# *Early RINA prototyping and deployment in the IRATI project, and future research in the PRISTINE and IRINA projects*

i2cat

**Sergi Figuerola,**
Innovation and Technology Director
(sergi.Figuerola@i2cat.net)

**13th annual Global LambdaGrid Workshop October 4th, 2013**

glif
Global Lambda Integrated Facility

IRATI
INVESTIGATING RINA

# Agenda

- **What is RINA**
- Why researching RINA
- Flow of research and development activities
- EC-funded RINA research
  - IRATI
  - PRISTINE (in negotiations)
  - IRINA (in negotiations)

# RINA is an..

**Innovative approach** to computer networking using inter-process communications (IPC), a set of **techniques for the exchange of data among multiple threads** in processes running on one or more **computers connected to a network.**

Ref. : J. Day: “Patterns in Network Architecture: A Return to Fundamentals, Prentice Hall, 2008.

**The RINA principle:**

Networking is not a layered set of different functions but rather a single layer (DIF) of distributed IPC’s that repeats over different scopes.

# RINA Architecture

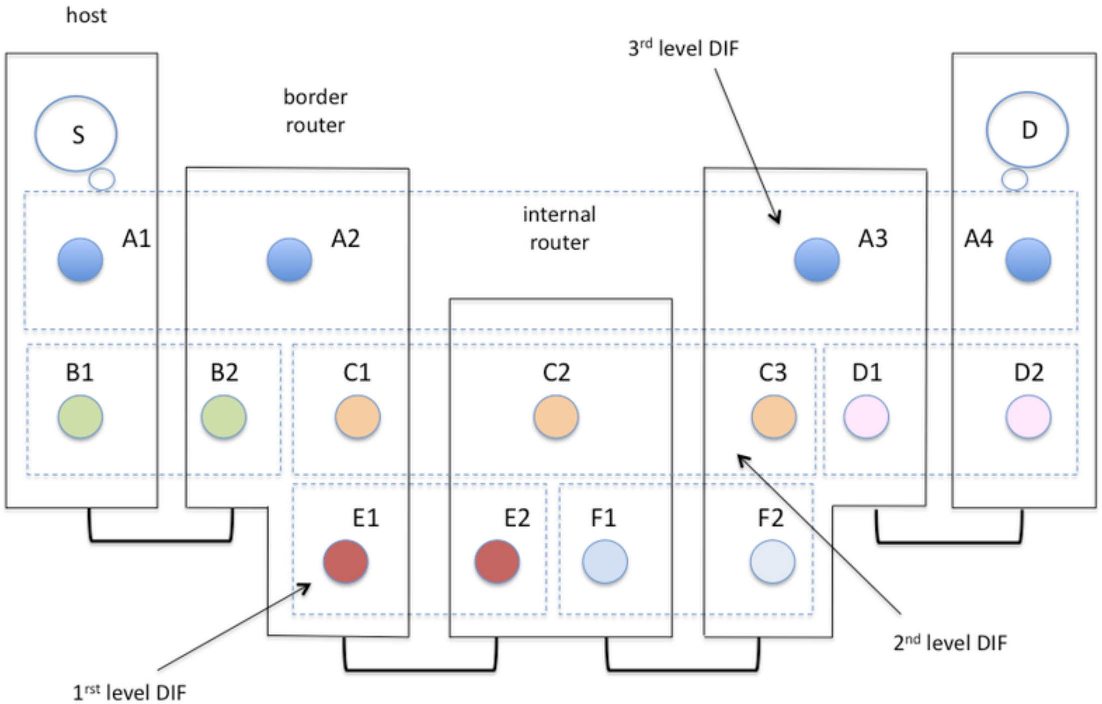


- A structure of recursive layers that provide IPC (**Inter Process Communication**) services to applications on top
- There's a **single type of layer** that **repeats** as many times as required by the network designer
- **Separation of mechanism from policy**

- **All layers have the same functions**, with different scope and range.
  - Not all instances of layers may need all functions, but don't need more.
- A **Layer** is a **Distributed Application that performs and manages IPC** (a Distributed IPC Facility –DIF-)
- This yields a theory and an **architecture** that **scales indefinitely**,
  - i.e. any bounds imposed are not a property of the architecture itself.

# Naming and addressing in RINA

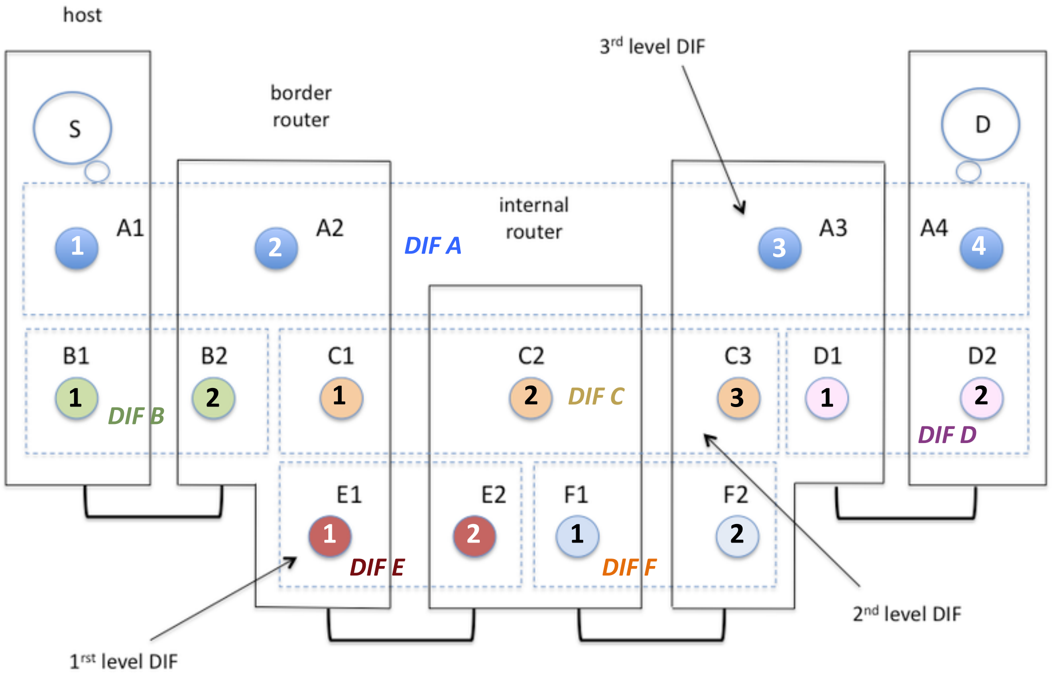


- All ***application processes*** (including IPC processes) ***have a name that uniquely identifies them*** within the application process namespace.
- In order to facilitate its operation within a DIF, ***each IPC process within a DIF gets a synonym*** that may be ***structured to facilitate its use within the DIF*** (i.e. an ***address***).

- The scope of an address is the DIF, ***addresses are not visible outside of the DIF***.
- The Flow Allocator function of the DIF finds the DIF IPC Process through which a destination Application process can be accessed.
- Because the ***architecture is recursive, applications, nodes and PoAs are relative***
  - For a given DIF of rank N, the IPC Process is a node, the process at the layer N+1 is an application and the process at the layer N-1 is a Point of Attachment.

