ERASER: Experimenting with real application-specific QoS guarantees in a large-scale RINA demonstrator

Dr. Jordi Perelló

Universitat Politècnica de Catalunya (UPC) – BarcelonaTech

FEC4 Meeting

Brugge (Belgium), October 8-10, 2018

Outline

- Experiment description
- Project results
- Business impact
- Feedback to Fed4FIRE+

EXPERIMENT DESCRIPTION

RINA & QTA-Mux

- **Clean-slate network architecture for the Future Internet**, based on a single type of layer called DIF (i.e., Distributed IPC Facility)
- Multiple **DIF instances** are always **stacked one on top of another**, each one providing IPC services to the DIFs or Apps above
- DIFs implement the same two protocols, **customizable via programmable policies**
- The **QTA-Mux** policy is the responsible for **the RINA QoS support**
- However, its performance has only been evaluated on single DIF scenarios so far

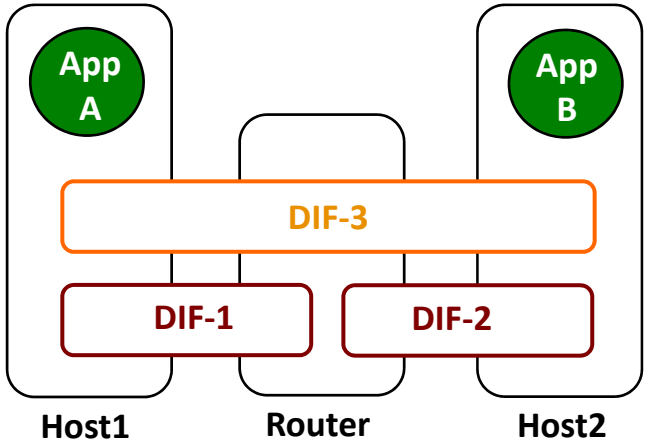

ERASER objectives

1. To define **QTA-Mux deployment scenarios** and **QoS Cubes** to be enforced in a **large-scale RINA network** scenario with multiple DIFs stacked one on top of another
2. To emulate a **realistic large-scale RINA network** scenario injecting **synthetic traffic flows** reproducing **heterogeneous applications**
3. To **evaluate the RINA QoS support** by measuring the **QoS metrics perceived** by synthetic traffic flows end-to-end
4. To perform a **real HD video (1080p) streaming demo** to better show the RINA QoS support under high congestion

ERASER scenario over the F4F+ Virtual Wall

37-Node Metro-Regional RINA network:

