

IRATI - Investigating RINA as an Alternative to TCP/IP



*FIRE Engineering Workshop, Ghent, Belgium
November 6th, 2012*



Distributed
Applications
Networks
Area

Sergi Figuerola
Project coordinator
Director @ DANA
(sergi.figuerola@i2cat.net)
<http://dana.i2cat.net/>

<http://irati.eu>



Current situation

- “Middleboxes” e.g. NATs, firewalls, are mainly patches to solve other problems.
- Naming, addressing and routing: multi-homing, multicast and mobility are hard to support.
- Scalability: non-existent, e.g. unbounded router table size.
- Security: no systematic approach to security, secure each protocol or add boxes (firewalls).
- Quality of Service: it is a recurrent issue.
- Management: complex due to the large number of protocols.
- The Internet seems to becoming a patchwork of growing complexity.
- But the problems are not only in the new demands of today’s applications, but in the flaws of the Internet’s original design.

IRATI project at a glance

- What is the main goal?
 - To advance the state of the art of the RINA (Recursive Internetwork Architecture) reference model and specifications to enable implementations deployable in production scenarios.
 - Design and implement a prototype on top of Ethernet will enable experimentation and evaluation of RINA compared to TCP/IP and its migration path.

Who? 4 partners



5 activities:

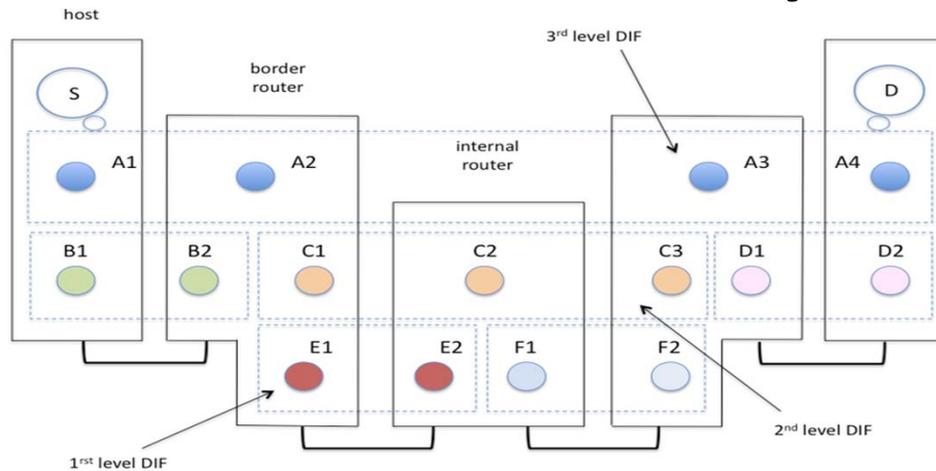
- WP1: Project management 
- WP2: Architecture, Use cases and Requirements 
- WP3: Software Design and Implementation 
- WP4: Deployment into OFELIA testbed, Experimentation and Validation 
- WP5: Dissemination, Standardisation and Exploitation 

Budget	
Total Cost	1.168.271 €
EC Contribution	897.060 € (76,8%)
Duration	2 years
Start Date	1 st January 2013
External Advisory Board	
Juniper Networks, ATOS, CISCO Systems, Telecom Italia	

Project Objectives

- Enhance the RINA architecture reference model and specifications, focusing on DIFs (Distributed IPC Facilities) over Ethernet
 - DIF: A collection of two or more Application Processes cooperating to provide Inter-Process Communication (IPC).
- RINA open source prototype over Ethernet for a UNIX-like OS.
- Validate experimentally the RINA theory and produce a set of requirements and comparison tests against TCP/IP.
- RINA prototype over Ethernet for JunOS SDK (Juniper MX480).
- Interoperability with other RINA prototypes over UDP/IP, which are under development by Pouzin Society, BU, I2CAT and TSSG.
- Provide feedback to OFELIA on experimentation with a non-IP based technology.

Why RINA?



Each DIF provides IPC services over a limited scope. First level DIFs operate on top of a physical medium, and their policies are optimized to deal with the particularities of the physical medium. First level DIFs provide IPC services to second level DIFs, etc.