# Architectural model

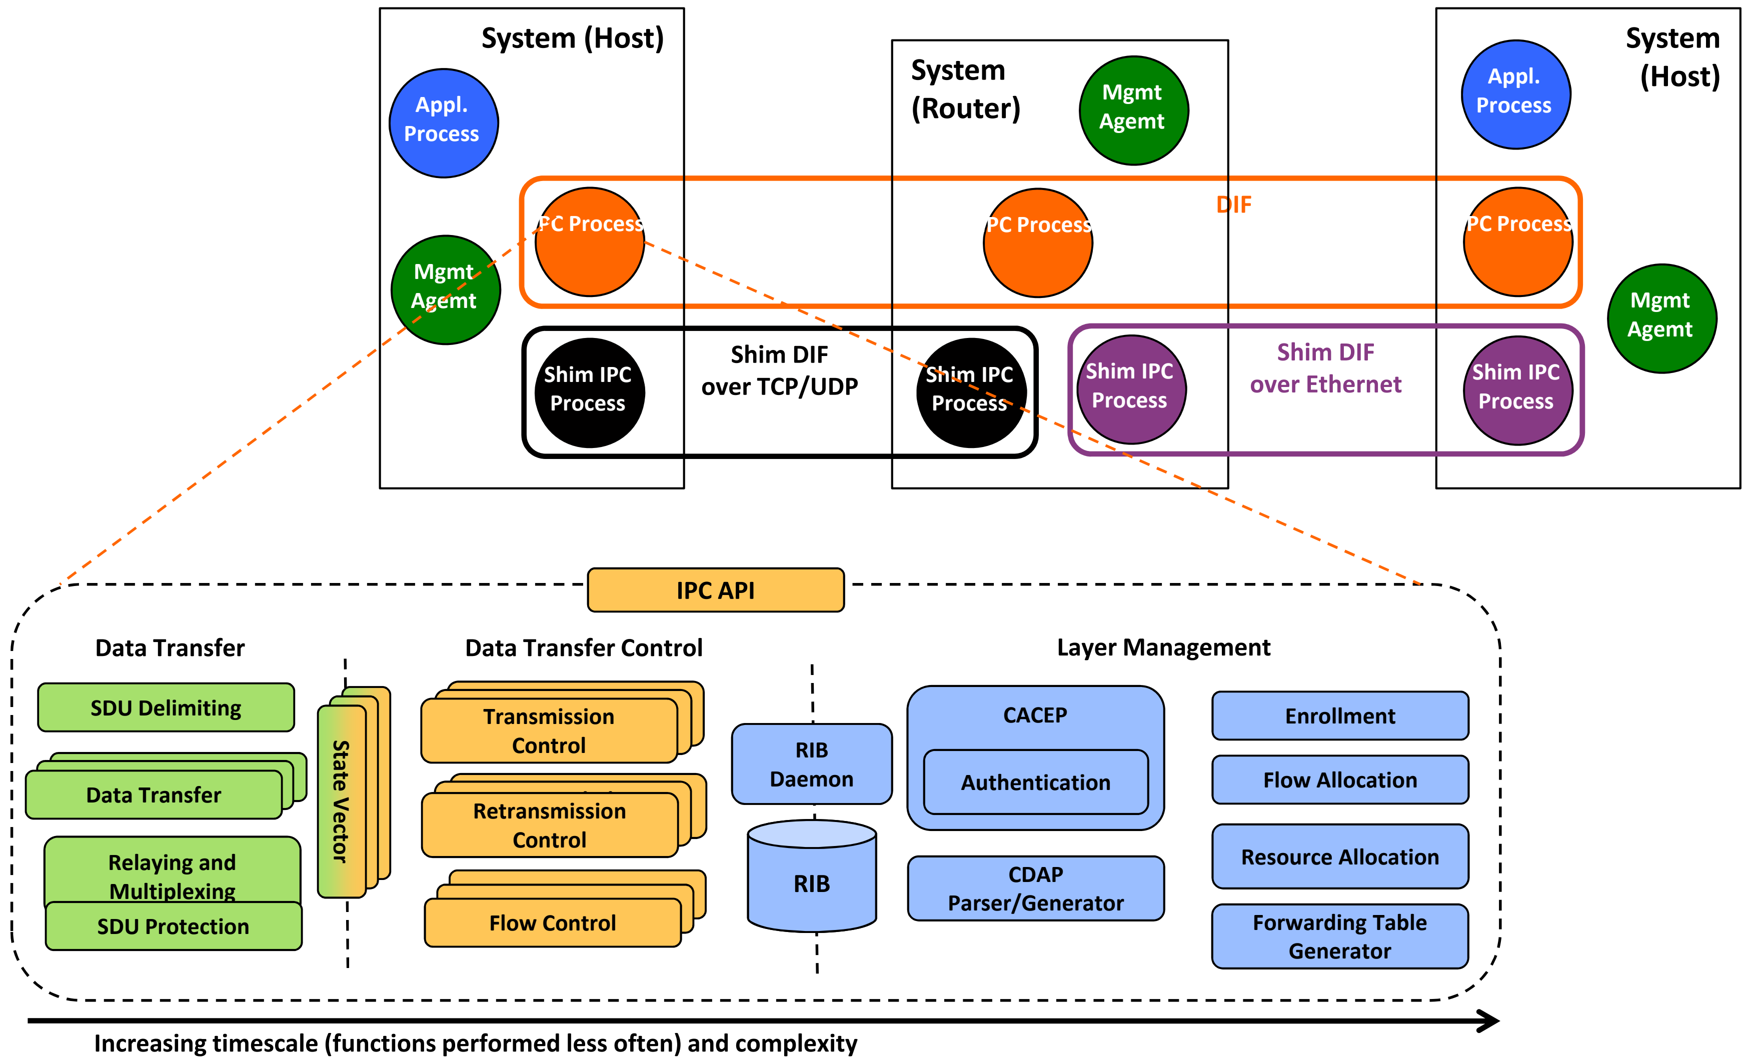

# Agenda

- What is RINA
- **Why researching RINA**
- Flow of research and development activities
- EC-funded RINA research
  - IRATI
  - PRISTINE (in negotiations)
  - IRINA (in negotiations)

# Why researching RINA (I)

- ***Architecture:***
  - <u>**Today**</u>: **5 layers**, layers "2.5", **layer violations**, "overlays", "virtual networks", "middleboxes" (NATs, firewalls, application-layer gateways) **Getting complex**!
  - <u>**RINA**</u>: **Repeating structure**, DIF (one type of layer, repeat as needed)
- ***Naming, addressing and routing***:
  - <u>**Today**</u>: **No independent application names**, no node names, just PoA names, **routing on PoAs** (multi-homing and mobility is hard to support)
  - <u>**RINA**</u>: **Complete naming & addressing**, routing on the node; **support for multi-homing and mobility without special protocols**. No need for global address space.
- ***Congestion control***:
  - <u>**Today**</u>: **Put in TCP**, not in the best place it could be, since **it maximizes the delay and variance** of the control loop (makes the system chaotic: self-similar traffic)
  - <u>**RINA**</u>: **Each layer can perform congestion control**, confining the effects of congestion to that layer. The delay and variance of control loops can be bound.

