ICT-317756

TRILOGY2

Trilogy2: Building the Liquid Net

Specific Targeted Research Project

FP7 ICT Objective 1.1 The Network of the Future

D4.1 - Public Record of Trilogy 2 Academic Workshop

Due date of deliverable: 31 December 2013

Actual submission date: 4 January 2014

Start date of project	1 January 2013
Duration	36 months
Lead contractor for this deliverable	UC3M
Version	v1.0, 3 January 2014
Confidentiality status	Public

Abstract

This deliverable presents the dissemination activities performed during the first year of the Trilogy 2 project. It includes standardization activities (ETSI, IETF), the organization of two academic workshops (IEEE Infocom Global Internet Symposium and ACM Conext Hot MiddleBox workshop), the publication of academic papers and other dissemination activities.

Target Audience

The target audience is anyone interested in the activities performed by the Trilogy 2 project to disseminate the project results.

Impressum

Full project title	TRILOGY2: Building the Liquid Net
Title of the workpackage	WP4 Dissemination, Standardisation and Exploitation
Editor	Marcelo Bagnulo Braun, UC3M
Project Co-ordinator	Marcelo Bagnulo Braun, UC3M
Copyright notice	© 2014 Participants in project TRILOGY2

1 Executive Summary

This deliverable describes the dissemination activities performed by the Trilogy 2 project during its first year. These activities include the organization of two academic workshops (Global Internet Symposium collocated with IEEE Infocom and the HotMiddleBox workshop collocated with ACM context), the participation of standardization activities (including the ETSI NFV and the IETF), the publication of several academic papers and some other dissemination activities.

List of Authors

Authors	List of document authors
	Pedro Andrés Aranda Gutiérrez
	Diego López
	Fco. Javier Ramón Salguero
	Bob Briscoe
	Marcelo Bagnulo
Participants	List of author institutions
	Telefónica, Investigación y Desarrollo
	BT
	UC3M
Work-package	WP4 Dissemination, Standardisation and Exploitation
Security	PUBLIC (PU))
Nature	R
Version	1.0
Total number of pages	35

Contents

List of Figures

List of Tables

Acronyms

BCP Best Current Practice.

ConEx Congestion Exposure.

COTS Commercial Off-The-Shelf.

ECN Explicit Congestion Notification.

ETSI European Technical Standards Institution.

EXP Experimental.

HW Hardware.

IESG Internet Engineering Steering Group.

IETF Internet Engineering Task Force.

INF Informational.

IntArea Internet Area.

IRTF Internet Research Task Force.

ISG Industry Specification Group.

MANO Management and Operation WG.

MPTCP Multipath Transmission Control Protocol (TCP).

NF Network Function.

NFV Network Functions Virtualisation.

NFVI NFV Infrastructure.

NFVO NFV Orchestrator.

NOC Network Operator Council.

OAM&P Operations, Administration, Maintenance and Provisioning.

PEG Performance Expert Group.

REL Reliability & Availability.

RFC Request For Comments.

SDN Software Defined Network.

TCP Transmission Control Protocol.

TCPM TCP Maintenance & Minor Modifications.

TSV Transport Services.

VNF Virtualised Network Function.

WG Work Group.

2 Introduction

This deliverable details the dissemination activities performed by the Trilogy 2 project during its first year. It is structured in the following way. We first present in section 2 the reports of two academic workshops we organized during this first year. In section 3 we present the standardization activities and in section 4 we provide details of the papers we have produced. Section 5 includes other dissemination activities of the project.

3 Academic workshops

The Trilogy 2 project organized two academic workshops during the first year, namely, the Global Internet Symposium and the Hot MiddleBox workshop. The two workshops are complementary in multiple ways: the Global Internet Symposium is a well established workshop that has high visibility and regularly attracts good quality papers. The focus of workshop is the global nature of the different Internet technologies, so it is quite broad. The Hot Middlebox workshop is a new workshop and its scope is much more narrow to technologies that are particularly relevant to the Trilogy 2 project. Both of them are held in conjunction with top level conferences, namely the Global Inter Symposium is held with IEEE Infocom and the Hot Middlebox workshop is held in ACM Conext. The Trilogy 2 project sponsored both workshops. We provide a report for each of them next.

3.1 Global Internet Symposium Report

The 16th Global Internet Symposium took place in Torino, Italy, on the 19th of April 2013, in conjunction with IEEE INFOCOM 2013. We include the Call for papers below.

There were 28 papers submitted to the workshop. Of these submissions, 13 papers have authors from institutions from the Asia Pacific region, 12 papers have authors from European Institutions, 5 have authors from North American institutions and 2 papers have authors from institutions based in Africa.

There were 13 papers accepted from the 28 reviewed, for an acceptance ratio of 46%. Each paper received reviews from at least three TPC members. We include the list of the PC members below. The accepted papers were organized into four sessions on Global Datasets, Inferring the Shape of the Internet, New Mechanisms and New Tools. We include the detail of the program below.

There were around 20 attendees at the workshop. Given the large number of workshops organized in conjunction with IEEE INFOCOM in 2013, and the timing on the last day of the conference, we consider the attendance to be reasonable. The workshop ran smoothly, with good feedback and discussion from the audience.

3.1.1 Program Committee

The program committee of GI13 had two chairs, one of them (Marcelo Bagnulo) part of the Trilogy 2 project. The committee had other 23 members of which 3 were part of the Trilogy project (Olivier Bonaventure, Saverio Niccolini and Costin Raiciu).