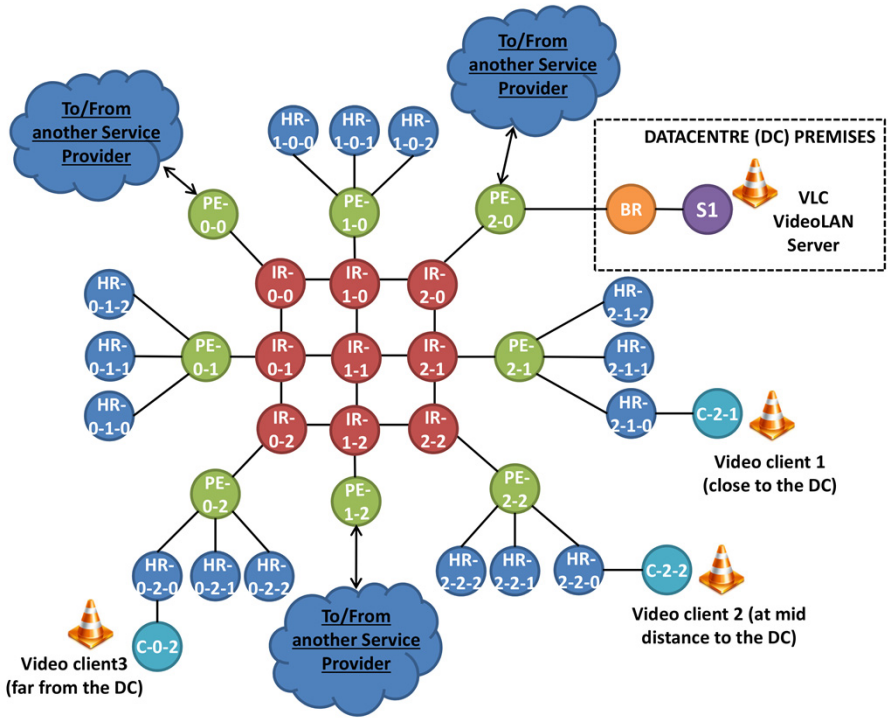

Configured DIFs

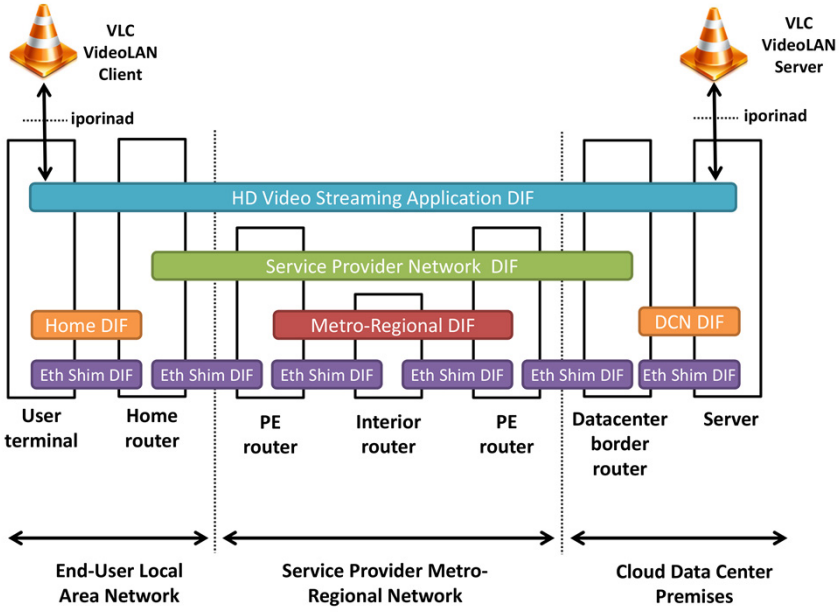

QTA-Mux deployment scenarios

- **At both SPN and MR DIFs / Only at the SPN DIF**
- Offered **QoS Cubes** & **synthetic application flow** characteristics:

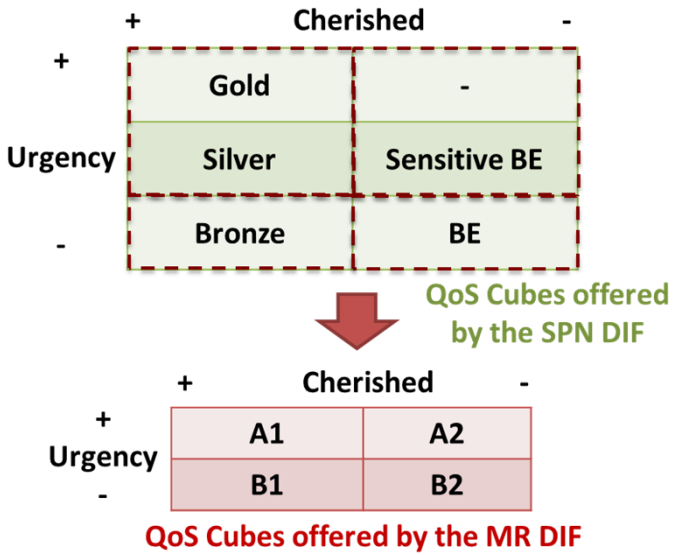

Application flow	Traffic distribution	Traffic direction	Details	Assigned to QoS Cube
HD Video Call	CBR	Bidirectional	CBR bitrate: 1.5 Mbps	Gold
Online Gaming	ON-OFF	Bidirectional	ON-OFF period avg. duration: 4s – 2s CBR bitrate during ON period: 4 Mbps	Silver
VoIP	ON-OFF	Bidirectional	ON-OFF period avg. duration: 3s – 3s CBR bitrate during ON period: 64 kbps	Sensitive BE
File Sharing	ON-OFF	Bidirectional	ON-OFF period avg. duration: 2s – 1s CBR bitrate during ON period: 5 Mbps	Bronze
Interactive traffic	Poisson	Bidirectional	Avg. bitrate: 2 Mbps	BE