# Why researching RINA (II)

- ***Scalability:***
  - Today: Limited due to the **fixed number of layers** in the architecture
  - RINA: **Recursion provides** a divide and conquer approach, the way to **scalability**
- ***Security:***
  - Today: **No systematic approach** to security, secure each protocol or add boxes in between to improve security (firewalls).
  - RINA: **Strong design dictates where security functions go in the architecture** (encryption, authenticaiton, access control). DIFs are securable containers.
- ***Quality of Service:***
  - Today: **Best effort** is the dogma, applications cannot express desired outcomes
  - RINA: **Each DIF is free to provide different QoS classes**, using different policies for resource allocation, routing and data transfer. Applications can request the desired characteristics for a flow (delay, loss, ordering, etc)
- ***Management:***
  - Today: **Complex**, reflecting the complexity in the architecture and the **high number of protocols**.
  - RINA: The **commonality** in the structure **simplifies management** by orders of magnitude

# Agenda

- What is RINA
- Why researching RINA
- **Flow of research and development activities**
- EC-funded RINA research
  - IRATI
  - PRISTINE (in negotiations)
  - IRINA (in negotiations)

# Flow of RINA R&D activities
(feedback between activities not shown for clarity reasons)

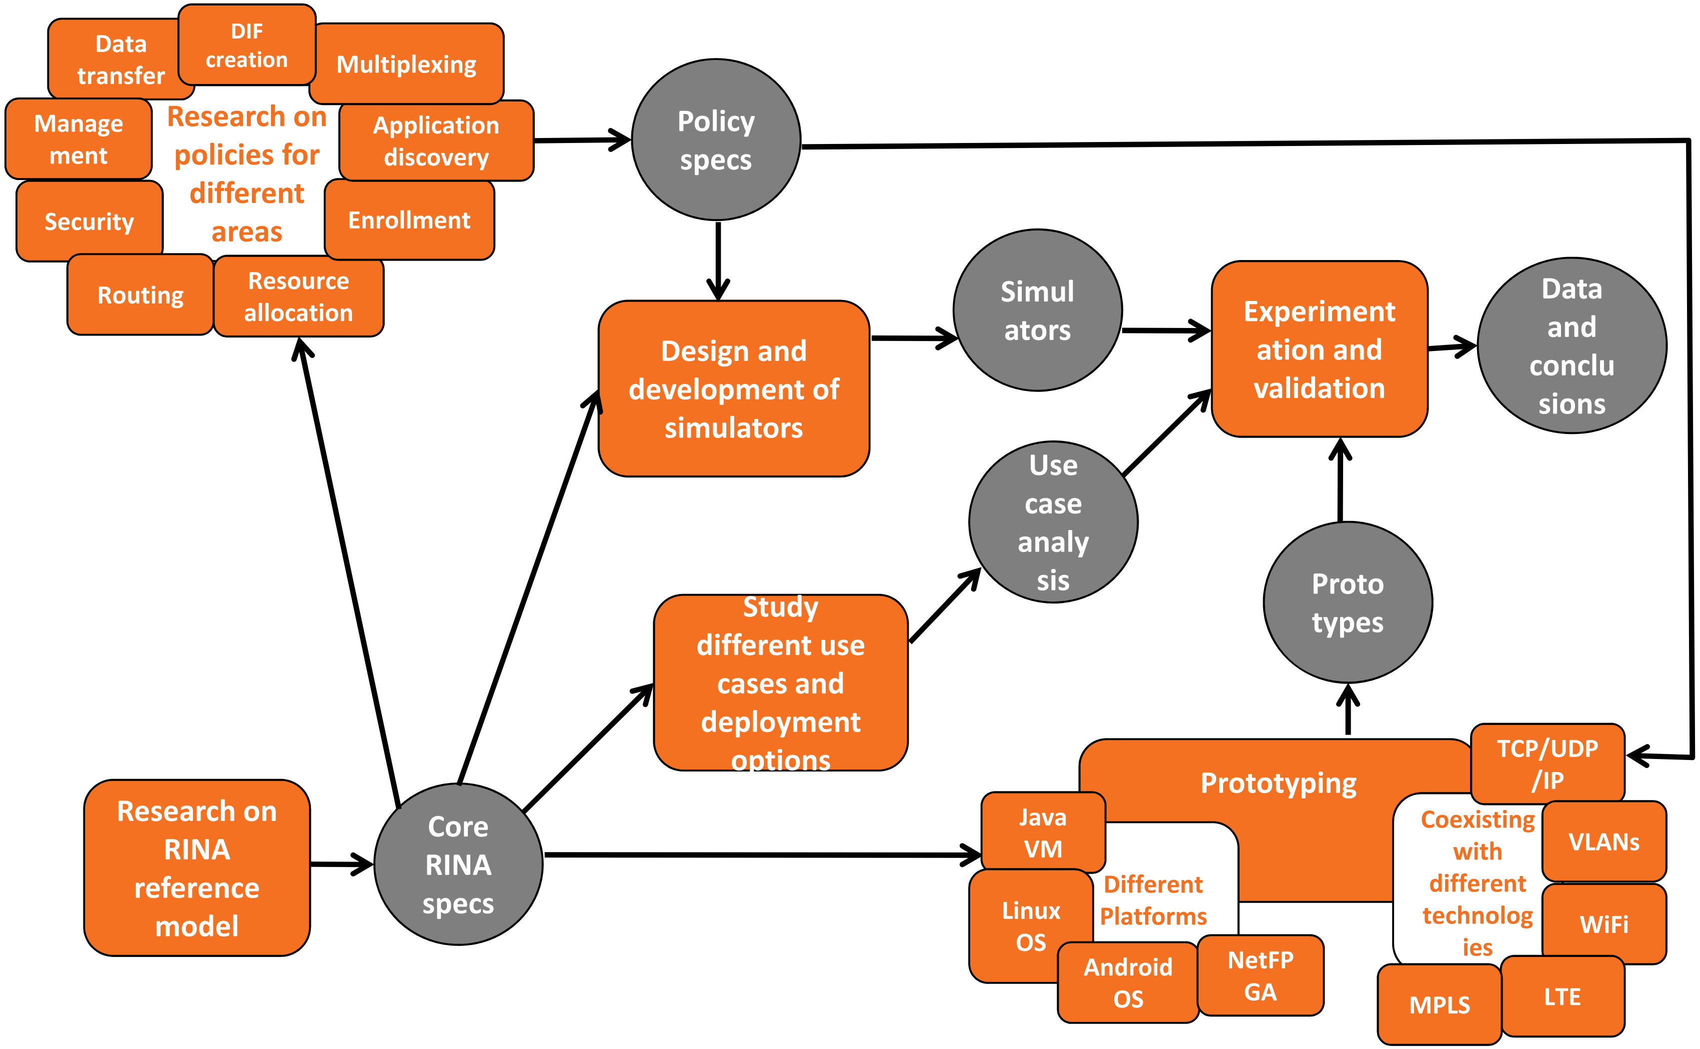

# Agenda

- What is RINA
- Why researching RINA
- Flow of research and development activities
- **EC-funded RINA research**
  - IRATI
  - PRISTINE (in negotiations)
  - IRINA (in negotiations)

# IRATI @ a Glance   http://irati.eu

IRATI — INVESTIGATING RINA

- What? Main goals
  - To **advance** the **state of the art of RINA** towards an **architecture reference model** and **specifications** that are **closer** to enable implementations deployable in **production scenarios**.
  - The **design and implementation** of a **RINA prototype on top of Ethernet** will enable the **experimentation** and **evaluation of RINA in comparison to TCP/IP**.