The details of the committee are as follows:

Technical Programm Committee Co-chairs:

- Marcelo Bagnulo (Universidad Carlos III de Madrid)
- Rod Van Meter (Keio University)

Program Committee

- Fred Baker (Cisco)
- Olivier Bonaventure (Universit catholique de Louvain)
- Luca Cittadini (Roma Tre University)
- Phil Eardley (BT)
- Marwan Fayed (University of Stirling)
- Wu-chang Feng (Portland State University)
- Pierre Francois (Imdea Networks)

- Polly Huang (National Taiwan University)
- Sukumal Kitisin (Kasetsart University)
- Mirja Kuehlewind (University of Stuttgart)
- Olaf Maennel (Loughborough University)
- Aniket Mahanti (University of Auckland)
- Daniel Massey (Colorado State University)
- Saverio Niccolini (NEC)
- Joerg Ott (Aalto University)
- Christos Papadopoulos (Colorado State University)
- Cristel Pelsser (IIJ)
- Colin Perkins (University of Glasgow)
- Costin Raiciu (University Politehnica of Bucarest)
- Pasi Sarolahti (Aalto University)
- Rade Stanojevic (Telefonica I+D)
- Stephen Strowes (Boundary)
- Joerg Widmer (Institute IMDEA Networks)
- He Yan (AT&T Labs - Research)

3.1.2 Technical Program

Session I: Global Datasets

Modeling and Characterization of Vehicular Density at Scale, Gautam Thakur (University of Florida, USA); Pan Hui (Deutsche Telekom Laboratories & Univeristy of Cambridge, Germany); Ahmed Helmy (University of Florida, USA)

On the Usage Patterns of Multimodal Communication: Countries and Evolution, Yi Wang (University of California, Riverside & Sprint Lab, USA); Michalis Faloutsos (University of California, Riverside, USA); Hui Zang (Sprint, USA)

DHT-based Traffic Localization in the Wild, Matteo Varvello (Bell Labs, Alcatel-Lucent, USA); Moritz Steiner (Bell Labs / Alcatel-Lucent, USA)

Session II: Inferring the Shape of the Internet

DRAGO: Detecting, Quantifying and Locating Hidden Routers in Traceroute IP Paths, Pietro Marchetta (University of Napoli, Italy); Antonio Pescap (University of Napoli Federico II, Italy)

The BGP Visibility Scanner, Andra Lutu (Institute IMDEA Networks & University Carlos III of Madrid, Spain); Marcelo Bagnulo (University Carlos III of Madrid, Spain); Olaf M Maennel (Loughborough University, United Kingdom)

Improving the Discovery of IXP Peering Links through Passive BGP Measurements, Vasileios Giotsas (University College London, United Kingdom); Shi Zhou (University College London, United Kingdom)

Session III: New Mechanisms

AliasCluster: A Lightweight Approach to Interface Disambiguation, Larissa Spinelli (Boston University, USA); Mark Crovella (Boston University, USA); Brian Eriksson (Technicolor, USA)

Applying Software-Defined Networking to the Telecom Domain, Georg Hampel (Bell Labs, Alcatel-Lucent, USA); Tian Bu (Bell Labs, USA); Moritz Steiner (Bell Labs / Alcatel-Lucent, USA)

Into the Moana Hypergraph-based Network Layer Indirection, Yan Shvartzshnaider (The University of Sydney & NICTA, Australia); Max Ott (NICTA, Australia); Olivier Mehani (NICTA, Australia); Guillaume Jourjon (NICTA, Australia); Thierry Rakotoarivelo (NICTA-Sydney, Australia); David Levy (University of Sydney, Australia)

COBRA: A Framework for the Analysis of Realistic Mobility Models, Gautam Thakur (University of Florida, USA); Ahmed Helmy (University of Florida, USA)

Session IV: New Tools

Network Emulation Testbed for DTN Applications and Protocols, Razvan Beuran (National Institute of Information and Communications Technology, Japan); Shinsuke Miwa (National Institute of Information and Communications Technology, Japan); Yoichi Shinoda (Japan Advanced Institute of Science and Technology, Japan)

Trade-off Analysis of Multi Topology Routing Based IP Fast Reroute Mechanisms, Selcuk Cevher (Karadeniz Technical University, Turkey); Mustafa Ulutas (Karadeniz Technical University, Turkey); Ibrahim Hokelek (TUBITAK BILGEM, Turkey)

Detecting Encrypted Botnet Traffic, Han Zhang (Colorado State University, USA); Christos Papadopoulos (Colorado State University, USA); Daniel Massey (Colorado State University, USA)

3.1.3 Call for Papers

The 16th IEEE Global Internet Symposium will be held in conjunction with IEEE Infocom 2013 in Turin, Italy on 19 April 2013. All relevant dates, location, and travel information are available from the IEEE Infocom 2013 website[1].

[1] http://infocom.di.unimi.it/index.php/giwksp.html

The IEEE Global Internet Symposium aims to provide a forum for researchers and practitioners to present and discuss advances in Internet-related technologies. The focus of the symposium is on experimental systems and on emerging future Internet technologies, and especially on scaling such systems to a global scale. The Program Committee encourages original submissions describing promising work in progress, speculations about the future of the Internet, and progressive position papers (which should be clearly marked as such).
The proceedings of the 16th IEEE Global Internet Symposium will be published on-line through the IEEE Xplore Digital Library after IEEE Infocom 2013 concludes.

Topics

Authors are invited to submit papers on issues, especially scaling issues, related to current and future Internet technology. Topics of interest include, but not limited to, the following:

- Understanding Internet protocols and applications at global scale Internet Measurements and Methodology
- Traffic measurement, analysis, modeling, and visualization
- Network architectures
- Handling Internet dynamics/heterogeneity (by applications and/or the network)
- Large-scale distributed Internet applications
- Privacy and/or security issues in the Internet
- Anomaly, intrusion and attack detection
- Routing (unicast, multicast, anycast, etc.)
- Flow management (fairness/sharing, congestion control, differentiated services, etc.)
- The Internet and wireless/mobile devices, as well as intermittent connectivity
- P2P networking and overlay networks
- Provisioning, monitoring, and management of IP services
- Content networking (caching, content distribution, content routing, content services, load balancing, etc.)
- Interactions of the smart grid systems with the Global Internet.
- Impact of data centers in the global Internet.

In all cases, GI should focus on a global scope of solution, i.e., how these issues change or become interesting when they are of global scope or global coordination.

3.2 CoNEXT Hot Middlebox Workshop Report

The 2013 ACM Workshop on Hot Topics in Middleboxes and Network Function Virtualization Hot Middlebox13 took place in Santa Barbara, California, on December 9th, 2013. It was the only accepted workshop for CoNEXT 2013 (except for the student workshop, which is be default always part of the conference.
The call for papers attracted 20 submissions, with the program committee accepting 8 from these, a 40% acceptance rate. The papers covered a wide range of topics, including middlebox traversal, platforms for network function virtualization and middlebox processing, verification, and the evolvability of the Internet in the presence of middleboxes.

