



ICT-317756 **TRILOGY2**

Trilogy2: Building the Liquid Net

Specific Targeted Research Project FP7 ICT Objective 1.1 The Network of the Future

D4.3 - Public record of the Trilogy 2 exploitation plan and its standardization workshop

Due date of deliverable: December 30, 2015 Actual submission date: December 30, 2015

Start date of project Duration Lead contractor for this deliverable Version Confidentiality status 1 January 2013 36 months Universidad Carlos III de Madrid v0.1 , 10.12.2015 Public

Abstract

This deliverable documents the Trilogy2 exploitation (plans and impact as of today), the Trilogy2 standardization workshop held in Heidelberg, Germany on the 30th of November 2015 as well as lists and highlights the project dissemination activities (publications, software, presentations at events, standardization activities and press releases).

Target Audience

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

Disclaimer

This document contains material, which is the copyright of certain TRILOGY2 consortium parties, and may not be reproduced or copied without permission. All TRILOGY2 consortium parties have agreed to the full publication of this document. The commercial use of any information contained in this document may require a license from the proprietor of that information.

Neither the TRILOGY2 consortium as a whole, nor a certain party of the TRILOGY2 consortium warrant that the information contained in this document is capable of use, or that use of the information is free from risk, and accept no liability for loss or damage suffered by any person using this information.

This document does not represent the opinion of the European Community, and the European Community is not responsible for any use that might be made of its content.

Impressum	
Full project title	TRILOGY2: Building the Liquid Net
Title of the workpackage	D.4.3 Public record of the Trilogy 2 exploitation plan and its standard-
	ization workshop
Editor	Saverio Nicolini, NEC
Project Co-ordinator	Marcelo Bagnulo Braun, UC3M
Copyright notice	© 2015 Participants in project TRILOGY2

Executive Summary

This deliverable documents the Trilogy2 exploitation (plans and impact as of today) and the Trilogy2 standardization workshop held in Heidelberg, Germany on the 30th of November 2015. The workshop format was a half day interim meeting of the IRTF NFV Research Group held immediately before the CoNEXT 2015 conference organized by NEC (http://conferences2.sigcomm.org/co-next/2015/) to give the opportunity to participants to attend a very important conference in the area of work of Trilogy2 as well as exploit the possibility of increasing the outreach of the workshop by capturing the attention of CoNEXT attendees. Additionally, this deliverable reports on the dissemination activities of members of the consortium during the period 2013-2015 as well as presenting its highlights (several publications at top tier events, great focus on running code, most of which open source with some upstream contirbution of the MPTCP code).

List of Authors

Authors	Saverio Niccolini, Felipe Huici, Olivier Bonaventure, Gino Carrozzo, Giacomo Bernini, Nicola Ciulli,
	Marcelo Bagnulo Braun, John Thomson, Julian Chesterfield
Participants	NEC, UCL-BE, NXW, UC3M, ONAPP
Work Package	D.4.3 - Public record of the Trilogy 2 exploitation plan and its standardization workshop
Security	Public (PU)
Nature	R
Version	v0.1
Total number of pages	36

Contents

Ex	ecutiv	e Sumn	nary	3
Lis	st of A	uthors		4
Li	st of F	igures		7
1	Expl	oitation	of Trilogy2 Foreground	8
	1.1	Multip	Path TCP	9
		1.1.1	Purpose	9
		1.1.2	T2 Foreground	9
			1.1.2.1 Upstreaming Multipath TCP into the main Linux kernel	10
		1.1.3	Potential impact	11
		1.1.4	Traction	11
			1.1.4.1 Multipath networks	11
			1.1.4.2 VRT	11
			1.1.4.3 Apple	11
			1.1.4.4 Tessares	11
			1.1.4.5 Intel	12
			1.1.4.6 Gigapath	13
			1.1.4.7 OVH	13
		1.1.5	Further Work	14
	1.2	Federa	tion and Cloud.net	15
		1.2.1	Purpose	15
		1.2.2	T2 Foreground	15
		123	Potential Impact	15
		12.3	Traction	15
		1.2.4	Further Work	16
	13	Minice		17
	1.5	131		17
		1.3.1	T apose	17
		1.3.2		17
		1.3.3		17
		1.3.4	Further Work	17
	1 /	Troper		10
	1.4			10
		1.4.1		10
		1.4.2		10
		1.4.5		19
		1.4.4		19
	15	1.4.3		19
	1.5			20
		1.5.1		20
			1.5.1.1 Purpose	20
			1.5.1.2 12 Foreground	20
			1.5.1.5 Potential impact	21
		150	1.5.1.4 Iraction	22
		1.5.2		23
		1.5.3	SUN and NFV MSc programme	23
			1.5.3.1 Purpose	23
			1.5.3.2 T2 Foreground	24
			1.5.3.3 Potential Impact	24
			1.5.3.4 Traction	24

tril**o**gy 2

			1.5.3.5	Further Wo	ork							 	 	 		. 24
		1.5.4	Telefoni	ca								 	 	 		. 24
			1.5.4.1	Purpose .								 	 	 		. 24
			1.5.4.2	T2 Foregro	ound							 	 	 		. 24
			1.5.4.3	Potential I	mpact .							 	 	 		. 24
			1.5.4.4	Traction .								 	 	 		. 25
			1.5.4.5	Further Wo	ork							 	 	 	•	. 25
2	T2 st	tandardi	zation wor	kshop												26
	2.1	Target	and scope	of the work	shop .							 	 	 		. 26
	2.2	Attend	lees									 	 	 		. 26
	2.3	Agend	a and Disc	cussion								 	 	 		. 26
	2.4	Outcon	me		• • • •							 	 	 	•	. 27
3	Diss	eminatio	on Activitie	S												28
	3.1	Public	ations - Ye	ear 1 - 2013								 	 	 		. 28
	3.2	Public	ations - Ye	ear 2 - 2014								 	 	 		. 28
	3.3	Public	ations - Ye	ear 3 - 2015 a	and beyo	ond.						 	 	 		. 29
	3.4	Presen	tation at e	vents, standa	rdizatio	n activ	ities,	press	s rele	eases	5.	 	 	 		. 30
	3.5	Softwa	are									 	 	 		. 31
	3.6	Summ	ary and H	ighlights of I	Entire Pi	roject						 	 	 	•	. 34
4	Con	clusions	i													35
Re	References 30									36						

List of Figures

1.1	Hybrid Access Networks with Multipath TCP	12
1.2	Architecture of the Gigapath service	13
1.3	Licence tree of VM3 components	18
1.4	VNF Pool Enabled Virtual CPE reference scenario	20
1.5	VNF Pool Enabled Virtual CPE components	21
1.6	Potential market interest on Virtual CPE solutions (source SDNCentral Virtual Edge Report	
	2015)	22

1 Exploitation of Trilogy2 Foreground

This section reports on the several exploitation activities of the Trilogy2 partners. It is noteworthy that several activities are supported by running code that has been published as open source, refer to the dissemination section for details on the software produced by Trilogy2 partners.

1.1 Multipath TCP

1.1.1 Purpose

The open-source implementation of Multipath TCP in the Linux kernel¹ is one of the software developed within the project. It is considered as the reference implementation for the Multipath TCP protocol [1].

1.1.2 T2 Foreground

This implementation has been released under the GPLv2 license that covers the Linux kernel and a lot of effort has been made to provide support to users and port it over different kernel releases. During the project lifetime, the following releases have been produced and distributed :

- version 0.86 for Linux kernel 3.5.7. This version, released in March 2013, included the following features
 - Creation of new subflows
 - Suport for the LIA and OLIA congestion control schemes
 - Handing traffic over from one subflow to another (resilience to failure)
 - Adding/Removing IP-Addresses for mobile hosts
 - Support for most middleboxes (segment splitting, coalescing, payload-rewriting)
- version 0.87 for Linux kernel 3.10. This version, released in July 2013, included the following features :
 - Hardware-offloading support (TSO/LRO)
 - TCP zero-copy support (sendfile and splice)
 - Support for NET_DMA
 - Several performance improvements
 - Support for NFS
 - Improvements to Middlebox support
 - Support for TCP-urgent pointer
- version 0.88 for Linux kernel 3.11. This version, released in October 2013, included the following features :
 - Support for the TCP_SYNCOOKIES and TCP_MD5SIG options
 - Modularized path-manager. The path-manager is now a loadable module, which allows to have a better-structured code and it is more easy to extend/add a path-manager. Three path managers were included : *default*, *fullmesh* and *ndiffports*
- version 0.89 for Linux kernel 3.14. This version, released in August 2014, included the following features :
 - Support for the wVegas congestion control scheme
 - Support for a new socket API to enable/disable MPTCP on a per-socket basis
 - Support for TCP keepalives
 - New binder path manager
 - Modularised scheduler infrastructure that allows to integrate different schedulers for different use-cases.
 - Proper support of TCP small-queues (TSQ) and TCP tail-loss probe (TLP)
 - Various performance improvements.
 - Code cleanup

¹Available from http://www.multipath-tcp.org

- version 0.90 for Linux kernel version 3.18. This version, released in September 2015, included the following features :
 - Support for the Balia congestion control scheme
 - TCP Fastopen support for Multipath TCP
 - More socket options (e.g., TCP_CONGESTION)
 - Support for stateless connection establishment (aka TCP SYN-Cookies)
 - Addition of MIB/SNMP counters for statistics-tracking and debugging
 - Various performance improvements and bug-fixes.
 - Code cleanup

The releases were mainly developed based on contributions from UCL-BE and Intel and input from external developers who have proposed extensions to the Multipath TCP implementation. The availability of a reference implementation that has been maintained and supported during the last three years has contributed to the adoption of Multipath TCP by various companies. Since Multipath TCP is open-source, we cannot track all the utilisations of the code developed within this project, but the section below summarises some of the known usages.