**Who? 4 partners**


**5 activities:**

- **WP1**: Project management (i2cat)
- **WP2**: Arch., Use cases and Req.
- **WP3**: SW Design and Implementation
- **WP4**: Deployment into OFELIA
- **WP5**: Dissemination, Standardisation and Exploitation

| **Budget** | |
|---|---|
| Total Cost | 1.126.660 € |
| EC Contribution | 870.000 € |
| **Duration** | **2 years** |
| Start Date | 1st January 2013 |
| **External Advisory Board** | |
| Juniper Networks, ATOS, Cisco Systems, Telecom Italia, BU | |

# IRATI contributions to RINA roadmap

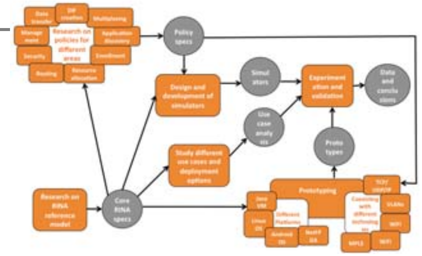

- **Reference model and core specifications**
  - **Detect** inconsistencies and **errors**
- **Research on policies for different areas**
  - **Routing** (link-state), Shim DIF over Ethernet VLANs (802.1q)
- **Use cases**
  - **Corporate VPNs** and **cloud networking**
- **Prototyping**
  - Initial implementation for **Linux OS** (user-space and kernel)
  - **Porting** of RINA implementation to **Juniper platforms**
- **Experimentation**
  - First experimental **analysis** of **RINA against TCP/IP** in similar conditions (focusing in LAN environments)

# Cloud/Network provider use case
(Introduction)

- RINA applied to **a hybrid cloud/network** provider
  - Mixed offering of connectivity (Ethernet VPN, MPLS IP VPN, Ethernet Private Line, Internet Access) + computing (Virtual Data Center)