3.2.1 Call for Papers

Modern networks increasingly rely on advanced network processing functions for a wide spectrum of crucial functions ranging from security (firewalls, IDSes, traffic scrubbers), traffic shaping (rate limiters, load balancers), dealing with address space exhaustion (NATs) or improving the performance of network applications (traffic accelerators, caches, proxies), to name a few. Such "network appliances" or "middleboxes" are a critical piece of the network infrastructure and represent, to a first-order approximation, the de-facto approach for network evolution in response to changing performance, security, and policy compliance requirements.
However, most of this functionality is implemented in costly, hard-to-modify dedicated hardware, making the network difficult to evolve or adapt to changing traffic requirements. Recent work seeks to address this issue by shifting network processing from a world of dedicated hardware to one built on software-based processing running on (sometimes virtualized and shared) platforms built on commodity hardware servers, switches, and storage. This vision of "software-based" network services enables new in-network functions to be rapidly instantiated, on-demand, and at places in the network where it is most needed, without having to modify the underlying hardware. The scope of this workshop focuses both on the design of the data plane to support advanced services as well as the control plane functions necessary to manage these advanced data plane functions. In some sense, this vision is complementary to ongoing efforts in the SDN community, where the focus has largely been on the control plane and assuming a commodity data plane.
While our workshop builds on the recent promise of realizing high-performance network processing on commodity hardware, many questions remain open:

- What are the best virtualization technologies for implementing high-performance network functions?
- What are the challenges when trying to push them to rates of 10Gb and beyond?
- How do we provide the best possible isolation, both in terms of software isolation but also performance?
- How do we ensure that middlebox modules from different entities running on the same platform are assigned to the available hardware in an optimal way?

- How can we provide quick instantiation of processing (in the order of milliseconds or less)?
- What control plane abstractions are necessary to manage such advanced and stateful services?

The HotMiddlebox workshop will serve as an avenue to showcase and discuss ongoing work from both academic and industry efforts in this space and to identify key challenges and potential solutions, with the ultimate goal of providing a roadmap for practical deployment in operational networks.
We encourage the submission of work-in-progress papers in the area of middlebox design, implementation, measurement, management, and deployment. We look for submissions of previously unpublished work on topics including, but not limited to, the following:

- Performance optimizations of network stacks on virtualized systems
- Verification of unknown code running on shared middlebox platforms
- Extensible software stacks for rapid implementation of new middlebox functions
- Mechanisms for migration of stateful middleboxes
- Resource allocation mechanisms for shared/virtualized middlebox platforms
- Integrating new software middleboxes into legacy networks
- Backend storage/memory architectures for middleboxes
- Management abstractions and policy language frameworks for middleboxes
- Experiences in deploying software-based middleboxes in operational networks
- Connections to Software-Defined Networking
- Deployment and use of middleboxes in the cloud
- Measurements of middleboxes in enterprise, ISP, and data center networks.
- Novel security, performance, and monitoring applications atop middleboxes
- Challenges for policy verification in the context of middlebox services.

Submission Guidelines: each submission must be a single PDF file no longer than six (6) pages in length (in two-column, 10-point format) including references, following the LaTeX style file. Papers should be submitted electronically via the submission site. Papers must include the author name and affiliation for single-blind peer reviewing by the program committee. Accepted papers will be published in the ACM Digital Library. Publication at Hot Middlebox is not intended to preclude later publication of an extended version of the paper. Authors of accepted papers are expected to present their papers at the workshop.
Important dates:

- Submission: September 6, 11:59pm EST
- Acceptance notification: October 6, 2013
- Camera ready: October 24, 2013
- Workshop date: December 9, 2013

3.2.2 Program Committee

Workshop Co-Chairs: Felipe Huici NEC Europe, UK, Vyas Sekar, Stony Brook University, USA

Technical Program Committee:

- Aditya Akella, University of Wisconsin, USA
- Katerina Argyraki, EPFL, SWITZERLAND
- Jon Crowcroft, Cambridge, UK
- Lars Eggert, NetApp, USA
- Norbert Egi, Huawei, USA
- Mark Handley, University College London, UK
- Giuseppe Lettieri, University of Pisa, ITALY
- Diego Lopez, Telefonica, SPAIN
- Anil Madhavapeddy, Cambridge, UK
- KyoungSoo Park, KAIST, SOUTH KOREA
- Peter Pietzuch, Imperial College, UK
- Sylvia Ratnasamy, UC Berkeley, USA
- Jennifer Rexford, Princeton University, USA
- Luigi Rizzo, University of Pisa, ITALY
- Anees Shaikh, IBM, USA
- Robin Sommer, ICSI, USA
- Nick Feamster, Georgia Institute of Technology, USA
- Julian Chesterfield, OnApp, USA

- Teemu Koponen, Nicira / VMware, USA
- Christian Maciocco, Intel Labs, USA
- Bob Briscoe, British Telecome, UK

3.2.3 Technical Program

10:00-10:15 *Welcome Message*

10:15-11:05 **TCP**

10:15-10:40 *Multipath in the Middle(Box)* Gregory Detal, Christoph Paasch, and Olivier Bonaventure (Universite catholique de Louvain)

10:40-11:05 *Evolving the Internet with Connection Acrobatics* Catalin Nicutar (University Politehnica of Bucharest), Christoph Paasch (Universite Catholique de Louvain), Marcelo Bagnulo (Univeridad Carlos III Madrid), and Costin Raiciu (University Politehnica of Bucharest)

11:05-11:55 **Platforms**

11:05-11:30 *Towards Minimalistic, Virtualized Content Caches with Minicache* Simon Kuenzer, Joao Martins, Mohamed Ahmed, and Felipe Huici (NEC Europe Ltd.)