1.1.2.1 Upstreaming Multipath TCP into the main Linux kernel

In order to gain even more visibility and coverage, the MPTCP kernel implementation must be integrated upstream. This subsections briefly describes the challenges that faced in trying to integrate MultiPath TCP (MPTCP) in the official Linux kernel repository. This is a non-trivial task because the MPTCP implementation is invasive and mixes the MPTCP and TCP code and thus substantially increases the complexity of Linux's kernel TCP/IP implementation.

We present several changes done to the MPTCP implementation that move the MPTCP code in a separate layer with the purpose of managing the complexity of the kernel implementation. To do so a series of problems must be solved: extract the MPTCP specific code from the TCP hotpath, create an MPTCP specific layer on top of TCP, pass MPTCP specific options to/from the TCP layer and have an efficient TCP fallback mechanism.

With the purpose of separating the MPTCP and TCP layers in mind, on obvious problem was that lot of the MPTCP code added relied on a few conditions that differentiated the two layers: if the current socket was a meta socket, if the other end was MPTCP-capable or if the current socket is a master socket. This means that, inside the TCP code, when one of the conditions above was true, the flow would be passed to an MPTCP function. This was solved by removing these conditional statements by doing the following: for each condition that could pass the control from TCP to MPTCP functions, create a suite of functions pointers in the tcp_sock kernel data structure, each pointer corresponding to a place where that condition would be checked. Afterwards, whenever that flag/condition changed its value, change the function pointer in tcp_sock to its TCP or MPTCP function. In order to isolate the MPTCP code even more, we created MPTCP specific request sockets as well as connection request operations. This allowed us to reduce request_sock to its original size. It also allowed us to remove the MPTCP code from the TCP stack part that deals with accepting new connections since we can now parse the MPTCP options and do the MPTCP checks in the MPTCP specific connection request operations after which we can call the TCP specific connection request operations. As a generalization of these changes we also removed code duplication between IPv4 and IPv6 connection request functions. This independent changes were positively received by the Linux kernel networking community and they were merged upstream.

The dedicated MPTCP Layer takes form by implementing a new protocol, IPPROTO_MPTCP that represent the MPTCP protocol. This way, we can implement MPTCP specific operations cleanly and separately from the TCP layer. Following this design, meta socket is created as soon as the socket is created. With this approach, the application needs to specifically open an MPTCP connection. However, it should be possible to redirect TCP socket creation calls to MPTCP socket creation calls in inet_create(), based on, for example, sysfs settings. It also means that fallback to TCP must be handled inside the MPTCP layer itself by forwarding calls from the MPTCP layer to the TCP layer. Starting with an MPTCP connection by default and allocating specific resources for it from the beginning gives us the possibility to make a clearer separation between the master socket and the meta socket and remove the need for cloning the meta-socket. In short, those changes enabled a simplified connect path and a whole new receive path.

1.1.3 Potential impact

Multipath TCP already has an impact on commercial products as discussed in the section below.

1.1.4 Traction

The impact of the Multipath TCP implementation in the Linux kernel already goes beyond the Trilogy2 partners as several commercial deployments and utilisations of this implementation have already started. Here are some of the already known ones.

1.1.4.1 Multipath networks

Multipath networks² is an Irish company that was the first to use Multipath TCP to bond two ADSL links or an ADSL and a wireless link. Their initial product relied on a modified home router that used the Multipath TCP kernel together with OpenVPN and an HTTP proxy. The router intercepts all TCP traffic, sends it to a server running in the cloud over Multipath TCP and the server uses regular TCP to reach the final destination. Unfortunately, the company went bankrupt and the service is not sold anymore as of this writing.

1.1.4.2 VRT

VRT is the Flemish TV broadcaster in Belgium. They have designed their own cars to allow web journalists to capture videos, edit them and upload them to the VRT head quarters. Videos are long files that require a large bandwidth to be uploaded quickly. To allow the journalists to send their video reports as quickly as possibly, the latest VRT car, called "The Beast"³, has been equipped with three types of antennas :

- one satellite antenna
- several 3G antennas with the corresponding SIMs
- several WiFi antennas

Once a video is ready, the server running in the car automatically starts all the available network interfaces and combines them thanks to Multipath TCP to upload the entire video to the VRT head quarters. This car has been used in production for more than a year at VRT.

1.1.4.3 Apple

Apple has started to use Multipath TCP on iPhones and iPads in September 2013 to support the Siri voice recognition application. Thanks to Multipath TCP, these mobile devices can better cope with losses and connectivity problems over the wireless interfaces. This deployment uses an implementation written by Apple's engineers that is now also included in MacOS. Apple's implementation of Multipath TCP does not include all the features of the protocol defined in [1] but it is fully interoperable with the Linux implementation.

1.1.4.4 Tessares

Tessares⁴ is a spinoff from UCL-BE that was created with funding from Proximus, the Belgian network operator and the VIVES investment fund. Its objective is to develop new innovative network services on top of Multipath TCP. The first product developed by this company is a solution for Hybrid Access Networks. Such an access network combines two different types of technologies, typically DSL and 3G/4G. It is illustrated in figure 1.1.

This solution is composed of two different network devices :

- The Hybrid CPE (HCPE)
- The Hybrid Aggregation Gateway (HAG)

The HCPE is a CPE device that is capable of using two separate access networks. It is typically a home router that has been extended with a 3G/4G interface. Tessares provides a tuned version of the Multipath TCP implementation in the Linux kernel that has been optimised for this platform. It also includes a Multipath

³See http://deredactie.be/cm/vrtnieuws/videozone/nieuws/cultuurenmedia/1.1961302

⁴See http://www.tessares.net



Figure 1.1: Hybrid Access Networks with Multipath TCP

TCP proxy that intercepts the TCP connections established by the devices in the home network and converts them into Multipath TCP connections. Thanks to the utilisation of Multipath TCP, the devices used in the home network can use both the DSL and the 3G/4G network. The Hybrid Aggregation Gateway terminates the Multipath TCP connection and converts them into regular TCP connections so that regular servers that have not been upgraded to support Multipath TCP can be contacted.

Tessares leverages both the Multipath TCP implementation developed within the Trilogy, CHANGE and Trilogy2 project as well as the Multipath TCP proxy that was designed within the CHANGE project.

The BroadBand Forum is working on several solutions to support Hybrid Access Networks. During the last Broadband World Forum in London, several companies have demonstrated solutions that include the Multipath TCP implementation in the Linux kernel : Tessares, SoftatHome, Sagemcom, Technicolor, Intel and Ericsson. Alcatel is also working on a solution based on Multipath TCP.

1.1.4.5 Intel

Intel and Ericsson are collaborating to help telecom carriers deliver higher speed broadband connectivity to more people and underserved regions by combining existing wired (Cable, DSL) and wireless network (LTE) infrastructure. ⁵.

In demonstrations at Broadband World Forum 2015, the two companies are showing an end-to-end access solution based on Multipath Transmission Control Protocol (MPTCP). The hybrid solution aggregates the bandwidth of both DSL and LTE transmission links to multiply data rates and create a single, powerful broadband pipe between the network service node and a customer's broadband gateway / router.

