other hand relies on Trusted Third Parties (TTP) or on endpoint configuration using trust "schemas" to provide this binding [13].

A TTP-based solution requires from all participating entities to agree on a (set of) TTPs, which in many cases may be cumbersome, or even involve non-negligible monetary cost. This problem becomes even more significant if security verifications should be performed by in-network nodes. Additionally, TTPs are a significant security threat since, in the general case, a TTP can generate certificates for any name prefix hence, a malicious TTP, or an attacker that has access to the private key of a TTP, can generate an arbitrary number of valid certificates. Finally, in TTP-based solutions it is hard to manage cases where a TTP must rotate its keys. If this rotation means that old keys are revoked, then all issued certificated must be re-issued. Additionally, all endpoints must be configured with the new keys of the TTP, which may require significant effort, especially if in-network devices must be reconfigured.

Using trust schemas and self-generated certificates, security properties similar to our solution can be achieved. However, a schema-based solution requires that either all content name prefixes are known, in order for the appropriate rules to be configured, or that a TTP can be used as a trust anchor for prefixes not covered by a rule. Therefore, in scenarios where content names are "discovered" (e.g., in Web-like, or search applications) a schema-based solution behaves similarly to a TTP-based solution.

Integration with the NDN software

Our solution can be easily integrated to the *Named Data Networking Forwarding Daemon* (NFD). In particular, our solution can be implemented as a module of *NFD's core forwarding functionality*. Upon receiving a Data packet "flagged" by our solution (e.g., published using a pre-defined prefix), NFD may parse the metadata from the packet payload and check if it is valid for the content name of the Data packet; if not, then it should discard it and send a *Negative Acknowledgement* (NACK) packet to the consumer in order to notify it that the Data packet was not valid, as well as to consume the *Pending Interest* state on the remaining on-path routers.

NFD can use our existing Python3-based implementation for performing metadata validation, using the *popen* (process open) command offered by the Linux kernel, which allows a C++ process to invoke a Python process and read the output of the latter through a pipe stream. This approach allows a seamless integration with the NFD code, in that only minor modifications are required.

6. Present and Foreseen TRL

Our foreseen TRL level for our final software was 4 (Technology validated in lab). Nevertheless, our software exceeded this level, since it was also validated with the Inter-

Planetary File System (IPFS), a relevant production environment. Hence our final software is in TRL 5 (Technology validated in relevant environment).

7. Exploitation, Dissemination and Communication Status

The project outcomes are expected to be of interest to both the "Decentralized Credentials" and "Information-Centric Networking" communities. For this reason, our exploitation, dissemination, and communication plan is targeting both communities.

We created a project website, https://mm.aueb.gr/scn4ndn, where all information about the project is made publicly available. In addition to a project description, the website hosts two video presentations (a short interview where Dr. Nikos Fotiou presents the project, and the NGIAtlantic.eu workshop at the NGI Policy Summit 2020 "Privacy and trust: trends in experiments in EU-US research and innovation", where the SCN4NDN project was presented). We also participated in the NGIatlantic.eu Show your Talent through our Twinning Lab: Third Open Call proposers webinar. Following the Interviews with NGIAtlantic researchers, Nikos Fotiou was also interviewed about the project by Radio Judaica in Belgium. Moreover, the project was presented to Vienna Digital Identity meetup[1] that had 53 attendees. The project was also presented to Dr. Dirk Trossen, Chief Network Architecture Research Engineer at Huawei, and possible joint research directions were discussed. Finally, the project website provides pointers to the software we are developing, which is freely available in GitHub (https://github.com/mmlab-aueb/scn4ndn).

Project members participate in the W3C Credentials Community Group (W3C-CCG) [6], as well as in the Decentralized Identity Foundation (DIF) [7]. All DID-related outcomes will be presented to both groups. Furthermore, project members participate in IETF's Information-Centric Networking research group (ICNRG) [8]. Project outcomes will also be made available and presented, if given the opportunity, to this group. Finally, we presented our proposal in the NFD development team telco, which is the main forum for NDN developers, on April 15th, gathering important feedback from the participants regarding the integration of our work with the NDN testbed codebase.

The project submitted in May a paper to the ACM ICN 2021 conference [12], the flagship ICN event, on the use of DIDs for content items in NDN. A paper describing how our scheme can be used to secure NDN routing was accepted and presented to the SARNET workshop of the IEEE HPSR Conference [10], while another paper proposing the extension of our scheme to secure mutable content in IPFS was accepted and presented to the DI2F workshop of the IFIP Networking Conference [11]. A talk proposal related to our project has also been submitted to Future of PI, a EuroS&P 2021 workshop. Finally, the project team is participating in the NGI TETRA #2 Summer bootcamp.