11:30-11:55 *FlowOS: A Flow-based Platform for Middleboxes* Mehdi Bezahaf and Abdul Alim (Lancaster University) and Laurent Mathy (University of Liege)

12:00-13:30 **Lunch**

13:30-14:20 **Verification**

13:30-13:55 *Verifiable Network Function Outsourcing: Requirements, Challenges, and Roadmap* Seyed Kaveh Fayazbakhsh (Stony Brook University), Michael K Reiter (UNC Chapel Hill), and Vyas Sekar (Stony Brook University)

13:55-14:20 *SymNet: Static Checking for Stateful Networks* Radu Stoenescu, Matei Popovici, Lorina Negreanu, and Costin Raiciu (University Politehnica of Bucharest)

14:20-14:45 **Coffee Break**

14:45-15:35 **Traversal**

14:45-15:10 *Are TCP Extensions Middlebox-proof?* Benjamin Hesmans, Fabien Duchene, Christoph Paasch, Gregory Detal, and Olivier Bonaventure (Universite catholique de Louvain)

15:10-15:35 *Analysis and Topology-based Traversal of Cascaded Large Scale NATs* Andreas Mller, Florian Wohlfart, and Georg Carle (TU Munich)

3.2.4 Participants

Below is a list of the workshop's participants:

- Norbert Egi Huawei Technologies Inc.

- Gregory Lauer Raytheon BBN Technologies
- Laurent Mathy Lancaster University
- Tom Barbette Universite de Liege
- Benjamin Hesmans ucl
- Andreas Mueller TU Munich
- Gregory Detal UCL
- Fabien Duchene ucl
- Nazim Ahmed University of Texas at Dallas
- Craig Russell CSIRO
- Richard Antiabong University of Texas at Dallas
- Christian Kreibich ICSI
- Abdul Alim Lancaster Universtiy
- Yukihiro Nakagawa Fujitsu Laboratories Ltd.
- Nanxi Kang Princeton University
- Seyed Kaveh Fayazbakhsh Stony Brook University
- Kazuya Tsukamoto Kyushu Institute of Technology
- Simon Kuenzer NEC Laboratories Europe Ltd.
- IZUMI MIZUTANI Hitachi, Ltd.
- Kevin Glavin Riverbed Technology
- Christoph Paasch UCLouvain
- Marcin Furtak Samsung Electronics Polska Sp. z o.o.
- Peter Steenkiste Carnegie Mellon University
- Vyas Sekar Stony Brook University
- Felipe Huici NEC Europe Ltd.

4 Standardization

The Trilogy 2 project has devoted a significant amount of effort in standardization of the project output. Trilogy 2 has focussed its effort in two standard organizations, namely, ETSI and the IETF, which we expand next.

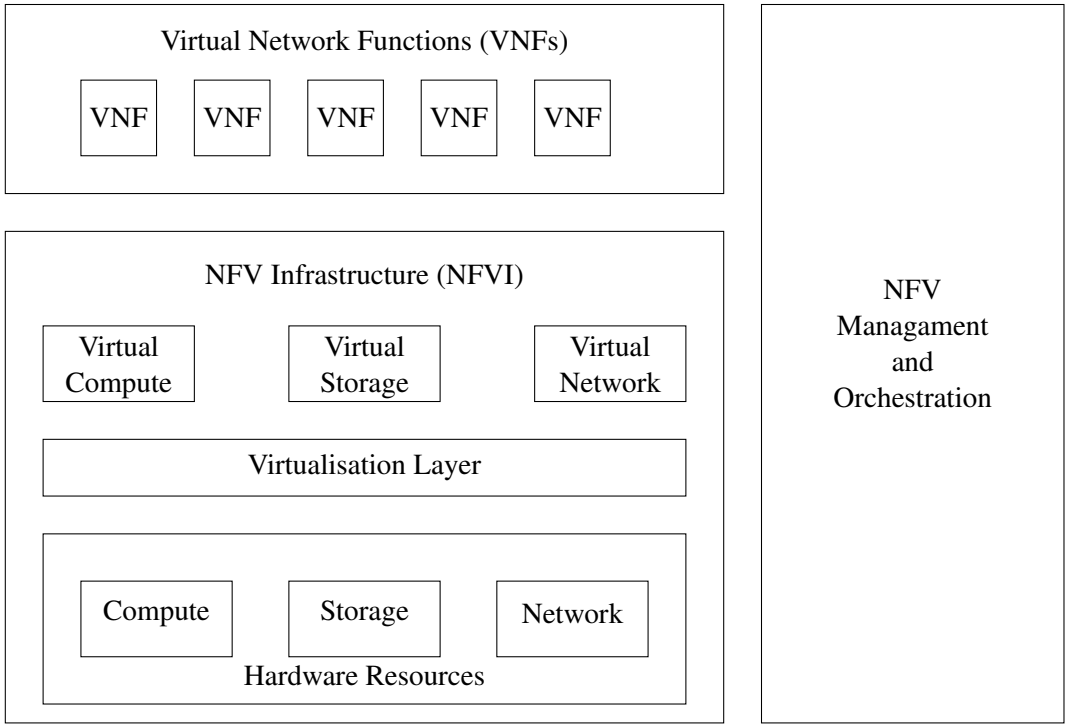

Figure 4.1: The NFV architecture framework.

4.1 Network Functions Virtualisation (NFV) activities at ETSI

4.1.1 The NFV Industry Specification Group (ISG)

In summer 2012, a small group of network operators, including the Trilogy 2 teams in BT and Telefónica, began to establish a new eco-system to develop the idea of providing Network Functions (NFs) in software on Commercial Off-The-Shelf (COTS) Hardware (HW) using virtualisation techniques. Later that summer, they had settled on the name Network Functions Virtualisation (NFV) and had written a white paper that demonstrated the strong consensus between the thirteen founding operators [16]. The operators published this in September 2012 at the same time as announcing their agreement with ETSI to establish an Industry Specification Group on NFV.

The aim of NFV is to provide improved flexibility of the network creation process and cost efficiencies in capital and operations compared to the traditional approach based on dedicated HW, as well as reducing the risk of vendor lock-in. In addition, the concept of Virtualised Network Functions (VNFs) offers increased scalability and decouples NFs from the physical resources they need. NFV also accelerates the service innovation cycle due to its underlying software-based service development and deployment paradigm.

The launch meeting was held in the first month of the Trilogy 2 project (January 2013) in Valbonne, France. It attracted 165 delegates, who agreed that the ISG should focus on identifying gaps in standardisation needed for NFV, with the aggressive aim of completing and shutting down the ISG within 2 years, having identified preferred standards development organisations to adopt each task. All this aligned well with the 3-year duration of Trilogy 2.