interoute
from the ground to the cloud

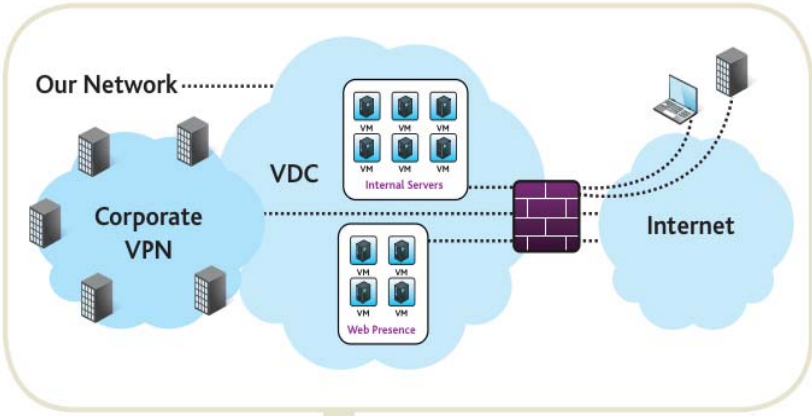


Access Network

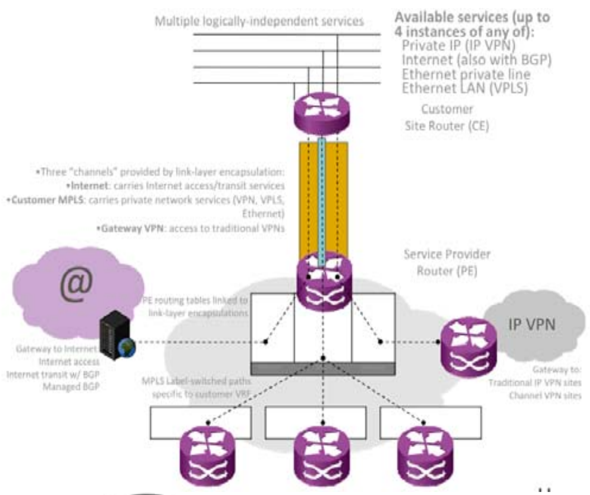


Wide Area Network

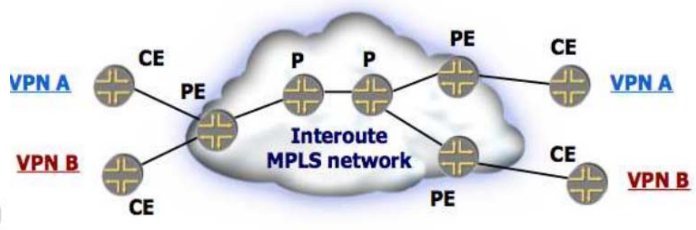


Datacenter Design

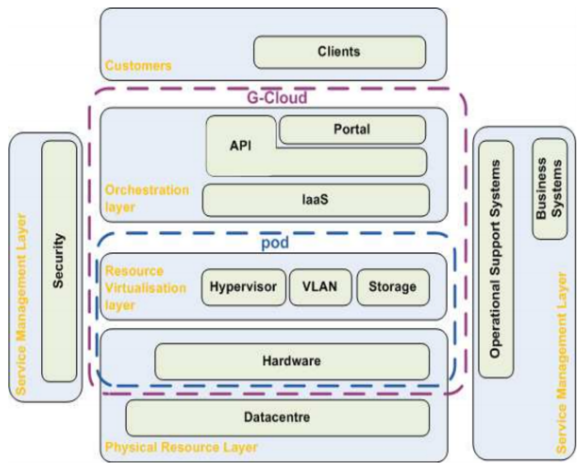

# Cloud/Network provider use case (Modeling)

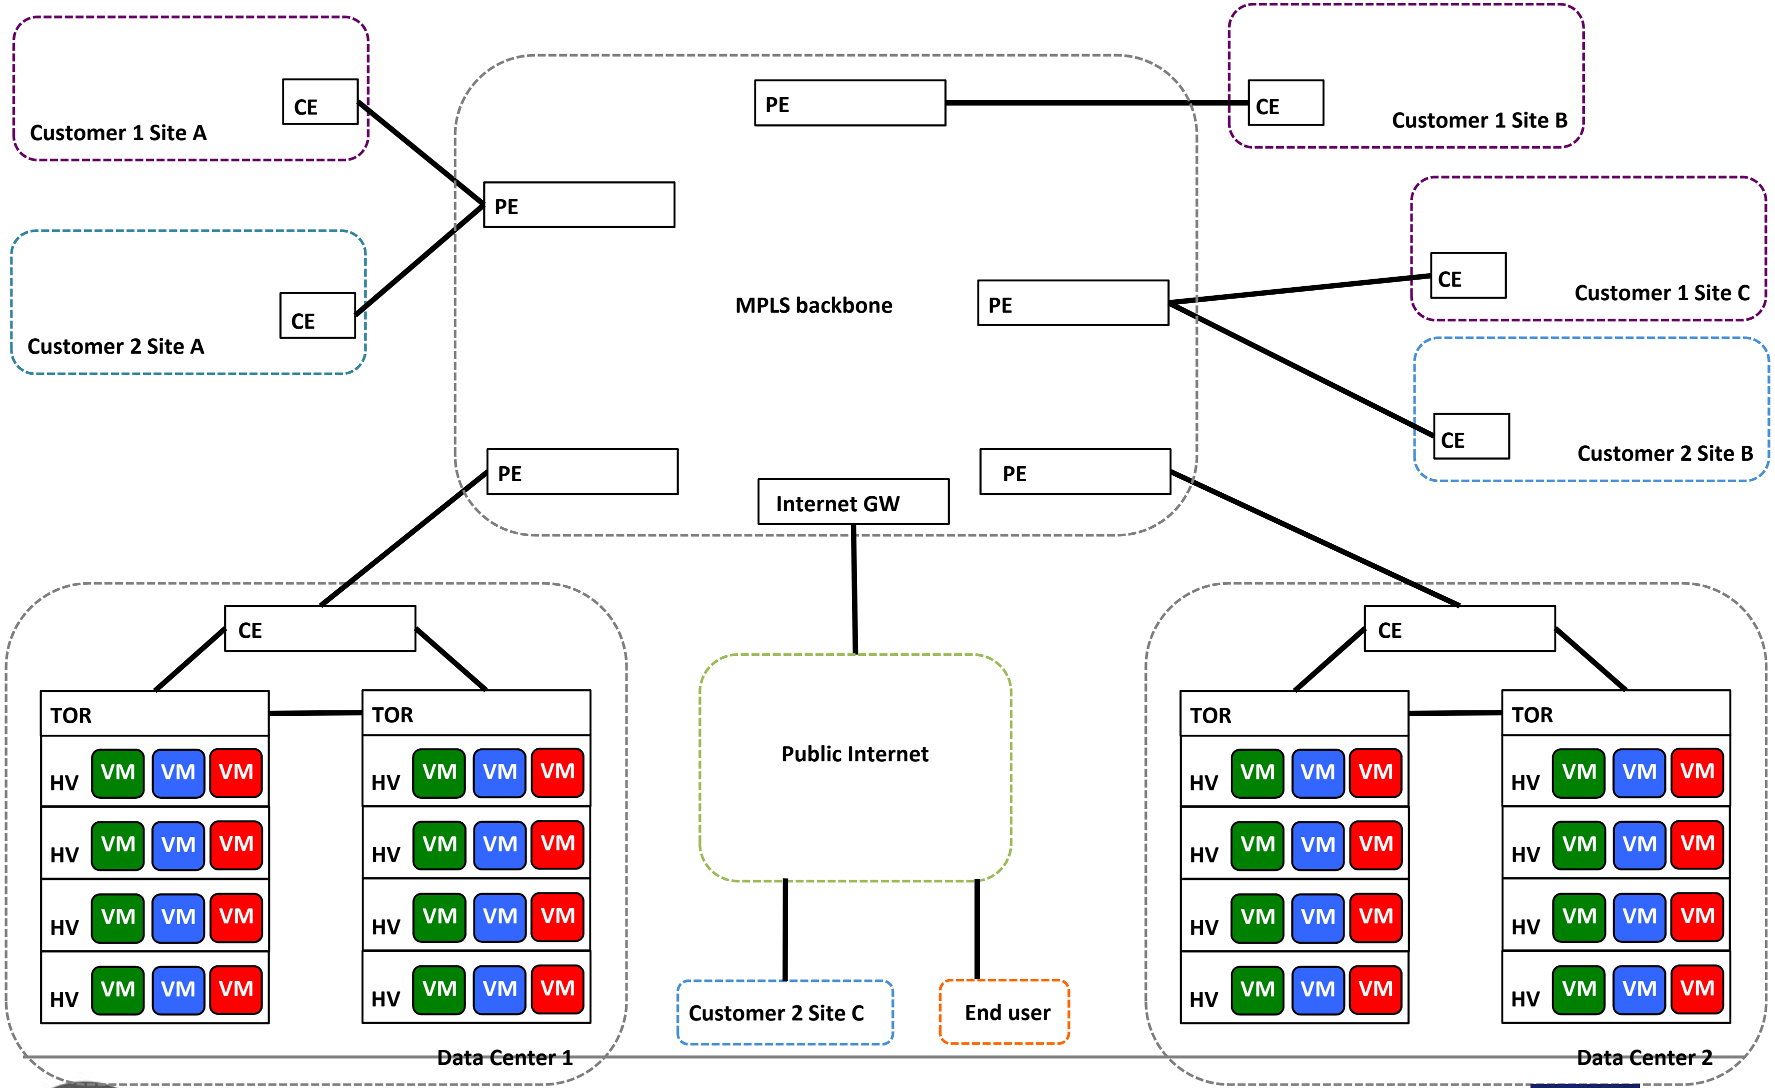

# Cloud/Network provider use case
(Applying RINA)

Scenario 1: Inter-DC

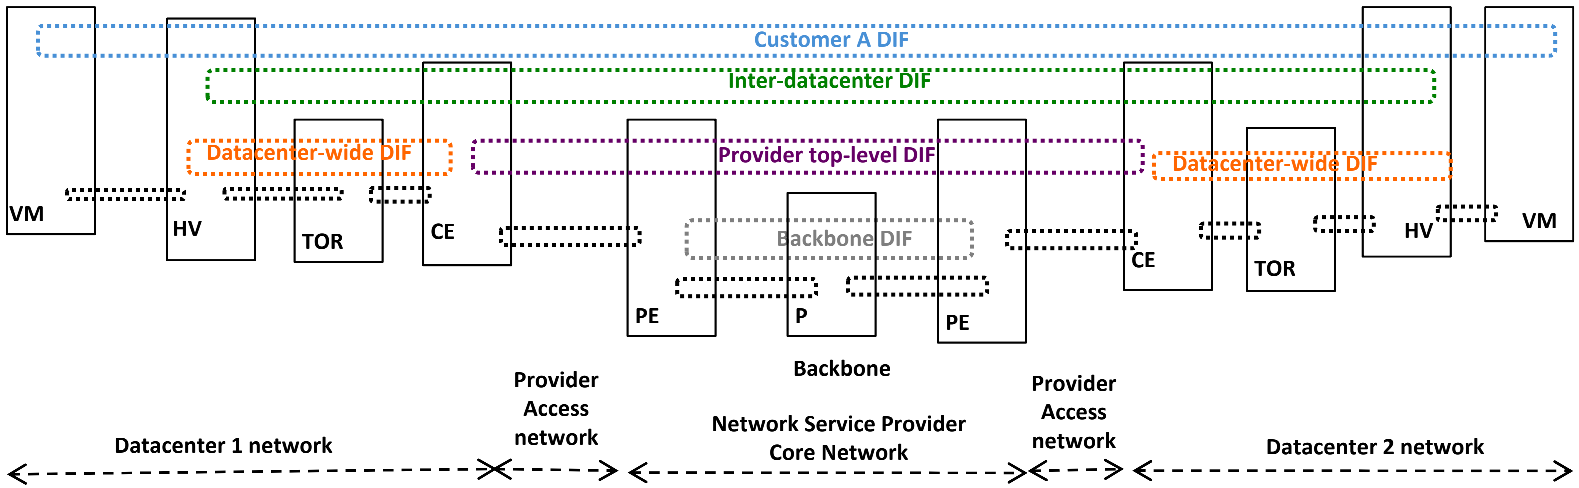


Scenario 2: DC-Customer

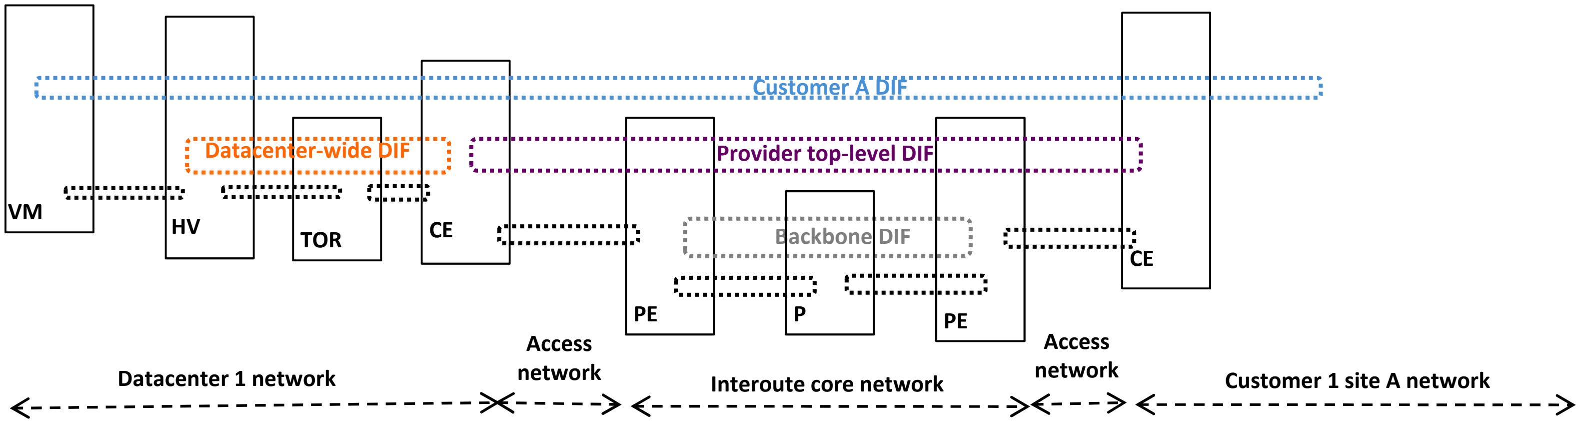

# Updated RINA specifications

- **New specifications** contributed to RINA during the first phase of the project
  - Shim DIF over Ethernet: Overlay RINA on top of IEEE 802.1q (VLANs)
  - Forwarding Table Generator based on Link-State routing technology (to compute the PDU Forwarding Table in medium-sized DIFs – 200 nodes aprox)
- **More updates to specs** foreseen during the next phases
  - Enrollment specification adapted to unreliable flows
    - Required to operate over Ethernet (also over UDP)
  - Concrete policies for data transfer
    - Sliding window flow control, retransmissions, rate-based flow control
  - Adaptation of routing update frequency to better support mobility

# Shim DIF over Ethernet

## General requirements

- The task of a **shim DIF** is to put a small as possible **veneer over a legacy protocol** to **allow a RINA DIF to use it unchanged**.
  - Not a RINA-conformant application. We are not trying to make legacy protocols provide full support for RINA.
  - Anything more should be provided by the first full DIF.