[1] https://www.meetup.com/Vienna-Digital-Identity-Meetup/events/276496501/

8. Impacts

With respect to the NGI initiative, our 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 ICN conference, and their participation in international working groups, such as IRTF's ICNRG, have already established a fruitful relationship that guarantees a successful collaboration. So far, we have collaborated with researchers from the Washington University in St. Louis, the Arizona State University, and the Tennessee Tech University (TTU), in order to extend the NDN testbed and co-operate in bug fixes and improvements of existing software. We made a git pull request to fix a bug in mini-ndn (https://github.com/named-data/mini-ndn), a package allowing NDN to run in the mininet emulator. Finally, we were asked by the NDN testbed maintainers to document our setup in a technical report, since the existing documentation has become obsolete due to changes in the testbed.

- Reinforced collaboration and increased synergies between the Next Generation Internet and the Tomorrow's Internet programmes.

Our 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, our DID-based content authentication mechanism can be applied in other networking and application contexts, such as the emerging Inter-Planetary File System (IPFS) [9], network routing advertisements and even IoT scenarios using CoAP; our work in this area was published in two conference papers. We are monitoring such discussions in the NDN list, where many people involved with Tomorrow' Internet are involved, and recently participated in a call on the NFD development list related to Verifiable Credentials (VCs). We are also discussing with Assistant Professor Susmit Shannigrahi from TTU the possibility of writing a joint paper on the evolution of NDN names with DIDs. We are also discussing with Dr. Dirk Trossen, Chief Network Architecture Research Engineer at Huawei, possible joint research directions.

- Developing interoperable solutions and joint demonstrators, contributions to standards

Our 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 technology is mainly developed under the umbrella of the IETF/IRTF. Both 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. Therefore, in addition to releasing our software, we presented our approach to the NFD developers during their regular teleconference, where we gathered important feedback.

- An EU - US ecosystem of top researchers, hi-tech start-ups/SMEs and Internet-related communities collaborating on the evolution of the Internet

We envision 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. We believe that research teams on both continents will benefit from this complementary partnership. The collaboration so far, which arose out of the needs to setup our experiments and update/fix existing tools, is quite promising.

9. Conclusion and Future Work

With the SCN4NDN project we experimented with the merger of Information-Centric Networking (ICN) and Decentralized Identifiers (DIDs). Our experiments not only allowed us to validate our hypothesis, i.e., DIDs enable advance communication paradigms improving at the same time security, but also they helped us improve our design. Through the interaction with the NDN community and DID researchers, we managed to shape a solution that on one hand can be easily integrated to NDN, and on the other hand implements a DID method with intriguing properties and many applications. For these reasons, we believe that SCN4NDN is only a first step towards new, exiting endeavors.

ICN-related future work

SCN4NDN is just a first step and many promising potentials of this approach still remain unexplored. In particular:

- SCN4NDN leveraged did:self only at the application layer. However, core network components can benefit from this approach, especially given their information-centric operation. For example, a serious attack against NDN (and ICN in general) is “cache poisoning” where an attacker fills caches with fake items. Using our approach, such an attack can be prevented.
- SCN4NDN tries to solve the problem of security key breaches only in “closed systems”, i.e., systems where “everybody talks with everybody.” However, a more general solution is required. SCN4NDN did not provide a solution for open systems since it did not want to make any assumptions about the ICN API of NDN. We can now state that by taking advantage of the NDN ICN API and its support for “versions” we can mitigate security incidents in open systems. Nevertheless, this possibility is yet unexplored.

- SCN4NDN uses public keys as content name prefixes. Although this simplifies the security operations it still requires a secure way for disseminating these public keys. Therefore, using the SCN4NDN approach there are cases that still require a PKI system. Certificateless Public Key Cryptography (CPKC) can be a solution to this problem. CPKC improves Identity-Based Encryption by solving its inherent key escrow problem (i.e., the fact that there is an entity—the Key Generator—that knows all secret keys).
- The SCN4NDN approach requires access to the whole content in order to perform content verification authentication. However, there are cases where it is desirable to hide parts of the content, e.g., because they include sensitive information. Consider for example the case of an electronic identity and the scenario where a service is interested in learning only the age of the user: a solution that allows hiding all fields but the age, but still makes it possible to verify the authenticity of the identity would increase significantly the privacy of the system. This problem could be solved by integrating Zero-Knowledge Proofs to our approach.