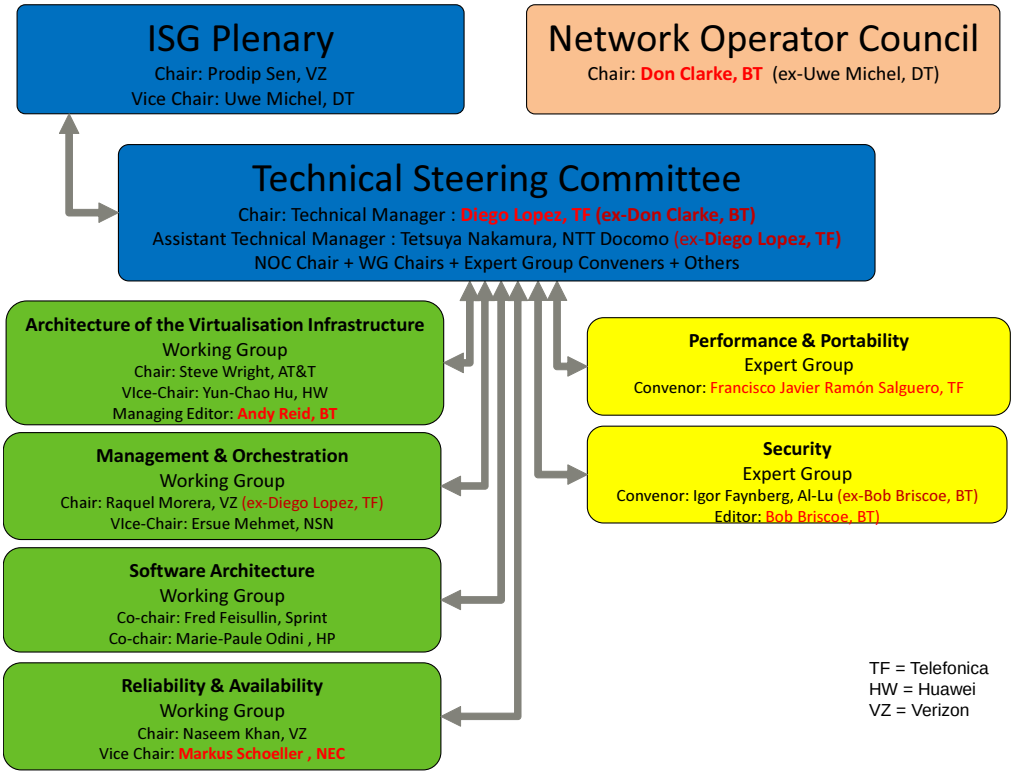

Figure 4.2: Organisation of the ETSI NFV Industry Specification Group (ISG)

Over the ensuing year, the working groups, expert groups and steering groups set up by the launch meeting have held an average of about 8 audio conferences per week (406 in 11 months), plus 3 further face-to-face plenaries with growing numbers of delegates (285 in Santa Clara, CA; 227 in Bonn, DE; and 345 in Sunnyvale, CA). As a result, a half a dozen stable draft specifications have been published that document the high level architecture (see Figure 4.1), use-cases, terminology, etc., and the ISG's requirements have started to be issued as liaison statements to other standards bodies. During this period, the number of signed up companies and institutions has grown to 170. ETSI's NFV home page [9] provides links to published documents, as well as related material such as groups working on Proofs of Concept.

4.1.2 NFV ISG Organisation

Figure 4.2 shows the organisational structure of the NFV ISG, and highlights officers of the ISG who are closely associated with Trilogy 2, including those who were initially voted into posts but who handed over their role during the first year.

The main work of the ISG is conducted in the 4 work groups and 2 expert groups shown, and described below. Major decisions are made by the whole ISG, while day-to-day technical co-ordination is overseen by the Technical Steering Committee.

ETSI provides a legal and administrative support framework within which industry specification groups can operate, but otherwise allows them considerable autonomy. For instance outputs do not have to be authorised by ETSI, which is one reason why the founding network operators originally chose ETSI as host for the

ISG. From the start, the network operators formed themselves into a Network Operator Council (NOC) that provides steer by articulating and clarifying requirements. The Network Operator Council (NOC) has no formal powers over the ISG, but it is available in an advisory capacity, representing the main customers of the ISG's work. The NOC now consists of 25 network operators, who recently published an update to the NFV white paper that originally launched the ISG.

4.1.2.1 The Architecture of the Virtualisation Infrastructure WG

The primary difference between NFV and traditional networking is that network functions are separated into low level infrastructure network elements (e.g. Ethernet switches and cabling) connected to compute infrastructure running Virtualised Network Functions (VNFs). This WG defines the architecture of the network and compute infrastructure that supports virtualisation, and identifies gaps where standardisation will be required. The NFV ISG's Architecture of the Virtualisation Infrastructure [2] is primarily concerned with providing an execution environment for VNFs and allocating physical resources. It is a direct derivative of the NFV aspects of the Trilogy 2 Architecture [1], which should be referred to for details.

4.1.2.2 The Software Architecture WG

Once running, each VNFs may have identical interfaces to its non-virtualised counterpart, which aids migration and allows the focus of new engineering effort to rest on the virtualisation environment. Nonetheless, virtualisation offers opportunities to decompose traditional network functions into component parts. This WG investigates such opportunities for restructuring network functions, aiming for re-use, scalability etc.

4.1.2.3 The Management and Operation WG (MANO) WG

NFV adds new capabilities to communications networks. In addition to functions provided by the current model of Operations, Administration, Maintenance and Provisioning (OAM&P), a new set of management and orchestration functions specific to the virtualised environment of NFV is needed.

Non-virtualised Network Functions (NFs) are often tightly coupled with the infrastructure they run on in legacy networks, resulting in vendor lock-in for Service and Network Providers, as well as inflexibility in allocation and re-allocation of resources between functions. NFV decouples the implementation of the NF from the resources it is deployed upon. These include computation, storage, and networking resources and are addressed as NFV Infrastructure (NFVI).

The decoupling exposes a new set of entities, the VNFs, and a new set of relationships between them and the NFV Infrastructure (NFVI), which did not exist in traditional networks. The Management and Operation WG (MANO)WG has incrementally worked out the new aspects introduced by the NFV process based upon those aspects of NFs that remained constant.