Ericsson is adding the protocol to the feature set of its MultiService Proxy (MSP) network node. Concurrently, Intel is incorporating the technology in a new DSLTE Home Gateway reference design based on the AnyWAN GRX 350. Intel AnyWAN platforms can be easily configured to support any incoming WAN data link, including all variants of Cable, DSL, FTTx and LTE. That simplifies design and implementation of product families for OEM/ODM customers and helps ensure that carriers can source solutions from multiple Customer Premises Equipment (CPE) vendors.

⁵See http://blogs.intel.com/technology/2015/10/dsl-lte-hybrid-access-ericsson-intel/



Figure 1.2: Architecture of the Gigapath service

1.1.4.6 Gigapath

Gigapath is a commercial service that was launched during the summer 2015 by Korean Telecom⁶. In Korea competition among network operators forces them to provide higher bandwidth mobile services. The cellular networks deployed in this country are among the fastest in the world, but this is still not sufficient. Gigapath allows smartphone uses to combine together their 4G and WiFi networks to reach bandwidths of 800 Mbps and more.

From a technical viewpoint, the solution deployed by KT combines Multipath TCP and the SOCKS protocol (see figure 1.2. Korean Telecom has convinced Samsung and LG Electronics to port the open-source Multipath TCP implementation in the Linux kernel on their high-end smartphones. As of December 2015, there are about half a dozen different smartphone models from these two vendors that include Multipath TCP. Each smartphone also includes a SOCKS client that intercepts all TCP connection establishments and redirects them to a SOCKS proxy running on one server managed by Korean Telecom. The SOCKS proxy uses the Multipath TCP implementation in the Linux kernel and terminates the Multipath TCP connection.

In July 2015, 5,000 users had subscribed to the Gigapath service. In November 2015, there were about 20,000 users.

1.1.4.7 OVH

OVH is a French cloud company that also provides DSL services. In September 2015, they announced a new product called "Overthebox"⁷. This product combines Multipath TCP and SOCKS proxies to enable users to bond different DSL lines together. In contrast with the SOCKS-based solution deployed by KT, OVH did not modify the enduser devices. Instead, they provide a device that is attached to the different DSL routers that need to be combined. This device acts as the default gateway in the home network and serves as the DHCP server. Its SOCKS client can then intercept all established TCP connections and convert them into Multipath TCP towards a SOCKS server running in the cloud. The SOCKS server terminates the Multipath TCP and creates a regular TCP connection to the final destination. In December 2015, OvertheBox is still in beta phase, but more than 300 users already participate in the beta and the commercial deployment is expected soon.

⁶See https://www.ietf.org/proceedings/91/slides/slides-91-mptcp-5.pdf

⁷See https://www.ovhtelecom.fr/overthebox/

1.1.5 Further Work

The examples of the previous sections show that Multipath TCP and its reference implementation in the Linux kernel already have an impact. The design of the protocol has been heavily influenced by the presence of various types of middleboxes on the current Internet and it is interesting to note that most of the commercial deployments use Multipath TCP on some forms of middleboxes. This was not initially expected since Multipath TCP was designed as an end-to-end protocol. This indicates two important directions for future work :

- Design and develop the tools that enable Multipath TCP to be used on both clients and servers. Currently, one of the main difficulties on the server side is to develop load balancing techniques that are compatible with Multipath TCP. Security tools like intrusion detection or firewalls that often protect servers also need to be reconsidered with Multipath TCP in mind.
- Design, realise and deploy an Internet architecture that exposes the middleboxes like routers are exposed in today's architecture and enable applications to efficiently interact with different types of middleboxes.

1.2 Federation and Cloud.net

1.2.1 Purpose

The OnApp Federation is a marketplace that allows global Cloud sites to pool resources. OnApp provides the ability for its customers to buy and sell within a Business-to-Business (B2B) marketplace. Within the OnApp UI there is an aggregate view of all available resources for a Buyer to purchase and resell. Within the same UI, a Seller can provide their local compute resources for sale.

Cloud.net is a virtual service provider that utilises the OnApp marketplace. Cloud.net buys and trades Virtual servers on a global scale directly to end-users as a retail channel. Cloud.net provides a single view for a user of the platform and allows the purchasing of resources. Users of the platform can buy compute (VM) resources around the world.

1.2.2 T2 Foreground

The initial research and development work that went into the Federation was supported by work carried out in Trilogy2. The testbed platform that formed the early basis for the Federation was proposed and built during the early stages of Trilogy2 and involved a number of partner sites that were configured to share resources. User roles and accounts were part of the set of primitives that were extended in order to support trading of resources in the wide-area. The operations that were allowed in between Cloud sites was first deployed on the testbed platform to ensure that different cloud sites could support them.

Cloud.net is a virtual service provider offering that allows for full exploitation of liquid resources to be combined with the business case. The control of liquidity through incentives and enforcement from Trilogy2 was tested through Cloud.net.

In line with the objective of making source code available where it doesn't interfere with the commercial exploitation, the dashboard of Cloud.net has been open sourced. The source code is available at https://github.com/OnApp/cloudnet.

1.2.3 Potential Impact

As described in the next section a commercial realisation has already been achieved for both the Federation and Cloud.net. SMEs and small cloud providers can benefit from a Federation platform. The barrier to entry for creating a fully-fledged global Cloud platform is great and limited to existing multi-national, technology companies. Cloud and service providers that want to extend their reach into other markets normally have to go through large CDN or platform provider solutions such as Akamai or AWS. Although there is competition in this space the Federation offers a real alternative that can help smaller sites provide similar services through contracts with smaller providers in the regions that they are interested. SME's can benefit in the lowering of prices and by working with specific providers in other markets can create bespoke solutions that are more suitable for their needs.

OnApp as an SME also benefits from providers utilising the Federation platform as it means more customers will utilise the core products from which revenue can be generated. Secondary revenue streams are also improved through the increased utilisation of associated services such as the CDN platform.

1.2.4 Traction

The Federation (http://www.onapp.com/federation) and Cloud.net (http://www.cloud. net) are real-world commercial exploitations of the concepts investigated in Trilogy2. These OnApp commercial products demonstrate the business case behind the LiquidNet. The findings have been interesting and reported in the associated Deliverables within Trilogy2.

The platforms are fully supported by OnApp and ongoing work is being carried out to ensure that the Liquid-Net concepts are carried forward. Some of the larger platform providers are utilising the Federation to build private, multi-site cloud offerings that connect two or more Clouds together for increased manageability and visibility into the consolidated resources.

OnApp has 3,000+ small to large service provider customers globally that captures about \$500M of retail value compute. The Federation has a presence in over 170 cities and is the largest with respect to compute, bandwidth and locations. The Federation also supports transactions between over 500 service provider buyers and sellers.

1.2.5 Further Work

Customers are now utilising the Federation to offer services that they would not be able to offer without Liquid resources. The ability to further extend offerings by utilising other companies hardware and infrastructure and existing business relationships in different regional locations has been of interest to cloud service providers, especially as a way for expanding into new regions without having the up-front costs normally associated with such an expansion.

The Federation has had uptake from customers but further work is needed to further incentivise usage of a LiquidNet from the buyer side. Commercially the selling side of Liquid resources has been a relatively simple value proposition for cloud resource owners but the buying has been the complex part. Through customer feedback on Cloud.net it has also been clear that end-users are not so interested in fine-grain statistics and metrics as would be expected. The service providers have used the metrics and found it useful for differentiating and checking their offerings with competitors. End-users though would much prefer certain performance tiers that are then guaranteed and monitored by a third-party such as OnApp. As such OnApp are looking to re-focus the Cloud.net offering into tiered solutions and monitoring the cloud service provider sites to ensure that they meet certain SLA and QoS/QoE performance criteria.

Research work has also been carried out into the viability of using the Federation to power small, developer test workloads through the Cloud.net platform. The results so far have demonstrated that the platform could be used for this new type of workload. In addition the Federation will be used as a basis for powering thousands of enterprise hybrid customers.

Services such as Disaster Recovery as a Service (DRaaS) are based on the Federated model and allow for new business propositions, see D3.3. Continual block migration over the wide-area allows a site to become available again after an outage on the primary site by starting the services on the fail over site. Other services that rely on the Federation are being planned, including VM migration over the wide-area that was initially worked on in the course of Trilogy2. As the platform becomes more mature, additional offerings and services will be developed, depending on what customers need.

Part of the work of Trilogy2 is being continued in the H2020-ICT-2014-2 Superfluidity project (Grant agreement 671566). The Trilogy2 concepts of the LiquidNet are being taken forward in the form workload mobility. Further exploration is being taken to see how workloads can be migrated across different computer instruction set architectures and be placed anywhere in the network. Some workloads will benefit from workloads moving to the edge whereas for others it may be beneficial to move workloads closer to the core. This research work will look more into the topology mapping and placement of 'Fluid' resources.