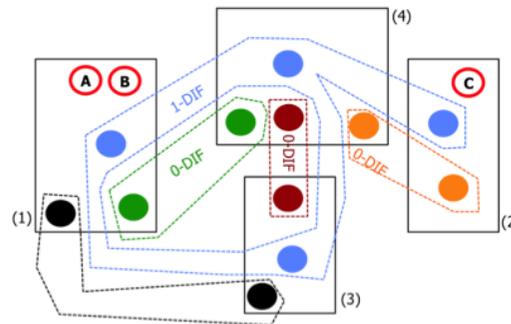
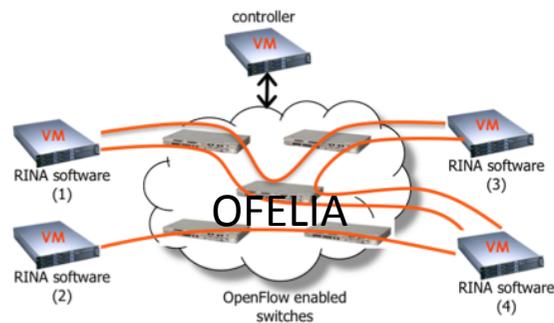
- An architecture based on a set of fundamentals: Networking is Inter-Process Communication (IPC).
- A complete naming and addressing schema.
- Simplicity: A single error and flow control protocol, a single application protocol.
- Inherently supports: Multi-homing, mobility and multicast, security (the layer is a securable container), Quality of Service support.
- Scalability: Recursion provides a divide and conquer approach, existence of policies that constrain the membership size of each DIF, use of hierarchical addresses.
- A new approach to current network architectures.

and... why OFELIA?

- It offers a programmable Ethernet network
 - Experiment with different network topologies beyond typical isolated integration tests (controlled L2 network through OpenFlow).
 - Create and modify the network topology dynamically to test different dynamic network conditions (e.g. link failures).
- It also offers geographical diversity and multi-domain environment
 - Deploy complex scenarios to simulate the behaviour of the RINA stack with current network technologies (migration path), or experiment with multiple RINA internetworks.
 - Test under real network conditions including delays introduced by long distance links.
- IRATI will provide an interesting use case for the OFELIA facility, the experimentation on non-IP based technology, and it will contribute to the development of the OpenFlow controller software used to dynamically setup and modify the Ethernet topologies used in the IRATI experiments.

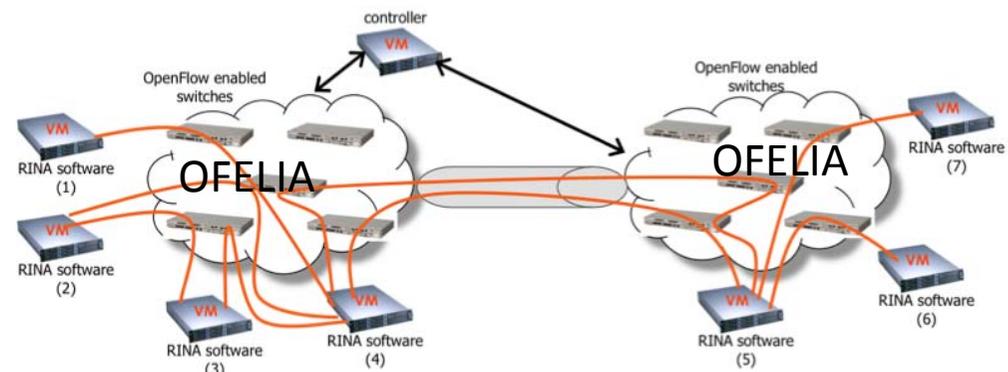
Phase 1: Basic Functionality - UNIX-like OS

- Validate the basic RINA functionality.
- Define the requirements of a RINA deployment within a local area network (weak security requirements, support of legacy applications, best-effort QoS, flat addressing scheme).
- The target platform will be the UNIX-like OS.