New research directions

One of our goals from the beginning of the SCN4NDN project was to design a solution with future directions in mind. Indeed, we had the chance to use the outcomes of SCN4NDN in other domains and experiment with the application of our solution for supporting mutable items in the Inter-Planetary File System (IPFS), for securing group membership in CoAP group communication, as well for providing security for semantic routing protocols. Our team will pursue follow-up activities towards these directions.

10. References (optional)

[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 at https://w3c-ccg.github.io/did-primer/

[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.

[4] Named Data Networking, “Named Data Networking Forwarding Daemon,” available at https://github.com/named-data/NFD

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

[6] W3C, “W3C Credentials Community Group Home Page,” available at https://www.w3.org/community/credentials/

[7] DIF, “Decentralized Identity Foundation Home Page,” available at https://identity.foundation

[8] IETF, “Information-Centric Networking Research Group (icnrg) charter,” available at https://datatracker.ietf.org/rg/icnrg/about/

[9] IPFS, “The InterPlanetary File System Home Page,” available at https://ipfs.io

[10] N. Fotiou, Y. Thomas, V. A. Siris, G. Xylomenos, G.C. Polyzos. “Securing Named Data Networking routing using Decentralized Identifiers,” Semantic Addressing and Routing for Future Networks (SARNET) workshop (part of the IEEE International Conference on High Performance Switching and Routing), 2021.

[11] N. Fotiou, V.A. Siris, G.C. Polyzos, “Enabling self-verifiable mutable content items in IPFS using Decentralized Identifiers,” Decentralizing the Internet with IPFS and Filecoin (DI2F) workshop (part of IFIP Networking Conference), 2021.

[12] N. Fotiou, Y. Thomas, V.A. Siris, G. Xylomenos, G.C. Polyzos, “Self-certifying names using Decentralized Identifiers,” submitted to the ACM Conference on Information-Centric Networking (ICN), 2021.

[13] Y. Yu., A. Afanasyev, D. Clark, K. claffy, V. Jacobson, L. Zhang, “Schematizing trust in named data networking,” in the proceedings of the 2nd ACM Conference on Information-Centric Networking (New York, NY, USA, 2015), ACM-ICN ’15

11. Glossary

AUEB	Athens University of Economics and Business (Coordinating partner)
CPKC	Certificateless Public Key Cryptography
CSU	Colorado State University
DID	Decentralized Identifier
DIF	Decentralized Identity Foundation
ICN	Information-Centric Networking
ICNRG	Information-Centric Networking Research Group (of the IRTF)
IPFS	InterPlanetary File System
JSON	JavaScript Object Notation
JWS	JSON Web Signature
NDN	Named Data Networking
NFD	Named Data Networking Forwarding Daemon
TTU	Tennessee Tech University
UofM	University of Memphis (US-based partner)
W3C-CCG	W3C Credentials Community Group

Deliverable 3: Part II

Financial and cost information

This part is to be treated as a consortium confidential deliverable, and access is restricted to consortium partners and EU commission operatives.

NGI ATLANTIC.EU

12. Workplan Progress and Travel Details

Note 1: In the project proposal the workplan was broken down into Activities, subdivided into tasks. We have used the term "Work Package" instead of "Activity" below. Since UofM is not compensated as part of this project, it is listed with no effort in the following tables.

Note 2: An extension was requested and approved, extending the project to 7.75 months, therefore the timeline and effort was extended for some work packages (marked with * in the tables below).

Work Package 1: Project management, dissemination and exploitation

WP 1, Task 1	Start month	1	Duration	7.75*	Total PM	**0.6***
Title	*Project Management*					
Partners involved				Person months		
Partner 1	*AUEB*			0.6		
Partner 2	*UofM*					
Goal: *Ensure the project follows its planning and that deliverables are on time.*						
Activities Description *During its entire duration, the project completed all tasks according to plan, including a 1.75 month extension to complete additional work. Three deliverables were produced (D1, D2 and D3) and delivered on time.*						