PROJECT RESULTS

Open Research Data in:

http://doi.org/10.5281/zenodo.1420391

QTA-Mux at both SPN & MR DIFs

- 75 bidirectional synthetic traffic flows injected into the SPN DIF; NIC capacities limited to reproduce 70, 80 and 90% offered load scenarios
- **Adequate QoS differentiation even under high congestion**:

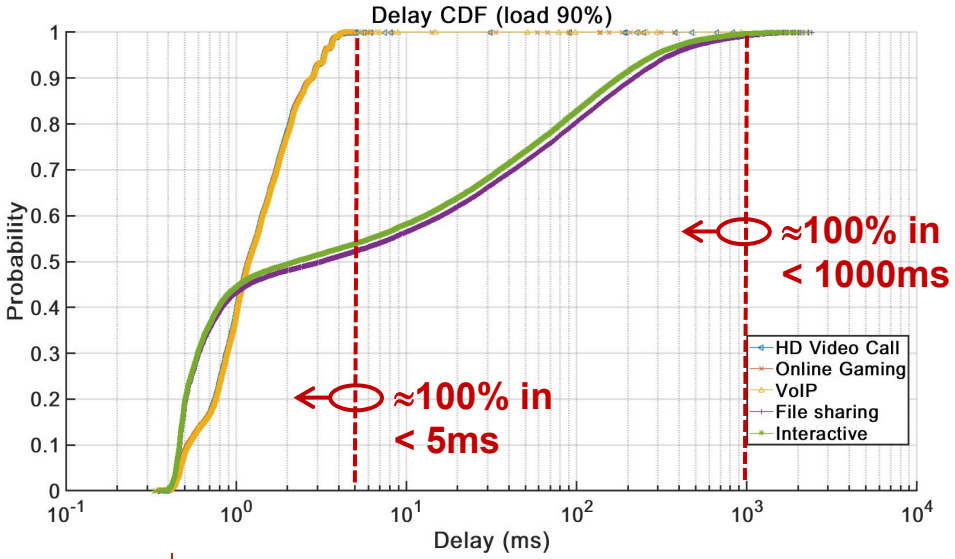

Perceived QoS metrics (90% load)

Traffic type	Min. latency (ms)	Max. latency (ms)	Avg. latency (ms)	Avg. packet loss (%)
HD Video Call	0.345145	830.5	1.483987	0.001659
Online Gaming	0.3626	375.76	1.468362	0.002339
VoIP	0.345375	924.95	1.469585	0.001342
File Sharing	0.354075	2408.85	72.310218	0.231222
Interactive	0.32893	2276	58.650972	3.469220

QTA-Mux at the SPN DIF only

- 75 bidirectional synthetic traffic flows injected into the SPN DIF; NIC capacities limited to reproduce 70, 80 and 90% load scenarios
- **QoS differentiation lost at the MR DIF** (default FIFO scheduling):