The NFV Orchestrator (NFVO), shown in Figure (4.1), has the role of managing the NFVI, orchestrating the allocation of resources needed by the VNFs, managing the new aspects of VNFs introduced by virtualisation, and coordinating the relationships between the NFVI and the VNFs. Work on the NFVO has so far produced

the Network Service and VNF life-cycle message flows as well as the information elements for NFV entities. To date (Dec 2013), the MANO WG is shifting its work towards NFV operational management focusing on fault and event management. Other aspects like capacity planning or migration are planned for the future.

4.1.2.4 The Reliability & Availability WG

The Reliability & Availability (REL) WG has focused on providing resilience to failure using redundancy and load-balancing between multiple VNF instances.

4.1.2.5 The Performance Expert Group (PEG)

Two critical and potentially conflicting aspects of the new environment built around the virtualisation technologies provided by NFV are performance and portability. It must be able to deploy VNFs over different types of NFVIs. However, VNFs may perform differently depending on the underlying infrastructure. Predictability is key in order to provide a reliable network infrastructure.

The NFV Performance Expert Group has produced a list of minimal features which templates should contain for the appropriate deployment of VNFs over an NFVI in order to meet these goals.

Additionally, the Performance Expert Group (PEG) has provided a set of recommendations on the minimum requirements which HW and hypervisor should have for an NFVI to be suitable for VNFs implementing different workloads (e.g. data-plane, control-plane, etc.) and a set of best practises regarding SW design and configuration, and Operating Systems.

4.1.2.6 The Security Expert Group

The Security Expert Group primarily provides a security review service to other WGs. It has also produced an NFV Security Problem Statement [**?**] that identifies new vulnerabilities that may arise as a result of the introduction of NFV, as opposed to existing security vulnerabilities, either in networking or in application virtualisation (cloud). As well as identifying potential new security problem, the document also provides a reference framework in which to articulate security problems.

A certain amount of the insight in the NFV security problem statement has come from the Trilogy 2 project, particularly from the security aspects of the Trilogy 2 architecture [1], which addresses a similar but wider problem space.

4.1.3 NFV ISG Outputs from Trilogy 2 Partners

4.1.3.1 Telefónica

Telefónica had been experimenting with virtualisation of NFs in the period prior to the launch of the ETSI NFV ISG, so it was natural to become one of the founding network operators.

At the launch meeting Diego Lopez, who had brought Telefónica into the Trilogy 2 project, was appointed to two senior positions: Assistant Technical Manager of the ISG and chair of the Management and Operation WG (MANO) WG. Also Telefónica's Francisco Javier Ramón Salguero was appointed to chair the Performance Expert Group (PEG). In a subsequent reshuffle, Diego became technical manager of the whole ISG,

chairing the Technical Steering Committee.

4.1.3.2 BT

In the years prior to the formation of the ETSI NFV ISG, BT's network virtualisation team had worked with partners to build a virtualised BRAS and conducted a full performance evaluation and an in-depth cost study. In June 2012, they published the surprising result that the 5-year total cost of ownership of a generic hardware solution was already safely lower than that of bespoke hardware with equivalent performance, even for the 'acid test' case of a BRAS.

The team's main contribution though was to recognise, articulate and act on the strategic significance to the industry of these results, which would not be realised if each operator continued to act merely as independent customers rather than using their collective buying power to focus the industry to take this direction seriously. The BT team persuaded a similar team in Verizon to seek out and draw in other operators such as Telefónica interested in forming what became the NFV ecosystem.

The BT team created an NFV white paper as a vehicle to draw out consensus, and retained the editorial role through to its publication [16]. At the launch meeting Don Clarke was appointed Chair of the ISG's Technical Steering Committee, and Andy Reid, having led the discussions on architectural direction [17], was appointed managing editor of all the infrastructure architecture documents. Also Bob Briscoe was appointed to chair and convene the Security Expert Group. In a subsequent reshuffle, Don Clarke passed the Technical Manager role to Diego Lopez (see Telefónica contributions above) and became chair of the Network Operator Council. Also Bob Briscoe passed on the role of Security Expert Group convenor, but retained editorship of the sole Security document.

The main ETSI NFV outputs that BT has contributed significant Trilogy 2 outputs into are:

- The NFV Infrastructure Architecture Overview (Editor: Andy Reid) [2] and the other outputs of the Architecture of the VIrtualisation Infrastructure WG;
- The NFV Security Problem Statement (Editor: Bob Briscoe) [3]

BT has also contributed numerous presentations, the most important being to the full plenary meetings or joint working group meetings as follows:

- Andy Reid (as rapporteur) "NFV Infrastructure Architecture" 16-Jan-13 ETSI NFV ISG plenary#1 (Valbonne, FR) to a plenary audience of 165 engineers/strategists;
- Bob Briscoe (as chair) "NFV Security Expert Group Status Update" Apr-13 ETSI NFV ISG plenary#2 (Santa Clara, CA) to a plenary audience of 285 engineers/strategists;
- Bob Briscoe (as rapporteur) "NFV Security Reference Framework" Jul-13 ETSI NFV joint sw-arch/security meeting (Bonn, DE) to an audience of 80 engineers;

- Bob Briscoe (as rapporteur) “NFV Security; Problem Statement ” 30-Oct-13 ETSI NFV ISG plenary#4 (Sunnyvale, CA) to a plenary audience of 345 engineers/strategists;

4.2 IETF Standardisation

The first Trilogy project led to the formation of two IETF working groups: Multipath TCP (MPTCP) and Congestion Exposure (ConEx). The Trilogy 2 project continues the work on the IETF, with special focus on these two working groups.

4.2.1 MultiPath TCP (MPTCP) Working Group

BT's Philip Eardley co-chairs the MPTCP WG. The main MPTCP protocol specification [10] is in the process of update [11], with changes limited to behavioural clarifications and new messages that can coexist with earlier implementations. 3 of the 4 authors of these documents are continuing this work as part of Trilogy 2 (Costin Raiciu, UPB, Mark Handley, UCLondon and Olivier Bonaventure, UCLouvaine).

The work performed by the Trilogy 2 project in MPTCP security has been adopted by the MPTCP working group. In particular, the following draft has been adopted as working group item and it is expected to be published as RFC in the short term: M. Bagnulo, C. Paasch, F. Gont, O. Bonaventure, C. Raiciu, Analysis of MPTCP residual threats and possible fixes, draft-bagnulo-mptcp-attacks-01, October 2013.