1.3 Minicache

1.3.1 Purpose

Minicache is a unikernel (i.e., minimalistic VM) running both under XEN and KVM that implements caching features for both VoD and live streaming. Minicache enables the construction of *ephemeral CDNs*, virtualized, on-the-fly CDNs that can be popped up on demand (e.g., for a single live stream such as a sporting event) wherever there are cloud-like deployments (e.g., cloud providers such as Amazon's AWS, micro-datacenters or OnApp's federated cloud).

1.3.2 T2 Foreground

Beyond virtual CDNs, the development of Minicache has also led to technologies that allow us to easily put together purpose-built VMs that can be instantiated in a few milliseconds while handling potentially high throughput (10Gb/s or higher).

1.3.3 Potential Impact

We believe Minicache could bring benefits to a number of players. End users would see improvements in video quality (e.g., better join times and buffering ratios) thanks to the very quick instantation times coming from the minimalistic and lean footprint of minicache; CDN operators could expand their service to cope with fluctuations in load and demand from specific regions by dynamically building a vCDN to complement their existing infrastructure; and network operators would derive additional revenue from acting as CDN operators and from renting out their infrastructure.

1.3.4 Traction

VMs with lean operating systems (also known as unikernels) are gaining traction (along with containers) in the industry. The development of Minicache helped NEC to have a good basis for this paradigm shift and provides a step forward in the unikernel space with respect to ClickOS/MiniOS. Minicache features a full fledged TCP/IP stack that is going to provide the basis for unikernel-based middleboxes that require such a stack (e.g., TCP proxies as opposed to stack-less middleboxes like traffic shapers or load balancers).

NEC's solutions in the space of traffic management and acceleration require such features: for example, NEC currently commercializes the Media optimizer, http://www.nec.com/en/global/solutions/nsp/tms/tms03.html, which requires TCP optimizations.

The possibility of providing agile TCP proxies that can be instantiated on-the-fly for accelerating single connections is something that Minicache can also enable.

1.3.5 Further Work

We're currently in the process of further optimizing Minicache on KVM, as well as conducting benchmarking for it on other single-board computers besides the Cubietruck. We're also investigating the possibility of reusing the TCP stack enhancements carried out as part of the Minicache work in order to build virtualized TCP proxy accelerators that can be instantiated on the fly.

1.4 Transparent Migration

1.4.1 Purpose

Virtual Machine Migration Manager (vM3) enables fast transparent migration of mobile application between devices. Applications are fired up in specialized environments build upon tiny virtual machines. The hypervisor of those virtual machines is the combination of KVM and Qemu. This assures reasonable performance and responsivness on Intel Mobile devices like tablets and phones. The main ingredient of transparent migration is the use of MPTCP's break-before-make features, that allows the applications TCP connection to smoothly survive the transfer.

1.4.2 T2 Foreground

The Virtual Machine Migrator Manager (vM3) application was researched and developed entirely during Trilogy2. One of the main purposes was to build a tool that would enhance user experience on Intel's mobile devices, while taking advantage of the MPTCP technologies.

To be inline with the goals of the Trilogy2 project, vM3 will be released as open-source. In order to release the code as open source, it must have a licence that is compatible with all internal components. Figure 1.3 shows the components of VM3 from a licensing point of view.



Figure 1.3: Licence tree of VM3 components

VM3 is an Android Application that enables running and migrating other applications as virtual machines. The migration part is done by taking advantage of the open source hypervisor Qemu. But an application is not of much use without graphical experience. Here is where the SDL library glues the video and audio of the VM back into the Android application. And, of course, being an Android Application, VM3 relies on the Android Framework.

The challenge is to establish the correct licence for VM3 and make sure that all components can coexist from a licensing point of view. Since zlib is a very permissive licence, this is practically is equivalent to have both GPL v2 and Apache v2 licences in the same project. Apache v2 is a more permissive licence that enables contributors to modify and re-licence a code. But GPLv2 is more restrictive. It has the characteristics of a copyleft licence, hence it obliges the user to not only release it's own code as open-source, it must also be licensed as GPLv2 (or compatible).

The solution is to release VM3 app as a GPLv2 application that uses an Apache v2 library - the Android Framework. As a result of the whole process, the code will be published on well known open-source sites

(like Github) and also on the OTC website - 01.org.

1.4.3 Potential Impact

This technology is to be open-sourced in order to enable easy adoption and to provide a base for further development. Also, as a result of the virtualization features that power up this technology, transparent migration of mobile device application will consolidate Intel's presence on the mobile market. Since Intel devices work better with other Intel devices, this impact the whole ecosystem, since applications can be migrated from phone/table to desktop/laptops and back.

Currently VM3 is ready to be tested directly by users as a beta version of the final project, but is far from a potential product. Instead of investing in transforming VM3 in classic closed source product, Intel intends to harness the power of open-source.

1.4.4 Traction

vM3 application has not been released yet, due to ongoing the open-source process described earlier. Nevertheless, besides Trilogy2, vM3 has been present in a number of internal Intel presentation and demo sessions. vM3 will be released open-source entirely. The goal of publishing this work is to provide better answers to real-world use-cases and also to speed up the adoption of the MPTCP protocol. The decision to make VM3 open source is a natural one since the team in Intel that developed it is part of Intel Open Source Technology Center.

1.4.5 Further Work

Further development and research commitment on this project depends on the tractions gained after the release. Besides the immediate step of open-source release, there is a medium-term plan of maintenance and updates.



Figure 1.4: VNF Pool Enabled Virtual CPE reference scenario

1.5 NFV

NFV topics have been widely addressed by Trilogy 2 through heterogeneous project activities, spanning from architecture definition to participation and contribution to standards, as well as to implementation of new liquidity tools. Several Trilogy 2 partners, including Telefonica, Nextworks, UC3M and BT have been active in carrying out NFV related activities during these three years. The following subsections provide an overview of the Trilogy 2 exploitable NFV outcomes.

1.5.1 VNF Pool enabled virtual CPE

1.5.1.1 Purpose

Operators are looking at their future DCs and PoPs as more dynamic infrastructures through which they can provide service differentiation and customization via on-demand self-service resource composition tools.

High elasticity to scale up and down network resources and automation are a key requirement NFV solutions are targeting today. Network functions traditionally integrated in hardware middleboxes at customer premises are progressively migrated to Virtual Network Functions running on servers.

The virtual CPE (vCPE) developed in Trilogy 2 by Nextworks targets the control and coordinated deployment and configuration of network resource pools for routing to allow user-driven on-demand extension of IP/MPLS VPN services into virtual hosting environments. With vCPE, user's VMs created in the IaaS platform of the cloud provider can be dynamically linked to the private IP/MPLS VPN service of the user.

1.5.1.2 T2 Foreground

In T2, Nextworks has developed a VNF Pool Enabled Virtual CPE prototype which specifically targets the Operator's datacenter scenario.

The Nextworks' vCPE is basically an additional VM under the control of the operator, deployed on demand to dynamically configure the customer VPN memberships and Virtual Routing and Forwarding (VRF) instances. Upon the request from a customer to buy some IaaS resources and include them into a private cloud accessible through its Layer 3 VPNs already in place, the vCPE is deployed on-demand as part of the IaaS to extend this VPN. This is implemented by running Open virtual Switches (OVS) in the data centre hypervisors, which allow to perform VPNs and IaaS service chaining via software, properly managed and controlled by an SDN and NFV enabled orchestration framework.

Key outcomes and components of the Nextworks' vCPE prototype are:

• OpenStack Cloud Management Platform



Figure 1.5: VNF Pool Enabled Virtual CPE components

- SDN controller (one for the whole datacenter) with a set of network applications to provide on-demand flow and routing tables configuration at the OVS and datacenter edge router
- Virtual CPE routing protocols engine (BGP, OSPF, etc.)

Details on the implemented vCPE functionalities are provided in deliverable D1.4.

1.5.1.3 Potential Impact

In a recent market survey done by SDNCentral LLC in 2015, majority of service provider indicated their intention to move to a virtualized and/or SDN-based architecture for their WAN connectivity solutions:

- 25% of the service providers who indicated they expect to deploy this new architecture plan to do so in the next 6 to 12 months;
- 25% expect to enter production in the next 12 to 18 months;
- 11% are currently deploying or expect to deploy vCPEs in the next 3 to 6 months.

The majority (53%) are planning on rolling out CPEs that provide virtualized applications that fit into the scenario targeted by T2 solution. An interesting 47% of service providers are exploring virtualizing the customer edge with a thin CPE with vCPE applications running at the customers Point of Presence (POP) or central office (see Figure 1.6).