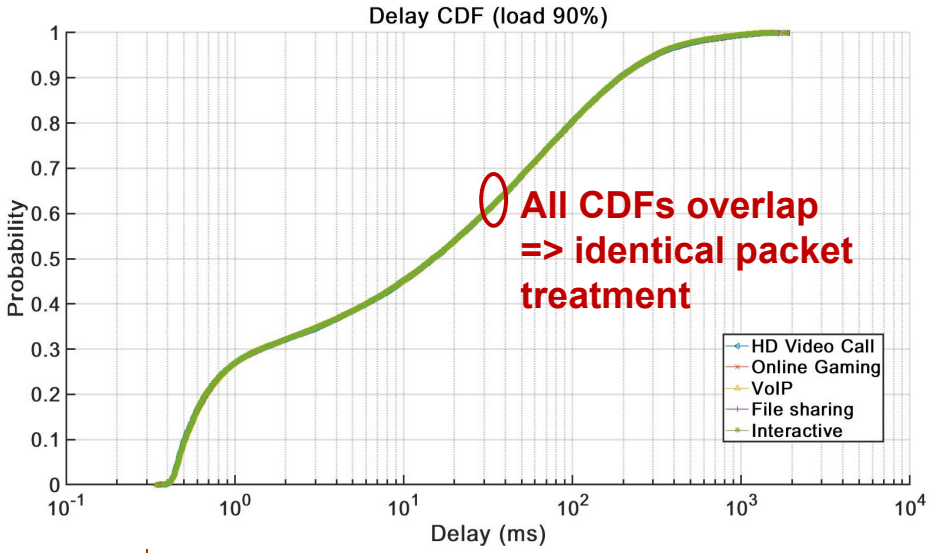

Perceived QoS metrics (90% load)

Traffic type	Min. latency (ms)	Max. latency (ms)	Avg. latency (ms)	Avg. packet loss (%)
HD Video Call	0.342600	1889.6	69.764187	0.139793
Online Gaming	0.343865	1893.9	70.126824	0.173666
VoIP	0.345180	1846.05	69.738218	0.150003
File Sharing	0.351630	1876.3	70.223907	0.155559
Interactive	0.346925	1865.15	69.441732	0.185567

Final HD video streaming demonstration

- 75 bidirectional synthetic traffic flows injected into the SPN DIF; NIC capacities limited to reproduce 80% and 90% load scenarios

Perceived QoE (80% load) from 1 (worst) to 5 (best)

QoS Cube assigned at HD Video Streaming DIF	S1 -> C-2-1 (1 hop at MR DIF)	S1 -> C-2-2 (2 hops at MR DIF)	S1 -> C-0-2 (4 hops at MR DIF)
Gold	5 (perfect quality)	5 (perfect quality)	5 (perfect quality)
Best Effort	2 (severe stuttering & frames lost)	1 (completely frozen)	2 (severe stuttering & frames lost)

Perceived QoE (90% load) from 1 (worst) to 5 (best)

QoS Cube assigned at HD Video Streaming DIF	S1 -> C-2-1 (1 hop at MR DIF)	S1 -> C-2-2 (2 hops at MR DIF)	S1 -> C-0-2 (4 hops at MR DIF)
Gold	5 (perfect quality)	3 (moderate stuttering & frames lost)	5 (perfect quality)
Best Effort	2 (severe stuttering & frames lost)	1 (completely frozen)	1 (completely frozen)

BUSINESS IMPACT

Value perceived

- Better **understanding of the RINA QoS support** through experimentation in a large-scale network scenario
- Better **positioning within the RINA research community** as a team expert in RINA, particularly in its QoS mechanisms and policies
- Gained **expertise on** available open-source **RINA implementations and tools** for future experiments and test-bed setups
- Gained **expertise on using Fed4FIRE+** for large-scale experiments in the future

Why Fed4FIRE+?

- Our group at the UPC did not own the **resources** needed for the ERASER large-scale experimentation activities
- **Virtual Wall** previously **used for RINA-related experimentation**
- Fed4FIRE+ Open Call 3 **offered funding** for conducting experiments
- Accepted experiments in Fed4FIRE+ Open Calls receive **support from the Patron(s)**, important for first-time users

Follow-up activities: ALLIANCE test-bed

- We plan to deploy a **small-scale RINA network test-bed** in the framework of the Spanish project ALLIANCE
- The **experience** with open-source **RINA implementations and tools gained in ERASER** will be **key to this end**
- Interesting **proposals yielding good results at small-scale** might be good candidates for **large-scale evaluation over Fed4FIRE+**

Follow-up activities: Other

- **Dissemination** of the **obtained results in ERASER**
  - **Target**: 6th International Workshop on the Recursive InterNetwork Architecture (RINA 2019), co-located with the 22nd Conference on Innovation in Clouds, Internet and Networks (ICIN 2019), February 2019, Paris (France)
- Participation in **new RINA-related proposals** to be submitted to upcoming H2020 calls around 5G and the Future Internet
- Inclusion of **ERASER experiment outcomes in Master and undergraduate lectures** at the UPC

FEEDBACK TO FED4FIRE+

Used resources

- **ERASER small-scale RINA scenario** for initial software integration and tests: 10 physical machines in Virtual Wall 2
- **ERASER large-scale RINA scenario** for the complete experiments & demonstration: 37 physical machines in Virtual Wall 1