- The **shim DIF** should **provide no more service or capability than the legacy protocol provides**.

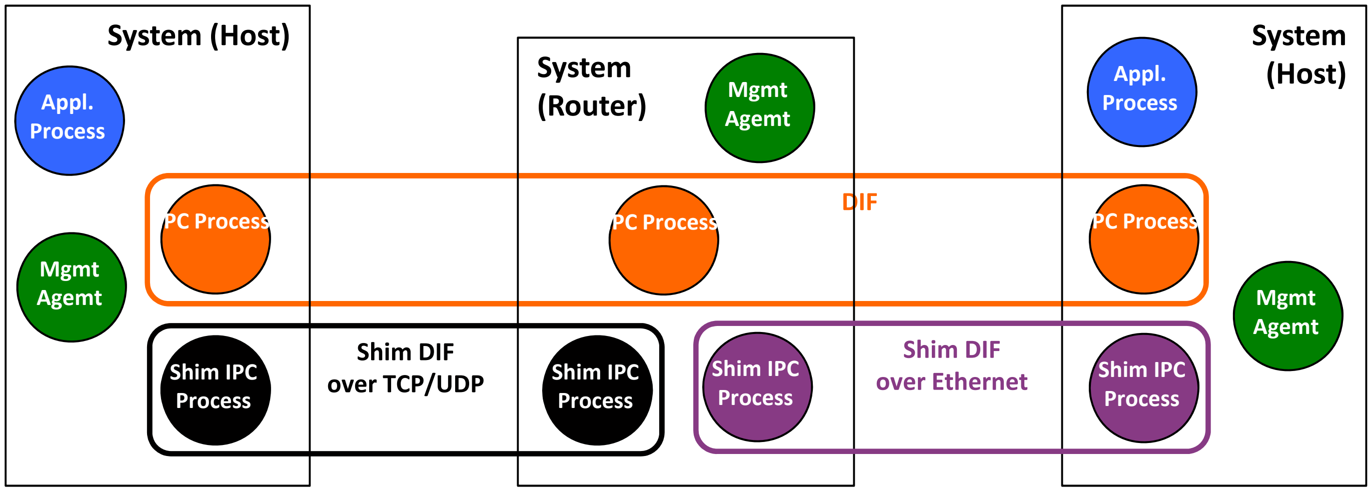

# High-level software architecture
## General requirements and choices

- **Linux** has been the chosen target platform for IRATI, due to
  - It is **widely used** in different contexts
  - Open source OS with a great community and documentation
- However the implementation aims to be as reusable as possible in similar environments
  - other UNIX-based Operating Systems
- The **implementation** targets both the **user-space** and the **kernel-space**, since
  - Low performance penalties have to be achieved for highly-frequent tasks (such as reading and writing data) -> Some components must be placed in the kernel
  - There is the need to access device driver functionalities in order to be able to overlay RINA on top of Ethernet (or other networking technologies in the future) -> Some components must be placed in the kernel

# Agenda

- What is RINA
- Why researching RINA
- Flow of research and development activities
- **EC-funded RINA research**
  - IRATI
  - PRISTINE (in negotiations)
  - IRINA (in negotiations)

# PRISTINE @ a Glance

- What? Main goals
  - To **design and develop an SDK for the IRATI RINA prototype, to unleash the programmability provided by RINA**.
  - To **use the SDK** to **design, implement and trial a set** of a **policies** to create **optimized DIFs** for each of the project **use cases: distributed cloud, datacenter networking** and **network service provider**.
  - To **design and implement** the first **RINA multi-layer management system**.