The document describes a roadmap for the MPTCP security work which has been adopted by the working group and will be performed in the near future. We expect that the Trilogy 2 work on MPTCP security will fit in the proposed roadmap.

In addition to the aforementioned document, the Trilogy 2 partners have performed the following presentations in the MPTCP WG:

- M. Bagnulo, MPTCP residual threats, Presentation at IETF MPTCP WG meeting, Jul 28, 2013
- M. Bagnulo, Connection acrobatics, Presentation at IETF MPTCP WG meeting, Nov 4, 2013
- C. Paasch, MPTCP Linux Kernel implementation - Status update, Presentation at IETF MPTCP WG meeting, Jul 28, 2013
- O. Bonaventure, Multipath TCP, tutorial at IETF'87, July 2013.

4.2.2 Conex Working Group

UC3M's Marcelo Bagnulo Braun continues to co-chair the ConEx WG, which published RFC6789 prior to the launch of Trilogy 2.

In addition, most of the Conex working group items have been led by Trilogy 2 partners, including the following documents.

"Byte and Packet Congestion Notification" has finally been authorised for publication as an RFC by the Internet Engineering Steering Group (IESG), having required many years of persuasion on its passage through two loops of the working group and IESG stages (BT's Bob Briscoe wrote the first draft in Jun 2007, before even the first Trilogy project). It is an architecturally important Request For Comments (RFC) that specifies

that the units used for the metric of bandwidth congestion will normally be bit per second not packet per second. And it explains that a congested resource should drop or mark packets with a probability independent of packet size.

"ConEx Concepts & Abstract Mechanism" is the main ConEx document that describes the architecture of ConEx and the security mechanisms, but in abstract form rather than the protocol specifics. It has passed two working group last calls and is waiting to be passed to the IESG.

"Network Performance Isolation in Data Centres using Congestion Policing" describes mechanisms to enforce limits on the ability of one data centre tenant to affect the performance of other tenants. It is intended to be an informational document describing a ConEx deployment scenario. However, the WG is first waiting to get the primary documents through, before moving on to deployment scenarios.

"Network Performance Isolation using Congestion Policing" explains why using congestion-rate as the metric for policing works so well, and simplifies networks compared to other metrics such as bit-rate or volume. It complements the other ConEx documents that use congestion policing. Again, this is intended to be an informational document, but it is being held until the primary ConEx documents are complete.

The main hold-up has been reaching agreement on a definition of congestion credit, which is used in the ConEx architecture to simplify audit (protocol integrity checking). Agreement was reached on the definition of credit in Nov 2013.

"Reusing the IPv4 Identification Field in Atomic Packets" is a placeholder that describes a way to encode ConEx in the IPv4 header for private networks such as data centres. ConEx is officially only being specified for IPv6 initially. Therefore this draft is not expected to be adopted by the Internet Area unless ConEx starts to be used widely in IPv4 for private networks.

BT's Bob Briscoe will be listed as a co-author on the next revision of *"Problem Statement and Requirements for a More Accurate ECN Feedback"*. This explains the problem that makes ConEx inaccurate when used over TCP with ECN not just drop. The TCP Maintenance & Minor Modifications (TCPM) WG wants this problem statement to go through the IESG before adopting any draft that proposes a solution.

"More Accurate ECN Feedback in TCP" specifies a preferred way to solve this accuracy problem with TCP's ECN feedback. Although the author list does not include anyone from Trilogy 2, it acknowledges that most of the scheme is based on the work of BT's Bob Briscoe.

Tthe following presentations of Trilogy 2 work at the IETF:

- Bob Briscoe presented *"ConEx Concepts & Abstract Mechanism"* to the IETF ConEx WG (Berlin, DE in Jul 2013); Audience: c.40 Researchers/Engineers;
- Bob Briscoe presented *"Network Performance Isolation in Data Centres"* & *"Congestion Policing"* to the IETF ConEx WG (Berlin, DE in Jul 2013); Audience: c.40 Researchers/Engineers;
- Bob Briscoe presented *"Network Functions Virtualisation"* [6] to the Internet Research Task Force

WG	Intended status	Progress	Authors	Title
MPTCP	INF	WG item	Marcelo Bagnulo, UC3M, Christoph Paasch (UCLouvain), O. Bonaventure (UCLouvain) , C. Raiciu (UPB) (& F. Gont)	Analysis of MPTCP residual threats and possible fixes [14]
TSV	BCP	RFC Editor Queue	Bob Briscoe, BT (& Jukka Manner, Aalto Uni)	Byte and Packet Congestion Notification [7]
ConEx	INF	WG draft	Bob Briscoe, BT (& Matt Mathis, Google)	ConEx Concepts & Abstract Mechanism [15]
ConEx	INF	individual draft	Bob Briscoe, BT (& Murari Sridharan, Microsoft)	Network Performance Isolation in Data Centres using Congestion Policing [8]
ConEx	INF	individual draft	Bob Briscoe, BT	Network Performance Isolation using Congestion Policing [4]
IntArea	EXP	individual draft	Bob Briscoe, BT	Reusing the IPv4 Identification Field in Atomic Packets [5]
TCPM	INF	WG draft	Mirja Kühlewind, Stuttgart Uni (& Richard Scheffenegger, NetApp)	Problem Statement and Requirements for a More Accurate ECN Feedback [13]
TCPM	EXP	individual draft	Mirja Kühlewind, Stuttgart Uni (& Richard Scheffenegger, NetApp)	More Accurate ECN Feedback in TCP [12]

Table 4.1: Trilogy 2 contributions to the IETF.

(IRTF)'s Software Defined Network (SDN) research group (Orlando, FL, US in Mar 2013); Audience: c.300 Researchers/Engineers/Strategists.

5 Academic papers

During this first year of the project, the Trilogy 2 project has strived to publish its results in relevant academic venues. Since the project is in its early stages and the results are in preliminary stage most of the selected venues were workshops. We detail next the list of published papers and the list of submitted papers.

