

**Federated E-infrastructure Dedicated to European Researchers
Innovating in Computing network Architectures**

Co-funded by the European Commission within the Seventh Framework Programme. Grant Agreement No. RI-213107

Deliverable JRA2.4: Final prototype testing

Version 1.3

Dissemination Level	Public
Contractual Date of Delivery	31 August 2010
Actual Date of Delivery	20 December 2010
Editor	Josep Pons – I2CAT
Contributors	Eduard Grasa – I2CAT Jean-Marc Uzé – Juniper Björn Rhoads – KTH (FEDERICA NOC)
Reviewers	Peter Szegedi – TERENA Mauro Campanella – GARR Martin Pots – Martel

Abstract

Deliverable DJRA2.4 “*Final prototype testing*” is a brief technical report on the latest IPsphere-MANTICORE interoperability prototype tests (initially planned as DJRA2.4 “*User controlled services in virtual infrastructures*” to be found in the FEDERICA Description of Work). It contains a summary of the efforts mainly done by the participants of the task TJRA2.3 in the final months of the FEDERICA project.

The main objective of the activity TJRA2.3 “*User controlled service provisioning*” was to investigate a potential collaboration framework between FEDERICA and IPsphere, and to validate its prototype (that was developed by the task participants and reported in DJRA2.2) for integrating the MANTICORE software (based on the IaaS Framework, a descendant of the UCLPv2 project) in the IPsphere model. The final prototype validation work, its objectives and results are described in this deliverable.

Executive summary

This document summarises the work which was completed in the last phase of the task TJRA2.3 “*User controlled service provisioning*” until the end of the extended project.

The objective of this task was related to the establishment of a technical collaboration framework between FEDERICA and IPsphere. The goal was the achievement by prototyping a gateway for integrating the MANTICORE [Man] implementation in the IPsphere model. DJRA2.2 “*Prototype for interoperability between IPsphere and MANTICORE*” was a prototype deliverable aimed to make a working FEDERICA-IPsphere interoperability prototype available. The prototype was already delivered on the 15th of May, 2009 (as Milestone MJRA2.1), and it can be downloaded from the FEDERICA Wiki:

https://wiki.fp7-federica.eu/pub/NA1/Deliverables/DJRA2.2_Prototype_v2.zip

The main technical challenge was to develop a prototype implementing the SMS Child Module of the IPsphere model. The SMS Child is eventually capable of interacting with other modules of the IPsphere and also with the MANTICORE IP Network service. The final step of this task summarised in this deliverable was to validate the prototype over the FEDERICA virtualised infrastructure by means of a slice interconnected to the i2Cat PoP. JRA2.3 requested a slice from FEDERICA (i.e., the joint research activity became an internal user of the FEDERICA infrastructure services) and performed the tests using the slice.

The validation tests have been done closely together with the SA2 (in particular, the NOC) partners with the goal that once the prototype becomes mature enough it can be incorporated into the service activates allowing the automated provisioning of interconnected slices by the NOC and eventually by the users themselves.

It can be concluded that the prototype, designed and developed following the IPsphere Forum specifications, has a real and positive application over experimental facilities like FEDERICA. It was also demonstrated that the IPsphere Framework is capable of interoperating with other Management Tools or Frameworks located in the layer between itself and the physical infrastructure.

Table of Contents

Executive summary	2
1 Introduction and objectives of the validation work	4
2 Validation scenario.....	5
2.1 FEDERICA Jra23 slice.....	6
2.2 I2Cat PoP	7
2.3 Interconnection between the two facilities	8
2.4 Complete scenario	9
3 Validation process and test results.....	11
4 Conclusions and further work	14
References.....	16

1 Introduction and objectives of the validation work

The IPsphere – FEDERICA interoperability prototype ‘s (described in Deliverable DJRA2.2) goal is to allow connection creations between FEDERICA slices using the IPsphere Framework [Ips], which is capable of negotiating end-to-end paths across several domains (may be represented by slices).

The prototype is based on the implementation of one of the main components of IPsphere, called SMS Child. This component has two main functions. The first one is to publish Element Offers originated from Element Templates to the IPsphere registry. The second function is to manage the element activation/deactivation requests received from another higher level IPsphere component: the SMS Parent.

The SMS Child also has to interact with a Management System (xMS) that can manage network elements and issue the required commands to the network devices in order to activate the requested Element. The xMS configures the different devices involved in the Connectivity Element (CE). The CE creates a path between the two edge machines connected to the slices. The routers of this path have to be capable of modifying the QoS parameters as the BW, for instance, in order to allow or filter flows of different types of content between their domain hosts.

Each domain can have a different but compatible management system. The SMS Child prototype (Fig. 1.) has to be capable of interacting with two different Management systems: MANTICORE and the Juniper SRC [Jun].