**Who? 15 partners**


WIT-TSSG, i2CAT, TID, Ericsson, NXW, Thales, Nexedi, Atos, BISDN, Juniper, Telecom SudParis, U Brno, UiO, CREATE-NET, iMinds

**7 activities:**

- **WP1**: Project management
- **WP2**: Use cases, req. analysis and programmable reference architecture
- **WP3**: Programmable performance-enhancing functions and protocols
- **WP4**: Innovative security and reliability enablers
- **WP5**: Multi-layer management plane
- **WP6**: System-level integration, validation, trials and assessment
- **WP7**: Dissemination, standardisation and exploitation

| **Budget** | |
|---|---|
| Total Cost | 5.034.961 € |
| EC Contribution | 3.337.000 € |
| **Duration** | **2.5 years** |
| Start Date | 1st January 2014 |
| **External Advisory Board** | |
| Cisco Systems, Telecom Italia, Deutsche Telekom, Colt Telecom, BU, Interoute | |

# PRISTINE contributions to RINA roadmap

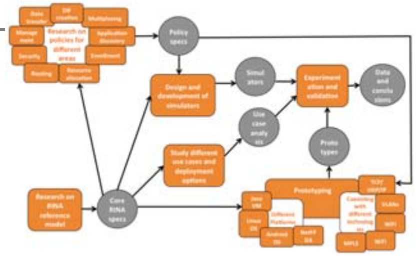


- **Reference model and core specifications**
  - Detect inconsistencies and errors
- **Research on policies for different areas**
  - **Congestion control**, **distributed resource allocation**, **addressing**, **routing**, authentication, **access control**, encryption, **DIF management**
- **Use cases**
  - **Decentralized cloud**, **datacentre networking**, network service provider
- **Prototyping**
  - **Build on IRATI implementation for Linux OS**. **Develop SDK** to allow easier customization, develop sophisticated policies with SDK. Prototype first DIF Management System
- **Experimentation**
  - **More realistic experimentation**, with more complex deployments, coexisting with **several technologies** at once (**IPv4, IPv6, Ethernet**), usage of business applications

# Use cases

- **Distributed cloud**
  - Decentralized cloud technology; customer's applications run in datacenters but also in servers from offices and home users.
  - Infrastructure interconnected through multiple ISPs, overall connectivity provided through overlay on top -> *Use RINA to provide this overlay*
- **Datacentre networking**
  - Evaluate RINA as a technology that allows **more dynamicity and tighter integration with applications** (dynamic instantiation of application-optimized VPNs)
- **Network Service Provider**
  - **Investigate benefits of RINA for NSP**: better network design, simpler management, DIFs that support different levels of QoS with stronger flow isolation, better security, programmability, etc.

# PRISTINE Infrastructure for trials

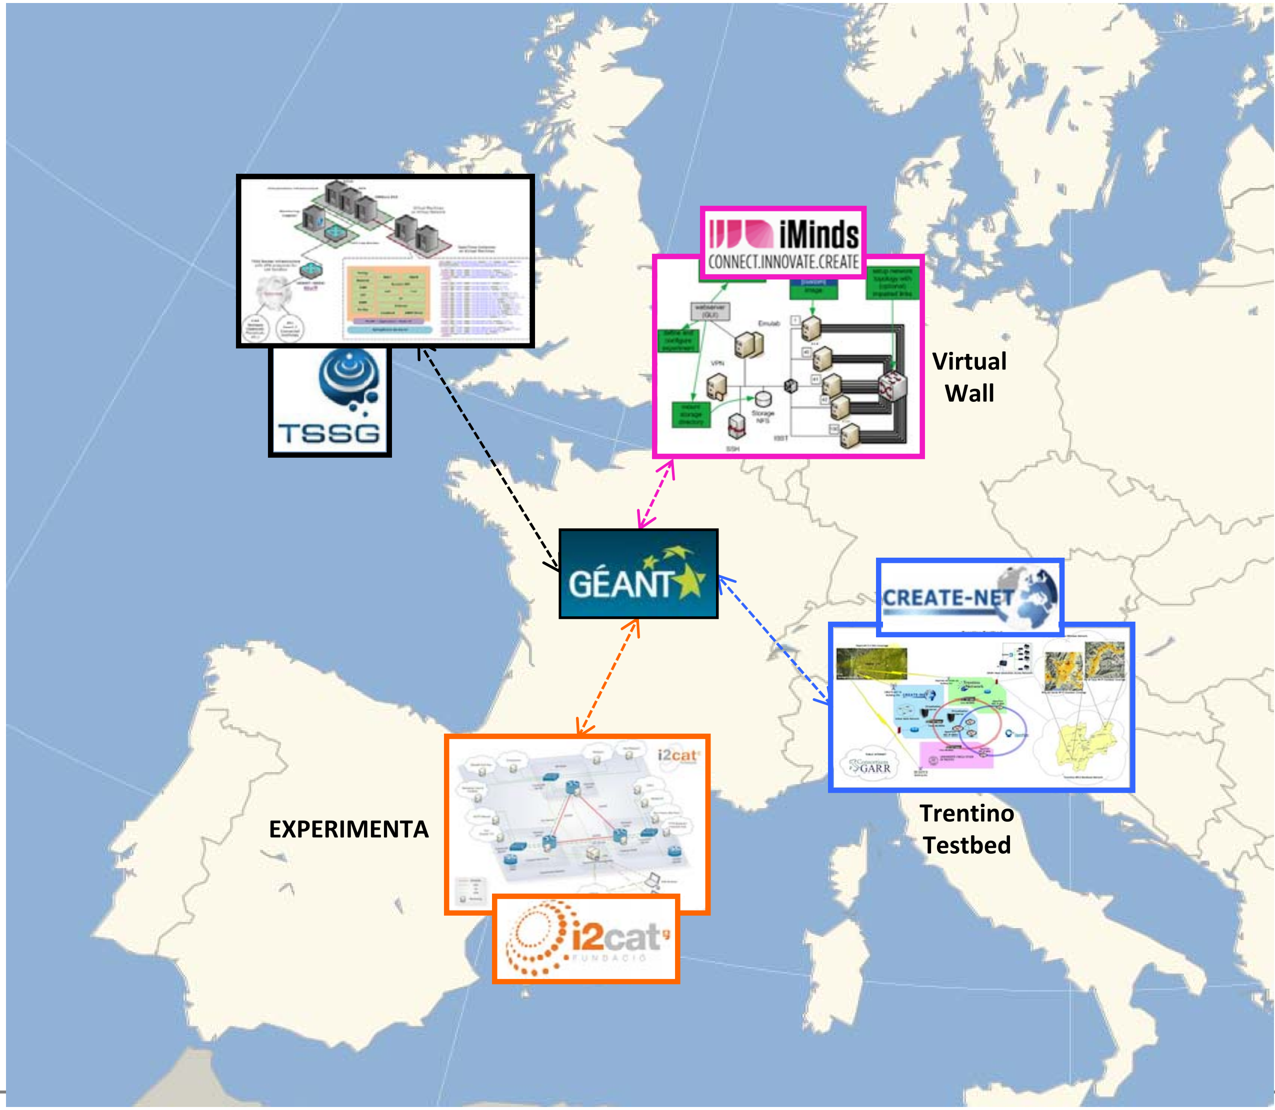

# Agenda

- What is RINA
- Why researching RINA
- Flow of research and development activities
- **EC-funded RINA research**
  - IRATI
  - PRISTINE (in negotiations)
  - IRINA (in negotiations)