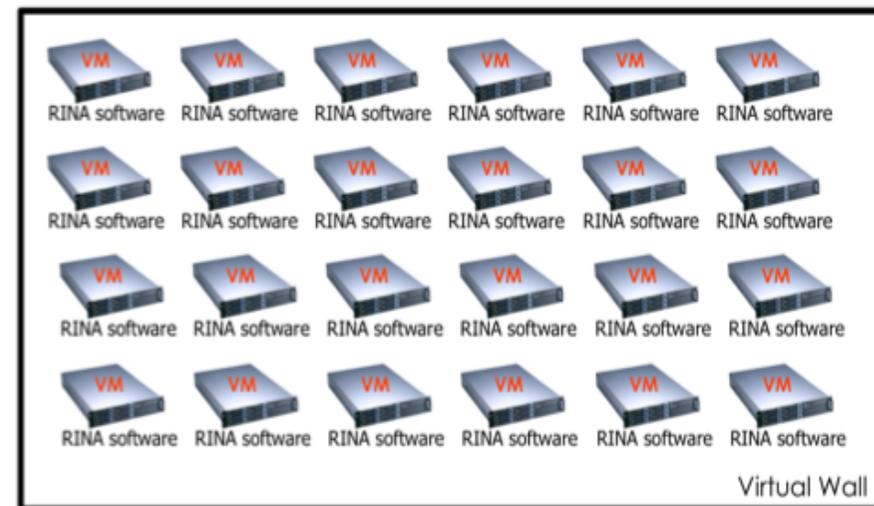
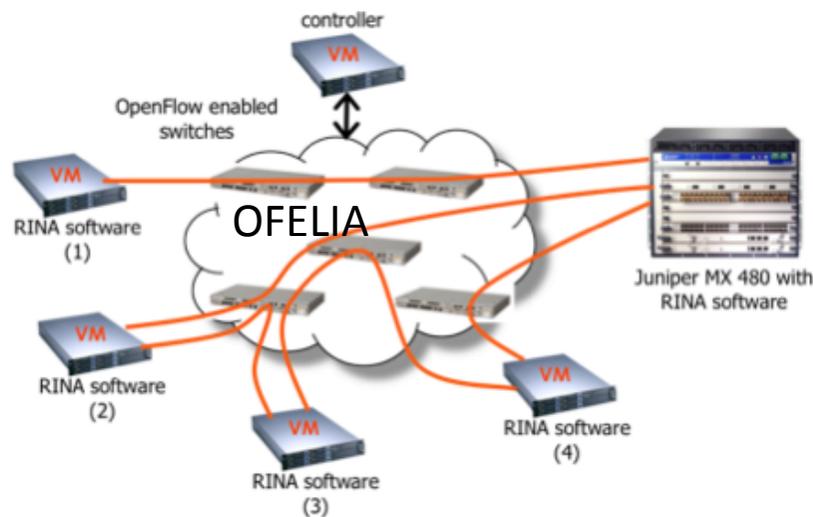
Single-island deployment with corresponding RINA DIFs

Multi-island experiment with several RINA internetworks



Phase 2: Scalability and JunOS

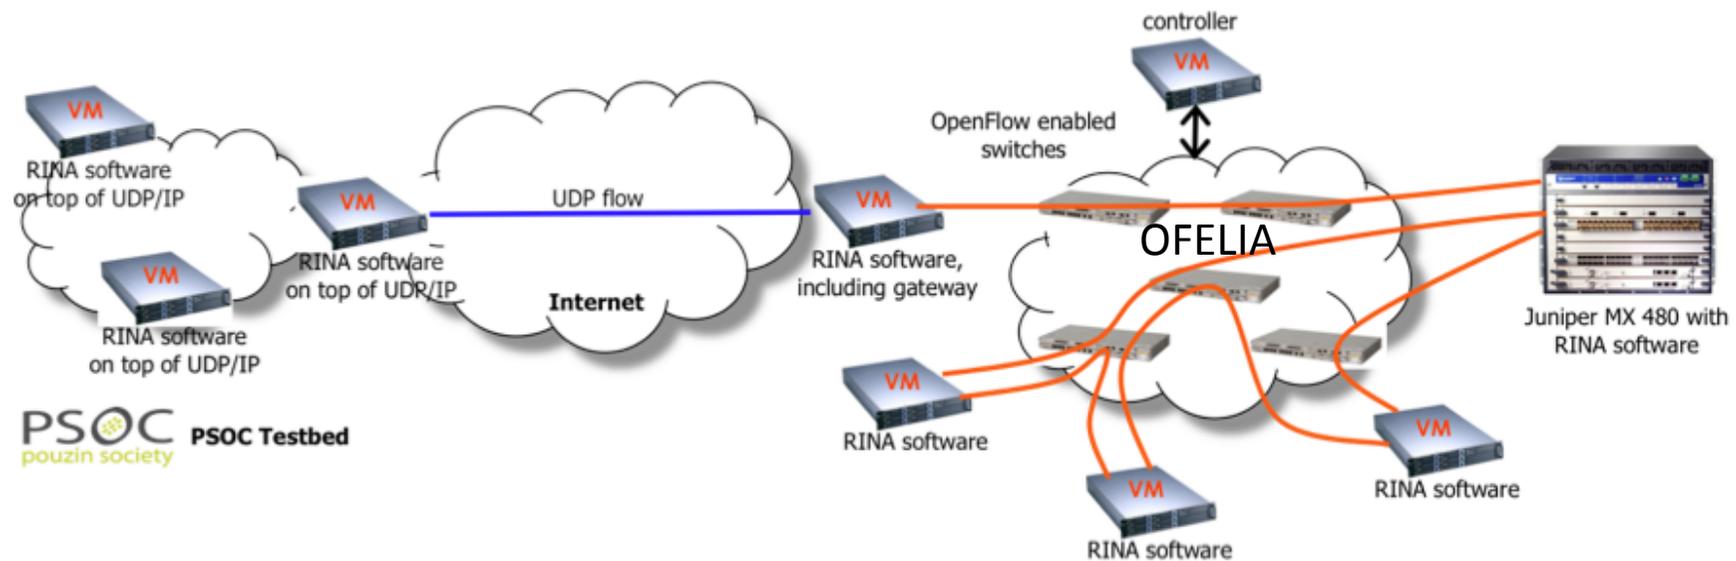
- Target different deployment scenarios
 - single network provider with different network hierarchies, different levels of QoS, multiple network service providers, etc...
- Assume that all the networks are either RINA or Ethernet capable (i.e. no IP).
- The UNIX- like OS and JunOS SDK will be the target platforms of this phase.