Published papers:

- Christoph Paasch, Ramin Khalili, Olivier Bonaventure, *On the Benefits of Applying Experimental Design to Improve Multipath TCP*, published in ACM CoNext 2013, Santa Barbara, California.
- Simon Kuenzer, *Towards Minimalistic, Virtualized Content Caches with Minicache*, published in HotMiddlebox'13 workshop (a Trilogy 2 sponsored workshop), collocated with ACM Conext 2013, Santa Barbara, California.
- Catalin Nicutar, Christoph Paasch, Marcelo Bagnulo and Costin Raiciu, *Evolving the Internet with Connection Acrobatics*, published in HotMiddlebox'13 workshop (a Trilogy 2 sponsored workshop), collocated with ACM Conext 2013, Santa Barbara, California.
- George Parisis, Toby Moncaster, Anil Madhavapeddy and Jon Crowcroft, *Trevi: Watering Down Storage Hotspots with Cool Fountain Codes*, published in ACM Hotnets'13, College Park, MD.
- Roberto Bifulco, Thomas Dietz, Felipe Huici, Mohamed Ahmed, Joao Martins, Saverio Niccolini and Hans-Joerg Kolbe, *Rethinking Access Networks with High Performance Virtual Software BRASes*, published in EWSDN 2013, BErlin Germany.

Submitted papers:

- A. Agache, R. Deaconescu, C. Raiciu, *Opportunistically Increasing Datacenter Network Utilization with GRIN*, Paper submission NSDI 2014.
- C. Nicutar, D. Niculescu, C. Raiciu, *Mobile Kibbutz*, Paper submission Conext 2013.

6 Other dissemination activities

In addition to the dissemination activities presented in the previous sections, the Trilogy 2 project has performed the following activities.

- A Multipath TCP Keynote by Olivier Bonaventure at the IEEE LAN/MAN Workshop in Brussels in April 2013
- A Tutorial on Multipath TCP and Its Use-Cases in the NorNet Testbed by Christoph Paasch at the NorNet Users Workshop in Oslo in September 2013
- A presentation on the concertation meeting organized by the EU Commission with the title ”Building the liquid network” by M. Bagnulo in October 2012
- A press release titled *”OnApp Joins Industry Leaders on Trilogy 2 ’Liquid Network’ Project”* by Julien Chesterfield, which was published in PingZine! [1], CDNadvisor [2], StorageNewsletter [3] and TheEnterpriseCloudSite [4].

[1] http://www.pingzine.com/onapp-joins-trilogy-2-project-to-enable-liquid-network-28526/
[2] http://www.cdn-advisor.com/onapp-becomes-part-of-trilogy-2-liquid-network-project/
[3] https://www.storagenewsletter.com/news/channel/onapp-trilogy-2
[4] http://www.theenterprisecloudsite.com/author.asp?section_id=3236&doc_id=268164&

Bibliography

[1] Initial Liquid Net Architecture. Deliverable ICT-317756 D2.1, Trilogy 2, December 2013.

[2] Network Functions Virtualisation; Part 1 Infrastructure Architecture; Sub-part 1 Overview. Group Specification ETSI GS NFV INF 001-1 v0.3.2, ETSI NFV ISG, November 2013. (work in progress).

[3] Network Functions Virtualisation; Security; Problem Statement. Group Specification ETSI GS NFV SEC001 v0.0.8, ETSI NFV ISG, December 2013. (work in progress).

[4] Bob Briscoe. Network Performance Isolation using Congestion Policing. Internet Draft draft-briscoe-conex-policing-00, Internet Engineering Task Force, February 2013. Work in progress.

[5] Bob Briscoe. Reusing the IPv4 Identification Field in Atomic Packets. Internet Draft draft-briscoe-intarea-ipv4-id-reuse-03, Internet Engineering Task Force, July 2013. (Work in progress).

[6] Bob Briscoe, Don Clarke, Pete Willis, Andy Reid, and Paul Veitch. Network Functions Virtualisation. Presentation: `http://www.ietf.org/proceedings/86/slides/slides-86-sdnrg-1.pdf`.

[7] Bob Briscoe and Jukka Manner. Byte and Packet Congestion Notification. Internet Draft draft-ietf-tsvwg-byte-pkt-congest-12.txt, Internet Engineering Task Force, May 2013. (Work in progress).

[8] Bob Briscoe and Murari Sridharan. Network Performance Isolation in Data Centres using Congestion Policing. Internet Draft draft-briscoe-conex-data-centre-01, Internet Engineering Task Force, February 2013. Work in progress.

[9] ETSI. Network Functions Virtualisation. Home page: `http://www.etsi.org/technologies-clusters/technologies/nfv`.

[10] Alan Ford, Costin Raiciu, Mark Handley, and Olivier Bonaventure. TCP Extensions for Multipath Operation with Multiple Addresses. Request for Comments 6824, Internet Engineering Task Force, January 2013.

[11] Alan Ford, Costin Raiciu, Mark Handley, and Olivier Bonaventure. TCP Extensions for Multipath Operation with Multiple Addresses. Internet Draft draft-ietf-mptcp-rfc6824bis, Internet Engineering Task Force, January 2013.

[12] Mirja Kühlewind and Richard Scheffenegger. More Accurate ECN Feedback in TCP. Internet Draft draft-kuehlewind-tcpm-accurate-ecn-02, Internet Engineering Task Force, June 2013. (Work in Progress).

[13] Mirja Kühlewind and Richard Scheffenegger. Problem Statement and Requirements for a More Accurate ECN Feedback. Internet Draft draft-ietf-tcpm-accurate-ecn-04, Internet Engineering Task Force, October 2013. (Work in Progress).

[14] F. Gont O. Bonaventure C. Raiciu M. Bagnulo, C. Paasch. Analysis of MPTCP residual threats and possible fixes . Internet Draft draft-bagnulo-mptcp-attacks-01.txt, Internet Engineering Task Force, 2013. (Work in progress).

[15] Matt Mathis and Bob Briscoe. Congestion Exposure (ConEx) Concepts and Abstract Mechanism. Internet Draft draft-ietf-conex-abstract-mech-08, Internet Engineering Task Force, October 2013. Work in progress.

[16] Network functions virtualisation - introductory white paper. `http://portal.etsi.org/NFV/NFV_White_Paper.pdf`, October 2012. SDN and OpenFlow World Congress, Darmstadt, Germany.

[17] Andy Reid. Track 1: Architecture of the Virtualization Infrastructure. Presentation to launch meeting of NFV ISG (ETSI-internal), Valbonne, Fr, January 2013.