WP 1, Task 2	Start month	2	Duration	6.75*	Total PM	**0.9***
Title	*Project dissemination and exploitation*					
Partners involved				Person months		
Partner 1	*AUEB*			0.9		
Partner 2	*UofM*					
Goal: *Disseminate project results via publications and other channels*						
Activities Description *The project first focused on constructing a website with all information related to the project, especially in video form, and pursued all publicity opportunities offered by NGIAtlantic, including blog posts and radio interviews. Links were established with US-based NDN testbed partners to collaborate on experiments and possible publications. The project participated in various groups to promote our work. 3 publications were prepared during the project (2 have been accepted).*						

Work Package 2: Integration of the proposed scheme to NDN

WP 2, Task 1	Start month	1	Duration	1	Total PM	**2**
Title	*Design and implementation of routing, based on Named Data Link State Routing Protocol*					
Partners involved				**Person months**		
Partner 1	*AUEB*			2		
Partner 2	*UofM*					

Goal:

Extend routing protocols and applications to incorporate the SCN4NDN approach.

Activities Description

As part of our testbed setup, we verified that content can be pulled from multiple sources by appropriately extending the end-user applications to recognize content replicas; using DIDs we can verify the links between content items published under different prefixes. Moreover, we automated (i) the process of registering "special purpose" prefixes, required by the Named Data Link State Routing Protocol to indicate "default gateways," (ii) the registration of static routes required by our experiments, (iii) the creation of virtual links (aka "faces"), and (iv) the advertisement and publishing of the appropriate certificates required for verifying the signatures included in the corresponding prefix advertisements.

WP 2, Task 2	Start month	2	Duration	1	Total PM	**2**
Title	*Design and implementation of signalling protocols for multipath/multisource*					
Partners involved				**Person months**		
Partner 1	*AUEB*			2		
Partner 2	*UofM*					

Goal:

Extend libraries and applications to incorporate the SCN4NDN approach.

Activities Description

We extended the ndn-python library and our applications to use metadata to link content published under different prefixes, using DIDs to verify that the content is valid. This allows using multipath and multisource transport in our experiments, without changes to the NDN testbed nodes or NFD.

WP 2, Task 3	Start month	2	Duration	6.75*	Total PM	5*
Title	*Implementation using NDN-libraries*					
Partners involved				Person months		
Partner 1	*AUEB*			5		
Partner 2	*UofM*					
Goal: *Implement the extensions designed by the previous tasks to NDN libraries*						
Activities Description *Using the ndn-python library as a base, we wrote applications that test the various scenarios envisioned, with extensions to handle multisource/multipath signalling and routing. During the extension we added delegation to DIDs and tested alternative methods of integrating our approach to the ndn codebase. We also built a semi-automated system for test execution on the NDN testbed. We have also contributed bug fixes to existing libraries (e.g., mini-ndn). Our software is available here: https://github.com/mmlab-aueb/scn4ndn*						

Work Package 3: Experiment design

WP 3, Task 1	Start month	3	Duration	1	Total PM	1
Title	*Test Preparation*					
Partners involved				Person months		
Partner 1	*AUEB*			1		
Partner 2	*UofM*					
Goal: *Prepare the experimental testbed and design detailed experiments*						
Activities Description *We have completed the testbed setup by integrating our nodes into the NDN-testbed and adding external systems serving as publishers and subscribers. Our testbed nodes are configured to be part of the NDN certification system (https://ndncert.named-data.net) and we have generated and installed the appropriate certificates required by our experiments. We have produced a testing plan detailing the experiments that we are going to conduct (available at: https://github.com/mmlab-aueb/scn4ndn).*						

Work Package 4: Experiment execution and results evaluation

Note 1: Due to the need to verify the correct operation of multipath and multisource transport using our scheme, we exchanged Tasks 4.1 and 4.2 in the work plan, that is, Task 4.2 started on M3 and Task 4.1 started on M4. This change is reflected in the following entries.

Note 2: After the extension was approved, all activities were extended to the end of the project, with additional effort in each activity.

WP 4, Task 1	Start month	4	Duration	4.75*	Total PM	**5.47***
Title	*Security related experiments*					
Partners involved				**Person months**		
Partner 1	*AUEB*			5.47		
Partner 2	*UofM*					

Goal:

Evaluate the security guarantees of our scheme

Activities Description

The task started in M4 and considered various security scenarios (network attacks, key breaches, key revocation, need to rotate keys), how they can be addressed, and what their performance impact is; during the extension we specifically focused on the impact of delegation to DIDs on security.