Single island with Juniper router and multiple RINA nodes within the Virtual Wall

Phase 3: IP gateway and interoperability

- Interoperability between RINA prototypes, developed outside of the project and deployed in a RINA network surrounded by an IP network.
- At this stage we will collaborate with the Pouzin Society through Boston University.



Interoperability between the PSOC and IRATI RINA prototypes

IRATI's contribution

- Enable a European leadership within the RINA exploration and research, first-hand experience on implementation and deployment of this new architecture while improving European competitiveness in Future Internet research.
- Involve several European institutions and disseminate the features, advantages and market opportunities that RINA enables to the industrial and academic European society.
- Bring up new business models and services as the result of the RINA adoption, involving all the current Internet stakeholders.

Technological Impact

- RINA is based on a solid fundamental theory of networking.
- Support heterogeneous physical media and heterogeneous applications through policy configuration.
- RINA's complete naming and addressing schema supports mobility and multihoming inherently.
- The separation of mechanism and policy in a repeating structure focuses advanced simpler programmability beyond the SDN initiative.
- Security: A layer is a securable container, inherently more secure before cryptographic methods are added.
- The DIF is the general theory of network virtualization.
- RINA scales indefinitely, limitations may be in physics, but not in the architecture (to be validated in IRATI experimentation).
- Greatly reduces capex and opex
 - Capex: Simplicity of RINA will reduce equipment costs. No need for firewalls, NATs and others.
 - Opex: The less middleboxes required, the less are the configurations to make, the fewer the errors to face.

Adoption process for RINA

- RINA provides many options for a phased adoption since it is not fixed to one place in the protocol stack.
- It can be applied below TCP/IP, to provide network service to TCP/IP hosts, or to connect IP networks, but also it can be used above IP, using existing IP links as middleware or a VPN.
- RINA networks can be alongside TCP/IP networks as an alternative or translating between applications
- Faux sockets API: A library providing the standard sockets interface but implementing the communication with the RINA stack
- Existing applications can use RINA DIFs without being modified, allowing seamless migration to RINA networks
- However a native RINA API is required to enjoy the full benefits (TBD)
- IRATI will design and implement the two libraries for the UNIX-like OS applications to interface with RINA

Dissemination Activities

- Dissemination is an activity of key importance, targeting at academia, industry, media and general public
 - Online presence
 - Online open-source software community
 - Scientific papers on high-impact magazines and conferences
 - A final workshop with dissemination of the results, live demonstrations and symposiums
 - Discussion and feedback through the External Advisory Board (ATOS, Juniper Networks, CISCO Systems and Telecom Italia)
 - Internal partner dissemination

Exploitation Activities

- Commercial exploitation
 - Exploitation strategies differ for the industrial (Interoute, Nextworks) and academic partners (iMinds, i2CAT), because of different targeting markets.
 - i2CAT considers the creation of spin-off.
- Research exploitation
 - Increase the know-how and IPR on the project related technologies, and eventually through patents.
 - Identification of RINA unexplored research topics and other technological challenges.
- Academic exploitation
 - At this step involve the Flemish and Catalan relevant universities into RINA research.
 - Work with universities to develop master and PhD level courses that teach the theory behind RINA.
 - i2CAT and iMinds expect to produce at least 4 PhD thesis.



Distributed
Applications
Networks
Area



Thank you
Moltes gràcies

contact: Sergi Figuerola
Project Coordinator

(sergi.figuerola@i2cat.net)

More information about

- IRATI: <http://irati.eu>
- RINA: <http://rina.tssg.org>, <http://pouzinsociety.org>