In this market scenario, a number of vCPE offers by big players like Brocade, Cisco, HP, Ericsson, etc. have just targeted the market.

Nextworks sees the value of its vCPE proposition mostly into three major target areas for business



Figure 1.6: Potential market interest on Virtual CPE solutions (source SDNCentral Virtual Edge Report 2015)

- *SDN/NFV training*. The interest towards SDN/NFV solution is generating a sustained demand for hands-on training on SDN and NFV tools (architectures, controllers, orchestrators) by both network operators who need to define their service architecture and vendors/integrators who need to shorten the time to enter the market of SDN/NFV products with their teams.
- *NFV solution development as 3-rd party partner*. Nextworks is turning the know-how generated in T2 into a valuable offer of software development consultancies to vendors and/or system integrators (as 3-rd party company), to complement customers' internal development teams in rapid and agile development of NFV solutions. This market is taking momentum mostly due to the technical talent crunch currently and the huge demand for specific coding and technical skills needed for the development of an NFV product from scratch.
- *Evolution of the Nextworks' network product portfolio*. Nextworks has an offer of multi-face router and firewall (NXwall) part of the Symphony suite for Smart homes and buildings. The vCPE developed in T2 can be used to evolve the NXwall appliance into a virtualized function to be used in large scale Symphony installations.

1.5.1.4 Traction

The Nextworks' VNF Pool Enabled Virtual CPE foreground is currently exploited by the company in the three market areas identified above. Summary highlights are:

- Area: KnowHow-Training.
 - Nextworks has delivered 3 expert-level training sessions on SDN/NFV to top customers in 2015 (1 big network vendor, 1 system integrator and consultancy firm, 1 Italian public research center), aimed at training network professionals on SDN/NFV architectures, tools, and software.
- Area: KnowHow-Consultancies.
 - Nextworks has won a contract with the European Space Agency (ESA) for the provisioning of Engineering Services on SDN/NFV technologies in the time-frame 2015-2017.
 - Nextworks has renewed contracts for a Senior Network Architect service through which experienced network and security engineers support design and evolution of datacenter, security and managed hosting services of the Earth Observation program of the European Space Agency (ESA). Focus of this work is on virtualization platforms, IaaS, and virtual network functions.
 - Nextworks has started discussions on industrial projects with two vendors (names covered by NDA) to be executed in 2016 for the implementation of NFV Management and Operation tools and a Transport SDN controller.

- Area: Products.
 - Nextworks R&D has started working on the virtualization of the NXwall function integrated into the Symphony product, based on the know-how generated in T2.
- Area: Standards.
 - Nextworks R&D has started working in NVF Research Group at IRTF with an Internet Draft on VNF Orchestration For Automated Resiliency in Service Chains (see draft-bernini-nfvrg-vnforchestration), co-authored with Telefonica, which is contributing to shaping NFV RG proposition on Policy based Resource Management.

1.5.2 Further Work

Based on T2 foreground, further work and research is planned on the NFV and VNF Pools. In particular,

- Nextworks has planned to work on consolidating the VNF Pool Enabled Virtual CPE in terms of data plane performances.
- Nextworks has started working (supported by other H2020 EU projects in 5GPPP program) on the extension of the VNF portfolio towards security functions (vFW, vIPS) and flexible chaining.
- Nextworks will continue the work in IRTF NFVRG trying to converge into an Informational RFC in 2016.

1.5.3 SDN and NFV MSc programme

1.5.3.1 Purpose

As result of the knowledge acquired through the Trilogy 2 project, UC3M, in collaboration with Telefonica and other industrial partners (Ericsson in particular) has launched two new graduate degrees in Network Function Virtualization (NFV) and Software Defined Networking (SDN). UC3M is offering two new post-graduate degrees: a MSc degree in NFV/SND of 14 months of duration and a Specialist degree in NFV/SDN of 7 months of duration.

These programes will enable IT professionals to apply the NFV and SDN technologies to the design and operation of telecommunication networks. The high level goals of the programes is to allow students to:

- Design network elements and network services for NFV and SDN,
- Plan and operate networks based on NFV and SDN technologies
- Understand the evolution of network products and solutions embracing the new NFV/SDN paradigm.

The courses included in the MSc and specialist programs are the following:

- Virtualization Technologies for NFV
- Cloud Management and Orchestration for NFV
- Data Center Communications
- Network Function Virtualization (NFV)
- Software Defined Networks (SDN)
- SDN, NFV and Virtualization Lab
- Applications of Software Based Networking Technology to Real Networks
- Software Based Network Security
- Business Models for Software Based Networks
- Operation and Management of Software Based Networks

- Advanced Laboratory of Software Based Networks
- Internship
- MSc Thesis

For the MSc degree must take both years. The first year provides a specialist degree.

The MSc and specialist degree are oriented to the needs of the industry and as such half of the professors teaching the courses are from industry. Among them, we can find partners from the Trilogy 2 project, in particular Diego Lopez and Luis Contreras from Telefonica. Other professors come from Ericsson, EMC, Dell and UC3M.

1.5.3.2 T2 Foreground

The courses included in the programe cover several topics mastered during the Trilogy 2 project execution. In particular, the programe includes courses on NFV, virtualization, cloud, SDN and data center networking, which have been deeply developed in the project.

1.5.3.3 Potential Impact

We expect that the impact of these graduate programes will be significant in preparing the industry and the work force for the new generation networks that will adopt these new paradigm of programmable networks. To the best of our knowledge, this is the first graduate degrees on these topics in the world.

1.5.3.4 Traction

We have launched the programe early in 2016 with a lot of success. The programe has started on november 2016 and have sold out all the seats offered (40 seats). Out of the 40 seats offered, 38 are taken by people from the industry i.e. people who are already working and are looking for learning about these new technologies because their work requires them so.

1.5.3.5 Further Work

We are already planning the second edition and just a month after starting the first edition, we already have 20 seats sold for the next edition.

1.5.4 Telefonica

1.5.4.1 Purpose

Network Function Virtualisation(NFV) involves the move of many network functions from running on specialist hardware to instead running as software on commodity hardware. Many people believe it is the most important way in which the computing industry is converging with the telecommunications industry. NFV technologies enable the teleco operators to move network functions to commodity software. This should lower CAPEX and OPEX and provide more agility in Service Provision.

1.5.4.2 T2 Foreground

From an operators perspective, a key development has been standardisation at the ETSI NFV ISG, where BT and Telefonica have played foundational and leading roles. Trilogy2 has been instrumental for this participation. In addition, T2 has developed several technologies and virtual network function components (VNFCs) that are likely to help operator deployments. One example is that Telefonica has developed the MPLS-enabled control plane, which has been shared with Nextworks in the context of Trilogy2 and been made available to the community through GITHUB.

1.5.4.3 Potential Impact

BT and Telefonica believe that NFV will have a major and transformational impact on the telecoms industry. It should lower CAPEX and OPEX. It should make it easier to scale up and to achieve resilience, through the simple activation new software instances. Several of the use cases make the equipment deployed on the customers premises much simpler (for example replacing routers with switches), whilst the software running the complex services is centralised where the operator can more readily access it if anything goes wrong. However, perhaps the main benefit is that it should provide more agility in Service Provision - new network services can be deployed more rapidly - at the speed of a software release, rather than the speed of a hardware roll out - and new services can be 'composed' from existing ones more easily.

The Telefnica Group is aiming at a comprehensive plan for network virtualisation⁸, which will include significant impact in CAPEX and OPEX savings through the use of NFV technologies in the Network. TID will play an essential role by transfering insights gained from Trilogy2 activities in this activities.

1.5.4.4 Traction

There is much interest in NFV, both inside BT and Telefnica, and with our vendors and our customers. There is an emerging ecosystem of VNF providers, which are delivering their products to operators, and customers who want the benefits it brings.

The work of both companies includes trials (jointly with vendors), market analysis, standardisation especially at ETSI NFV ISG, as well as study of the key remaining challenges.

BT already has an NFV product. As part of BT Cloud Connect we have deployed virtual Riverbed SteelHead application acceleration technology in the core of our global network, which provides superior application performance for customers accessing BTs own cloud services as well as other leading Software-as-a-Service (SaaS) offerings. It is the first time Riverbed technology has been deployed as an NFV service.⁹

1.5.4.5 Further Work

BT and Telefonica will continue our heavy involvement and participation in the ETSI NFV ISG. As the ISG evolves in time, we will seek the positions that will be most instrumental in the creation of a Telco-friendly ecosystem based on the liquidity concepts developed by T2.