WP 4, Task 2	Start month	3	Duration	5.75*	Total PM	**6***
Title	*Experiments related to support for aliases*					
Partners involved				**Person months**		
Partner 1	*AUEB*			6		
Partner 2	*UofM*					

Goal:

Evaluate the applicability and uses of aliases under our scheme

Activities Description

This task, started one month earlier (M3), compared the resilience and performance of single source and multisource transport; this is a continuation from the test setup and the validation of the testing plan, but with performance measurements. During the extension, we also tested the impact of delegation to DIDs.

WP 4, Task 3	Start month	6	Duration	2.75*	Total PM	**4***
Title	*Experiments related to support for mutability*					
Partners involved				**Person months**		
Partner 1	*AUEB*			4		
Partner 2	*UofM*					

Goal:

Evaluate the applicability of our scheme to mutable content

Activities Description

The task started in M6 and focused on the support for mutable items and their integration with the NDN philosophy (where changed items need to also change name). Following feedback from the NDN community we modified our approach to secure the root namespace of content, as opposed to individual names. We also applied our scheme to IPFS which also assumes immutable content.

No trips were taken during the reporting period, due to restrictions related to the COVID-19 pandemic. Communication between partners and third parties has been made exclusively via video conferencing tools. In particular, we held a bi-weekly call between UofM and AUEB researchers and a regular weekly call among all AUEB researchers, in addition to informal interactions. We participated in NGI events also via video conferencing tools.

13. Funds Utilisation Report

Cost Title	Amount	Description
Personnel cost(s)	127,676.92 €	A total of 26.97 PMs were spent, divided into tasks and personnel according to the following table.
Other Direct Costs – Travel only	0.00 €	No trips were made due to COVID-19 related restrictions. The allocated costs were moved to personnel costs for the extension.
Total Direct Costs	127,676.92 €	
Indirect Costs	31,919.23 €	Charged at 25% on direct costs
Total Costs	159,596.15 €	The extra amount (on top of the 150,000.00 € awarded) will be covered by AUEB funds.
Received Amount	58,907.29 €	
Remaining Amount	91,092.71 €	***Claimed*** remaining amount

The following table shows in detail how the effort was distributed across tasks and what each researcher contributed to each task; the final column shows total effort per WP/activity, the final row shows total effort per researcher.

	WP/Activity	G. Polyzos	G. Xylomenos	V. Siris	N. Fotiou	Y. Thomas	I. Pittaras	Total
1	**Project management, dissemination and exploitation**							
1.1	Project Management	0.30	0.30					**0.60**
1.2	Project Dissemination & Exploitation	0.18	0.20	0.10	0.30	0.12		**0.90**
2	**Integration of the proposed scheme to NDN**							
2.1	Design and implementation of routing, based on NDLSRP	0.99		1.01				**2.00**

2.2	Design and implementation of signalling protocols for multipath/multisource		1.15	0.37		0.48		**2.00**
2.3	Implementation using NDN-libraries				1.45	0.80	2.75	**5.00**
3	**Experiment design**							
3.1	Test preparation				0.50	0.20	0.30	**1.00**
4	**Experiment execution and results evaluation**							
4.1	Security related experiments	1.00	0.83	0.72	1.35		1.57	**5.47**
4.2	Experiments related to support for aliases	0.40	0.50	0.50	1.60	1.00	2.00	**6.00**
4.3	Experiments related to support for mutability	0.30	0.50	1.00	0.92	0.78	0.50	**4.00**
	Total	**3.17**	**3.48**	**3.70**	**6.12**	**3.38**	**7.12**	**26.97**

On behalf of Athens University of Economics, I, George C. Polyzos, confirm that this funds utilisation report is in accordance with the contract already in place between the Athens University of Economics and Business—Research Center and the Waterford Institute of Technology under financial support to third parties from Article 15 of Grant Agreement number 871582 — NGIatlantic.eu. I confirm that this report also includes all the expenditures (limited to PM and travel) incurred by all EU partners in this project and adhere to all instructions contained in H2020 Annotated Model Grant Agreement[2]. These are referenced in sections 3 and 5 of the contract. I also confirm that any applicable VAT or tax payments on the amount due to the Grant Recipient shall be fully borne by the Grant Recipient.

Signed for and on behalf of

Athens University of Economics and Business

Director, Mobile Multimedia laboratory

Digitally signed by GEORGE POLYZOS
Date: 2021.06.30 10:10:03 +03'00'

……

Prof. George C. Polyzos

Patision 76, Athens 10434, Greece

NGI ATLANTIC.EU

[2] http://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/amga/h2020-amga_en.pdf