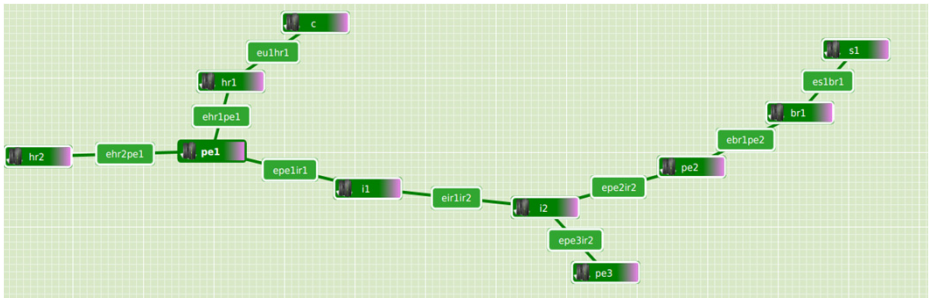

Small-scale ERASER scenario (JFed view)

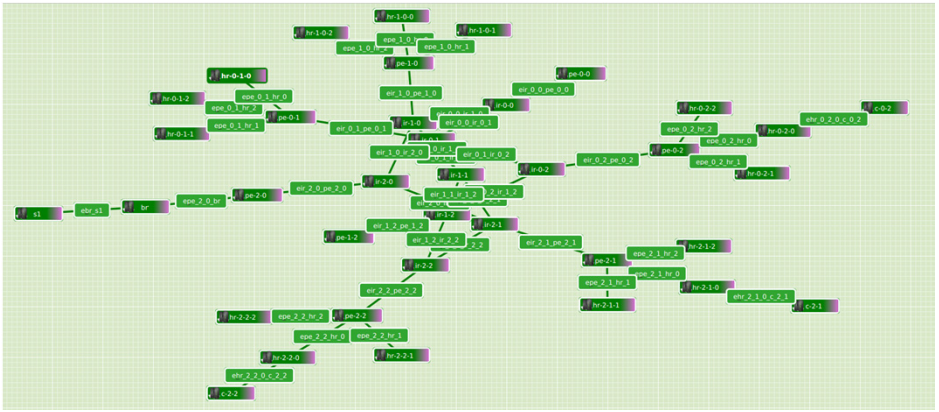

Large-scale ERASER scenario (JFed view)

Used tools (1/2)

- **JFed** has been initially used to configure the small-scale experiments, check physical machine availability, reboot nodes, etc.
- RINA-specific configurations (creation of DIFs, enrolment of nodes in DIFs, allocation of flows, etc.) are tedious to be performed manually
- Hence, we started using **Rumba** from the very beginning
  - Python framework to easily define large RINA scenarios and run scripted experiments, e.g., over the Fed4FIRE+ Virtual Wall
  - Developed in the H2020 ARCFIRE project
  - Presented and demonstrated at previous FEC2 and FEC3 meetings

Used tools (2/2)

- ERASER has also used more **open-source implementations and tools made available by the RINA research community**:
  - IRATI RINA Stack for OS/Linux
  - QTA-Mux scheduling policy
  - RINA-tgen (extended to generate CBR, Poisson & ON-OFF traffic)
  - rina-echo-time (ping-like application used to measure e2e latencies)
  - iporinad (used to tunnel IP video streaming traffic over RINA)
- **Win-win situation!** ERASER has also served to test them, reporting encountered bugs and required extensions to their original developers

Added value of Fed4FIRE+

- Large amount of **physical resources** available for experimentation
- Availability of a very significant number of **federated test-beds** with **diverse networking technologies**
- **Unified and easy setup of experiments** spanning several test-beds
- **Offered tools** to support and facilitate experimentation activities
- **Available documentation**


Co-funded by the European Union


Co-funded by the Swiss Confederation

This project has received funding from the European Union's Horizon 2020 research and innovation programme, which is co-funded by the European Commission and the Swiss State Secretariat for Education, Research and Innovation, under grant agreement No 732638.

WWW.FED4FIRE.EU