BT has several on-going activities, including in the areas of virtual CPE, virtual CDN, virtual routers, NFV infrastructure, the merging of cloud and VPN services, and so on.

Telefnica will continue the development and deployment of the liquid control plane for the vCPE developed during Trilogy2 as a reference implementation for different NFV related activities we participate in. The Quagga control plane is an integral part of in the following virtual network functions (VNFs) that we plan to use in our NFV Reference Laboratory to benchmark and test vendor VNFs:

- Virtual Carrier-Grade Network Address Translator (vCGNAT)
- Virtual CPE
- Virtualised Provider Edge Router (vPE)

TID has started other follow-up activities that provide opportunities to exploit Trilogy2 results. These include projects of the 5G public-private partnership (5GPPP) and other projects in the scope of the H2020 framework programme as well as the recently created 5TONIC laboratory¹⁰. The scope of this reference laboratory created by Telefnica and IMDEA in Madrid is to provide a research ground for the upcoming 5G networks. As in the case of Telefnica's NFV reference laboratory, TID will use the VNFs developed in the scope of Trilogy2 to kick-start activities in 5TONIC.

⁸http://www.lightreading.com/carrier-sdn/sdn-equipment/telefonica-unveils-aggressive-nfv-plans/ d/d-id/707882

⁹http://www.mynewsdesk.com/uk/bt/pressreleases/bt-accelerates-cloud-of-clouds-with-new-hosted-p ¹⁰http://5tonic.com/

2 T2 standardization workshop

2.1 Target and scope of the workshop

The target standardization forum has been the IETF/IRTF (this is inline with what we had identified in the DoW).

We therefore decided to organize a special interim session of the IRTF Network Functions Virtualization Research Group (NFVRG) as we deemed the best in scope for disseminating the activities of Trilogy2. The workshop has been organized by NEC with the collaboration of Telefonica I+D at the NEC Laboratories Europe premises in Heidelberg, Germany, with the objective of disseminating Trilogy2 ideas and foster their adoption into a wider community.

2.2 Attendees

The workshop featured a number of attendees on site plus several ones joining remotely via a webex session. The total number of attendees has been about 20 with people representing several partners of Trilogy2 (Nextworks, UCLouvain, UC3M, NEC and Telefonica I+D) and other EU projects (i.e., project UNIFY, https://www.fp7-unify.eu/), with representatives from Ericsson. Additionally the meeting was attended by the IRTF chair (Lars Eggert, NetApp) and by the chairs of the NFV Research Group (Ramki Krishnan, Dell). Addiitonal participants were from companies like Citrix, Deutsche Telekom, Huawei, Sandvine and ZTE.

2.3 Agenda and Discussion

The full proceedings are available at https://www.ietf.org/proceedings/interim/2015/12/01/nfvrg/proceedings.html. The meeting minutes are available at https://www.ietf.org/proceedings/interim/2015/12/01/nfvrg/minutes/minutes-interim-2015-nfvrg-2.

The discussion focused on the presentation of work that has been carried out in the context (and with acknowledgment) of the Trilogy2 project:

- VNF orchestration and automated resiliency with VNFPOOL in service chains G. Bernini (Nextworks) joint work with Telefonica I+D
- Leveraging Segment Routing for Service Function Chaining D. LeBrun (UCLouvain)
- NFVI PoP network topology: Problem statement J. Brenes (UC3M)

The first contribution of the Trilogy2 project (VNF orchestration and automated resiliency with VNF-POOL in service chains, reference document available at https://tools.ietf.org/html/draft-bernini-nfvrg-vnforchestration-00) presented an orchestration framework for automated deployment of highly available VNF chains. Resiliency of VNFs and chained services is a key requirement for operators to improve, ease, automate and speed up services lifecycle management. The proposed VNFs orchestration framework is also positioned with respect to current NFV and Service Function Chaining (SFC) architectures and solutions.

The discussion on this contribution focused on the possible adoption of this work as RG item recognized by the RG as being an important piece of work fitting to the charter of the group. Additional discussion was made to possibly integrate this work with other work done currently int he group in order to increase the possibility of adoption.

The second contribution of the Trilogy2 project (Leveraging Segment Routing for Service Function Chaining) presented how to leverage IPv6 segment routing for achieving service chaining. The presentation reported about implementation and an initial evaluation.

The discussion on this contribution focused on the fact that there should be an update of the work in order to be considered for RG adoption as RG item. In order to be more useful for the community, the work should highlight and give guidelines and best practices on how to leverage IPv6 segment routing for achieving service chains instead of merely reporting on an implementation and its evaluation.

The third contribution of the Trilogy2 project (NFVI PoP network topology: Problem statement, reference document available at https://tools.ietf.org/html/draft-bagnulo-nfvrg-topology-00) presented considerations for the design of the interconnection network of an NFVI PoP. An NFVI PoP is (according to the ETSI NFV specifications) defined as a single geographic location where a number of NFVI-Nodes are sited where an

NFVI-Node is a physical device deployed and managed as a single entity providing the NFVI functions required to support the execution environment for VNFs. In other words, an NFVI PoP is the premises where the processing, storage and networking resources (i.e. servers and switches) used to execute the network virtual functions (VNFs) are deployed and pooled. The servers and switches in a NFVI PoP will be interconnected forming the NFVI PoP interconnection network. The goal of this presentation was to explore the different design considerations for the NFI PoP interconnection network topology, including design goals and constrains.

The discussion on this contribution focused on the fact that the community is interested in hearing more insights into the evaluation of these design considerations via simulations and analytical models. Once this is done, the discussion can focus on possible adoption of this work as RG item.

The content of the other discussion was related to possible merging of activities coming from the UNIFY project:

- VNF Benchmark-as-a-Service (VBaaS) R. Szabo
- Recursive Monitoring Language in Network Function Virtualization (NFV) Infrastructure C. Meirosu

The discussion on how to integrate such contributions with contributions from the Trilogy2 project and with the overall activity of the NFVRG will continue on the mailing list in the next months.

2.4 Outcome

Overall the workshop was well attended and stimulated interesting discussions on the impact of the Trilogy2 initiated contributions and there are realistically high chances that such work can be taken up by the IRTF as belonging to their charter of work, especially the work on VNF orchestration and automated resiliency with VNFPOOL in service chains. Concrete proposals were made as to how merge such work with current work being already adopted by the NFVRG. This workshop also gave the opportunity to disseminate to a larger audience sone of the Trilogy2 foreground and to discuss possible joint contributions between the Trilogy2 and the UNIFY project.

3 Dissemination Activities

This section lists all scientific (peer reviewed) publications relating to the foreground of the project divided by year of publication as well as a list of other dissemination activities (presentation at events, standardization contributions, press releases, software).

3.1 Publications - Year 1 - 2013

• Title: On the Benefits of Applying Experimental Design to Improve Multipath TCP

Where: ACM CoNEXT 2013

Authors: Christoph Paasch, Ramin Khalili, Olivier Bonaventure

• Title: Towards Minimalistic, Virtualized Content Caches with Minicache

Where: ACM HotMiddlebox 2013 (CoNEXT workshop)

Authors: Simon Kuenzer, Joao Martins, Mohamed Ahmed, Felipe Huici

• Title: Evolving the Internet with Connection Acrobatics

Where: ACM HotMiddlebox 2013 (CoNEXT workshop)

Authors: Catalin Nicutar, Christoph Paasch, Marcelo Bagnulo and Costin Raiciu

• Title: Trevi: Watering Down Storage Hotspots with Cool Fountain Codes

Where: ACM Hotnets 2013

Authors: George Parisis, Toby Moncaster, Anil Madhavapeddy, Jon Crowcroft

• Title: Rethinking Access Networks with High Performance Virtual Software BRASes

Where: European Workshop on SDN 2013

Authors: Roberto Bifulco, Thomas Dietz, Felipe Huici, Mohamed Ahmed, Joao Martins, Saverio Niccolini, Hans-Joerg Kolbe

3.2 Publications - Year 2 - 2014

• Title: Multipath TCP

Where: ACM Queue

Authors: Christoph Paasch, Olivier Bonaventure

• Title: HACK: Hierarchical ACKs for Efficient Wireless Medium Utilization

Where: Usenix ATC 2014 - Best paper award

Authors: Lynne Salameh, Astrit Zhushi, Mark Handley, Kyle Jamieson, Brad Karp

• Title: Experimental Evaluation of Multipath TCP Schedulers

Where: ACM Capacity Sharing 2014 (SIGCOMM workshop)