# IRINA @ a glance

- What? Main goals
  - To make a **study of RINA against the current networking state of the art and** the most relevant **clean-slate architectures** under research.
  - To perform a **use-case study** of how **RINA** could be better used in the **NREN scenario,** and **showcase a lab-trial** of the use case
  - To **involve the NREN and GEANT community** in the different steps of the project, in order to to get valuable feedback

**Who? 4 partners**


**5 activities:**

- **WP1**: Technical coordination and interaction with GEANT3+
- **WP2**: Comparative analysis of network architectures
- **WP3**: Use case study and lab trials
- **WP4**: Dissemination and workshop organization

| Budget | |
|---|---|
| Total Cost | 199.940 € |
| EC Contribution | 149.955 € |
| **Duration** | **18 months** |
| Start Date | 1st November 2013 |

# IRINA contributions to RINA roadmap

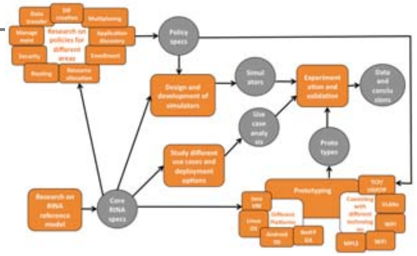

- **Reference model and core specifications**
  - **Compare with other clean-slate architectures**
- **Use cases**
  - Research network operators (NRENs and GEANT environment)
- **Prototyping**
  - Little adaptations to the IRATI prototype (Linux OS), to be able to trial the **use case in the lab**
- **Experimentation**
  - Focus on the **requirements of NRENs**

# GEANT and NRENs use case

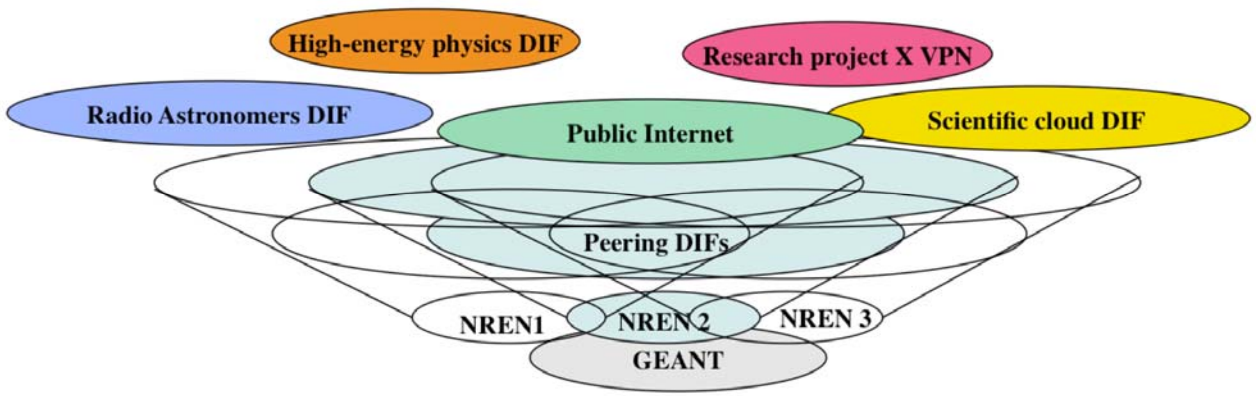


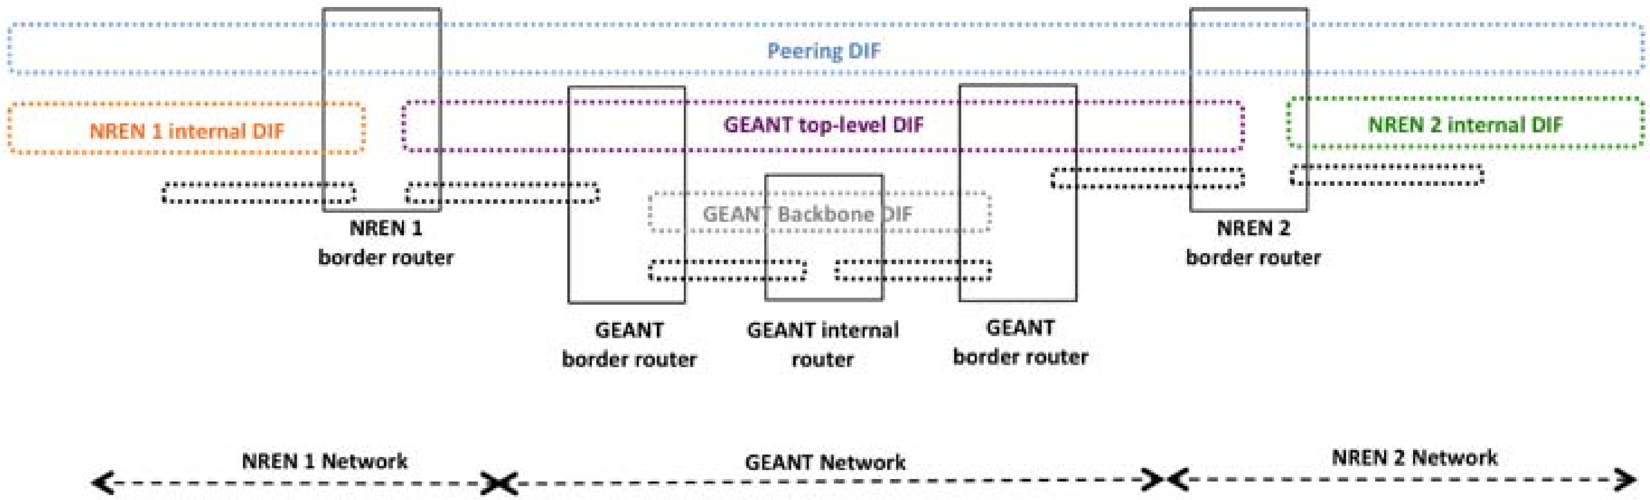

*Many Thanks !*

*Moltes gràcies !*

**Sergi Figuerola,**
Innovation and Technology Director
(sergi.Figuerola@i2cat.net)

**Eduard Grasa,**
RINA research line leader
(eduard.grasa@i2cat.net)

**13th annual Global LambdaGrid Workshop October 4th, 2013**

**http://www.i2cat.cat**
**http://dana.i2cat.net**
**http://irati.eu**