Authors: Christoph Paasch, Simone Ferlin, Ozgu Alay, Olivier Bonaventure

• Title: FUBAR: Flow Utility Based Routing

Where: **ACM Hotnets 2014** Authors: Nikola Gvozdiev, Brad Karp, Mark Handley

• Title: Using Cooperation for Low Power Low Latency Cellular Connectivity

Where: ACM CoNEXT 2014

Authors: Catalin Nicutar, Dragos Niculescu, Costin Raiciu

3.3 Publications - Year 3 - 2015 and beyond

• Title: Opportunistically Improving Datacenter Nerwork Utilization with GRIN

Where: Usenix NSDI 2015

Authors: Alexandru Agache, Razvan Deaconescu, Costin Raiciu

• Title: Towards Wifi Mobility without Fast Handover

Where: Usenix NSDI 2015

Authors: Andrei Croitoru, Dragos Niculescu, Costin Raiciu

• Title: InNet: Enabling in-network processing for the masses

Where: Eurosys 2015

Authors: Radu Stoenescu, Vladimir Olteanu, Matei Popovici, Mohamed Ahmed, Joao Martins, Roberto Bifulco, Filipe Manco, Felipe Huici, Georgios Smaragdakis, Mark Handley, Costin Raiciu

• Title: Shaping the Linux kernel MPTCP implementation towards upstream acceptance

Where: Netdev01, March 2015

Authors: Doru Gucea, Octavian Purdila

• Title: Analysis and Enhancements to Probabilistic Caching in Content-Centric Networking Where: The Computer Journal - Oxford Journals

Authors: J. Garcia-Reinoso, I. Vidal, D. Diez, D. Corujo and R.L. Aguiar

• Title: C3: Cutting Tail Latency in Cloud Data Stores via Adaptive Replica Selection Where: Usenix NSDI 2015

Authors: Lalith Suresh, Marco Canini, Stefan Schmid, Anja Feldmann

• Title: A first look at real Multipath TCP traffic

Where: Traffic Measurement and Analysis (TMA) 2015

Authors: B. Hesmans, H. Tran Viet, R. Sadre, O. Bonaventure

• Title: Transparent reallocation of control functions in IMS deployments

Where: IEEE Communications Magazine 2016 (to appear)

Authors: Jaime Garcia-Reinoso, Ivan Vidal, Paolo Bellavista, Ignacio Soto, Pedro Aranda

 Title: Informing Protocol Design Through Crowdsourcing: the Case of Pervasive Encryption Where: ACM SIGCOMM Workshop on Crowdsourcing and crowdsharing of Big (Internet) Data (C2B(I)D), Aug 2015, London, UK. Best paper Award

Where: Also published in ACM Computer Communication Review, Volume 45 Issue 5, October 2015

Authors: AM Mandalari, M Bagnulo, A Lutu

• Title: Observing Real Smartphone Applications over Multipath TCP

Where: IEEE Communications Magazine: Network Testing Series, March 2016 (to appear)

Authors: Quentin De Coninck, Matthieu Baerts, Benjamin Hesmans, Olivier Bonaventure

• Title: Observing Real Multipath TCP Traffic

Where: Computer Communications (accepted for publication)

Authors: Viet-Hoang Tran, Quentin De Coninck, Benjamin Hesmans, Ramin Sadre, Olivier Bonaventure

• Title: SMAPP : Towards Smart Multipath TCP-enabled APPlications

Where: Conext'15, December 2015

Authors: Benjamin Hesmans, Gregory Detal, Sbastien Barr, Raphal Bauduin, Olivier Bonaventure

• Title: A First Analysis of Multipath TCP on Smartphones

Where: PAM'2016, 2016

Authors: Quentin De Coninck, Matthieu Baerts, Benjamin Hesmans, Olivier Bonaventure

At the time of writing there are also 2 papers under review at IEEE transaction on Networking (A Game-Theoretic Method to Detect I-BGP Oscillations by Jose M. Camacho, Alberto Garca-Martnez, Francisco Valera, Yuedong Xu) and IEEE Communications Magazine (Transparent reallocation of control functions in IMS deployments by Jaime Garcia-Reinoso, Ivan Vidal, Paolo Bellavista, Ignacio Soto, Pedro Aranda).

3.4 Presentation at events, standardization activities, press releases

Trilogy2 disseminated the foreground of the project at large with several activities, out of which we report the most significant:

• 2013

Keynote and tutorials on Multipath TCP by Olivier Bonaventure and Christoph Paasch at IEEE LAN/MAN Workshop (Brussels, Belgium), IETF 87 (Berlin, Germany) and NorNet Users Workshop (Oslo, Norway)

Contributions on Multipath TCP by Marcelo Bagnulo, Cristoph Paasch, Olivier Bonaventure, Costin Raiciu, Toby Moncaster and Jon Crowcfroft at IETF 87 (Berlin, Germany) and IETF 88 (Vancouver, Canada)

Contribution to the creation of a new IETF Working Group on Network Service Chaining by Diego Lopez at IETF 87 (Berlin, Germany)

Press releases by Julian Chesterfield on Trilogy2 that received broad press coverage at PingZine, CDNadvisor, DataCenterSolutions, StorageNewsletter, TheEnterpriseCloudSite

Several presentations on NFV Infrastructure Architecture, Management and Orchestration, Performance and Portability, Security by Andy Reid, Diego Lopez, Javier Ramn and Bob Briscoe at ETSI ISG NFV meetings (Valbonne/France, Santa Clara/USA, Bonn/Germany, Sunnyvale/USA)

IETF RFC 7141 on Byte and Packet Congestion Notification by Bob Briscoe

Contributions to the draft group specifications of the ETSI ISG NFV Infrastructure and Security by Andy Reid and Bob Briscoe

Contributions on Multipath TCP by Cristoph Paasch and Olivier Bonaventure at IETF 89 (London, UK)

Contributions on TCP extensions by Bob Briscoe and Toby Moncaster at IETF 90 (Toronto, Canda) and IETF 91 (Honolulu, USA)

Talk on Trilogy2 contributions to the XEN virtualization framework by Filipe Manco at the XEN Developer Summit (Chicago, 2014)

Virtualizing the Network Edge: Virtual CPE for the datacenter and the PoP, G. Bernini, G. Carrozzo, P.A. Gutierrez, D. Lopez, poster at EUCNC2014 conference, Bologna, Italy, June 23-26 2014.

• 2015

Presentation on Internet architecture by Bob Briscoe at the joint ETSI ISG NFV meeting and IAB workshop

NFVI PoP network topology: Problem statement, M. Bagnulo, Internet draft, October 2015.

^{• 2014}

NFVI PoP network topology: Problem statement, M. Bagnulo, presentation at NFVRG at the IETF in Yokohama, July 2015

Informing Protocol Design Through Crowdsourcing: the Case of Pervasive Encryption, presentation ar HOPSRG at the IETF in Yokohama, July 2015

Informing Protocol Design Through Crowdsourcing: the Case of Pervasive Encryption, presentation at the joint ACM/IRTF RAIM workshop in Yokohama, July 2015.

VNF Orchestration For Automated Resiliency in Service Chains, G. Bernini, V. Maffione, P.A. Gutierrez, D. Lopez, Internet Draft, 3 July 2015.

VNF Orchestration For Automated Resiliency in Service Chains, P.A. Gutierrez, presentation at NFVRG at the IETF93, 22 July 2015.

VNF Pool Orchestration For Automated Resiliency in Service Chains, G. Bernini, V. Maffione, P.A. Gutierrez, D. Lopez, updated Internet Draft, 11 October 2015.

VNF Pool Orchestration For Automated Resiliency in Service Chains, P.A. Gutierrez, presentation at NFVRG at the IETF94, 4 November 2015.

Multipath TCP Address Advertisement, O. Bonaventure, Internet draft presented at MPTCP WG, IETF93, July 2015

Improving Multipath TCP Backup Subflows, O. Bonaventure, F. Duchene, Q. De Coninck, M. Baerts, B. Hesmans, Internet draft presented at MPTCP WG, IETF93, July 2015

Experimental Multipath TCP option, O. Bonaventure, B. Hesmans, Internet draft presented at MPTCP WG, IETF93, July 2015, will be included in RFC6824bis

Supporting long TCP options in Multipath TCP, O. Bonaventure, Internet draft presented at MPTCP WG, IETF93, July 2015

Multipath TCP timestamp option, O. Bonaventure, Internet draft presented at MPTCP WG, IETF93, July 2015

Use Cases and Operational Experience with Multipath TCP, O. Bonaventure, C. Paasch, G. Detal, IETF WG Document, presented at IETF93, July 2015

Analysis of the TCP EDO option, O. Bonaventure, O. Tilmans, Internet draft, presented at the TCP WG, IETF93, July 2015

Multipath TCP MIB, F. Duchene, C. Paasch, O. Bonaventure, Internet draft, presented at the MPTCP WG, IETF94, November 2015

The evolution of Internet transport protocols, O. Bonaventure, Keynote at CNSM'15 conference, November 2015

3.5 Software

Trilogy2 has been heavily committed to running code in order to demonstrate the feasibility of the approaches proposed. Several of these pieces of software has been also released and maintained as open source.

• Title: Multipath TCP - Linux kernel patch

Description: Kernel patch to implement Multipath TCP.

Repository: http://www.multipath-tcp.org/

• Title: MPTCPTRACE

Description: Software to analyze MPTCP packet traces (started within another project, Trilogy2 contribution includes various improvements to scalability).

Repository: https://bitbucket.org/bhesmans/mptcptrace

• Title: MBDETECT

Description: Detection tool to validate the operation of MPTCP through specific middleboxes (entirely developed inside Trilogy2).

Repository: https://bitbucket.org/rbauduin/mbdetect

• Title: MP-BGP Support in Quagga

Description: Patch to make sure that MPLS-VPN NRLIs include the route target and the route distinguisher (this patch was done for the vCPE use case in the context of Trilogy2, Github includes the full Quagga code to build the control plane).

Repository: https://github.com/paaguti/quagga-v4vpn

• Title: MP-BGP Support in Scapy

Description: Patch to improve BGP-4 support in general and MPLS-VPN NRLIs in particular in Scapy 2.2 (this patch was done for the vCPE use case in the context of Trilogy2, Github includes the complete Scapy 2.2 code, a port to scapy 2.3 is in progress).

Repository: https://github.com/paaguti/scapy-com

• Title: Minicache

Description: Development of Minicache system (entirely developed inside Trilogy2).

Repository: currently closed source due to internal exploitation.

• Title: VNF pool enabled virtual CPE

Description: A set of integrated software pieces for the virtualization of CPE functions in the operator's datacentre, to let business customers incorporate new VMs in their MPLS-VPNs (entirely developed inside Trilogy2).

Repository: currently closed source due to internal exploitation.

• Title: Software BRAS

Description: Virtualized software BRAS using ClickOS (entirely developed inside Trilogy2).

Repository: currently closed source due to internal exploitation.

• Title: Virtual Machine Migration for Mobile

Description: An integrated solution for virtualization on mobile devices, allowing applications to be encapsulated in a VM and live migrated between devices. The software modules included in it are: a custom-designed linux-based VM with MPTCP support (VMSlot), a VMSlot manager, Android application for VMslot support, Linux application for VMSlot support, Qemu patches for Android integration (entirely developed inside Trilogy2).

Repository: currently in the internal process of releasing the code as open source. coming soon.

• Title: PVTCP/Shushi implementation

Description: Implementation of the PVTCP API for encrypted, reconfigurable datacenter I/O (entirely developed inside Trilogy2).

Repository: coming soon

• Title: Jitsu

Description: A DNS server that automatically starts unikernels on demand (entirely developed inside Trilogy2).

Repository: https://github.com/MagnusS/jitsu

• Title: Mirage-bench

Description: A framework for benchmarking client/server Mirage applications (entirely developed inside Trilogy2).

Repository: https://github.com/MagnusS/mirage-bench

• Title: mirage-bootvar

Description: Library for reading Mirage unikernel boot parameters on Xen (entirely developed inside Trilogy2).

Repository: https://github.com/MagnusS/mirage-bootvar

• Title: Mobile Kibbutz

Description: Software patch for Cyanogen-Mod, a kernel patch to MPTCP and user-level python scripts that together implement the mobile kibbutz, as presented in the CoNEXT 2014 paper (joint funding from Trilogy2 and Mobil4, a project funded by the Romanian government).

Repository:

• Title: Channel switching

Description: Software patch for the 802.11 softmac, applicable to Linux kernels, that enables dynamic switching between different channels with preconfigured slots to be spent on each channel. Python scripts that associate to APs on those channels (joint funding from Trilogy2 and Mobil4, a project funded by the Romanian government).

Repository: coming soon

• Title: Multi-wifi

Description: Python software that associated to all known APs on a preconfigured Wifi channel. Estimation of the downlink physical loss rates from the AP using the metric described in the NSDI 2015 paper ("Towards Wifi Mobility without Fast Handover"). Implementation of ECN marking to discourage the sender from using bad APs (joint funding from Trilogy2 and Mobil4, a project funded by the Romanian government).

Repository: coming soon

• Title: GRIN

Description: Patch to MPTCP kernel together with user-level scripts that assign addresses automatically in a local testbed, implementing GRIN (NSDI 2015) (entirely developed inside Trilogy2).

Repository: available on request, no plans for open source.

• Title: Federated Market

Description: VM resource market (compute) across the wide-area. Uses Ruby on Rails, fog.io, Python (research activities supporting the commercial product being developed in Trilogy2).

Repository: http://onapp.com/federation/ (commercial product)

• Title: Cloud.net

Description: UI-front end to the Federated Market. Uses Ruby on Rails (research activities supporting the commercial product being developed in Trilogy2).

Repository: https://jager.cloud.net/

• Title: Cloud Liquidity - DRaaS

Description: Use-case for block storage (and later VM migration) across the wide-area. Storage level uses C code. Web UI front end uses Ruby on Rails. Control path code uses Python (research activities supporting the commercial product being developed in Trilogy2).

Repository: commercial offering

• Title: HACK

Description: NS3 and SoRa implementations of HACK (entirely developed inside Trilogy2). Repository: www0.cs.ucl.ac.uk/staff/a.zhushi/hack/

• Title: FlowParser

Description: A library to do online/offline flow capture and header fields compression (entirely developed inside Trilogy2).

Repository: https://github.com/ngvozdiev/flowparser

• Title: Pathfinder

Description: An OpenFlow 1.3 controller with sFlow support that implements FUBAR (entirely developed inside Trilogy2).

Repository: https://github.com/ngvozdiev/pathfinder

• Title: FUBAR

Description: An implementation of the algorithm presented in the HotNets 2014 paper (entirely developed inside Trilogy2).

Repository: available on request, no plans for open source.

• Title: tracebox

Description: Software to analyze middlebox interference through a network (started within the CHANGE project, Trilogy2 contribution includes support for TFO and EDO options and various improvements)

Repository: https://www.tracebox.org

3.6 Summary and Highlights of Entire Project

Trilogy2 strategy was to focus of publications with high impact as opposed to disseminate massively many tiny output pieces. We achieved several publications at top tier events that will leave a huge impact on the community: 3 publications at Usenix NSDI, 2 at CoNEXT, 2 at Hotnets and 1 respectively at Eurosys and Usenix ATC with its best paper award. Such achievement is a really outstanding and rare achievement: 3 publication at top tier events per year.

All publications has been supported by a huge effort in developing running code, most of which has been published and maintained as open source (where exploitation plans of the individual companies did not conflict with this).

Additionally Trilogy2 has achieved large dissemination in standards (IETF/IRTF and ETSI ISG NFV) as demonstrated by the continuous output of its partners at standardization meetings including the publication of an IETF RFC.

4 Conclusions

Trilogy2 has achieved noticeable impact in terms of exploitation of the foreground results:

- The open source implementation of multipath TCP in the Linux kernel is widely used commercially by several companies worldwide as detailed in section 1.1 and one of the company (Tessares) is a spin-off of one of the Trilogy2 partners (UCL-BE);
- Federated cloud storage has been used in commercial production by OnApp;
- Minicache is widely used internally in NEC for developing middleboxes that require a lightweight TCP/IP stack;
- The MPLS VPN solutions are part of the commercial offerings of Nextworks;
- The NFV solutions developed in the project are used in several pre-commercial projects at Telefonica.

Trilogy2 has supported, through continuous output, the standardization activities at the ETSI ISG NFV thanks to leadership provided by BT and Telefonica, as well as contributing to IETF activities thanks to the leadership of UC3M and UCL-BE.

Last but not least, the dissemination activities of members of the consortium have been impressive in terms of output at top tier conferences: three publications at Usenix NSDI, two at CoNEXT, two at Hotnets and one respectively at Eurosys and Usenix ATC, the latter with a best paper award. We believe that this average of three publications at top tier events per year is a rare achievement.

Bibliography

[1] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. TCP Extensions for Multipath Operation with Multiple Addresses\$. Technical Report 6824, IETF Secretariat, January 2013.