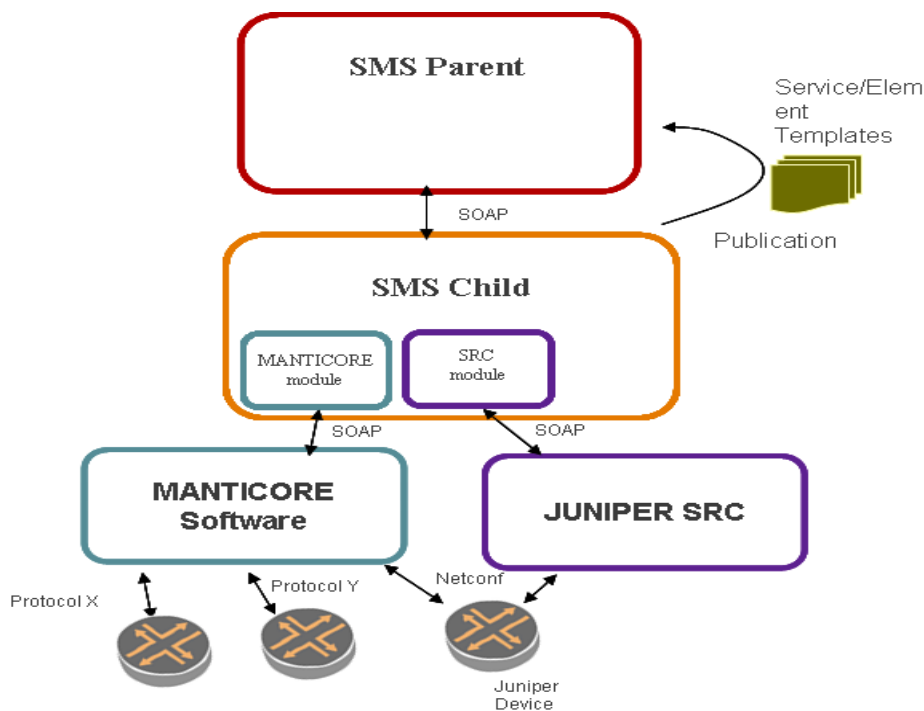


Fig. 1 Interoperability prototype general architecture

The final objective of the SMS Child implementation was to build a SMS Child capable of publishing and activating the elements of the i2Cat-FEDERICA slice by interoperating with the MANTICORE xMS. But this interoperability with MANTICORE was effective at a second phase. Prior to this, two scenarios were used to validate the implementation of the SMS Child and both used the Juniper SRC as xMS.

In the second phase, an enhancement to the MANTICORE software was done in order to add new functionalities related to the application of filters and policers on the router interfaces. After installing MANTICORE in one of the networks scenario, the SMS Child could talk with the IP Network service of MANTICORE in order to request information about the network for creating the elements, and to send the activation requests for reserving them. The test bed in this phase was completed with Deutsche Telekom's network in Berlin connected to the i2CAT Labs through a GRE tunnel for the Data Plane and with a simulator of an SMS Parent developed as part of the current prototype for the IPSphere Control Plane. MANTICORE tools were used as xMS and the tests were done using only equipment from i2CAT Labs and simulating a scenario with two networks using the virtualization capabilities of a Juniper M7i router.

In the Phase II scenarios, the prototype was validated by performing a modification of the QoS parameters of the end-to-end path when specific element activations were performed. Once the elements had been activated, the bandwidth of the path was increased and a premium quality video could be streamed perfectly.

After this second phase, the effort of the TJRA2.3 task related with this activity was focused on the final validation of the interoperability prototype using a FEDERICA slice and another network, consisting of FEDERICA equipment located at the i2Cat Point of Presence (PoP). Basically these final tests attempt to transmit premium content traffic after the activation of two IPSphere based elements from one host of the FEDERICA slice to a host belonging to the i2CAT network or vice versa.

The complete scenario used for this validation work is presented and explained in the following section of the document. In order to setup the validation scenario, a new FEDERICA slice has been requested from the User Policy Board (UPB) that was eventually delivered.

2 Validation scenario

In order to perform tests based on enabling paths (validated with content transmissions through them) between several hosts belonging to different provider infrastructures, we decided to use two sites: a FEDERICA slice and the FEDERICA equipment located at the i2Cat PoP, the latter playing the role of a different provider network.

2.1 FEDERICA Jra23 slice

The first step used to create this validation scenario was requesting a FEDERICA slice composed by a set of Virtual Machines connected to a pair of Routers. This slice, called JRA23, was requested formally from the UPB of the FEDERICA project.

A star topology was requested with several VMs with Windows XP running over them (in order to act as Host-Test machines) and a couple of them with Linux Ubuntu in order to deploy on these the MANTICORE software. Two Logical Routers that are partitions of Juniper Physical Routers acting as a single router were also requested. At that time, MANTICORE could only manage routers of this manufacturer therefore, Juniper Routers were required. The Jra23 slice can be seen in Fig. 2.

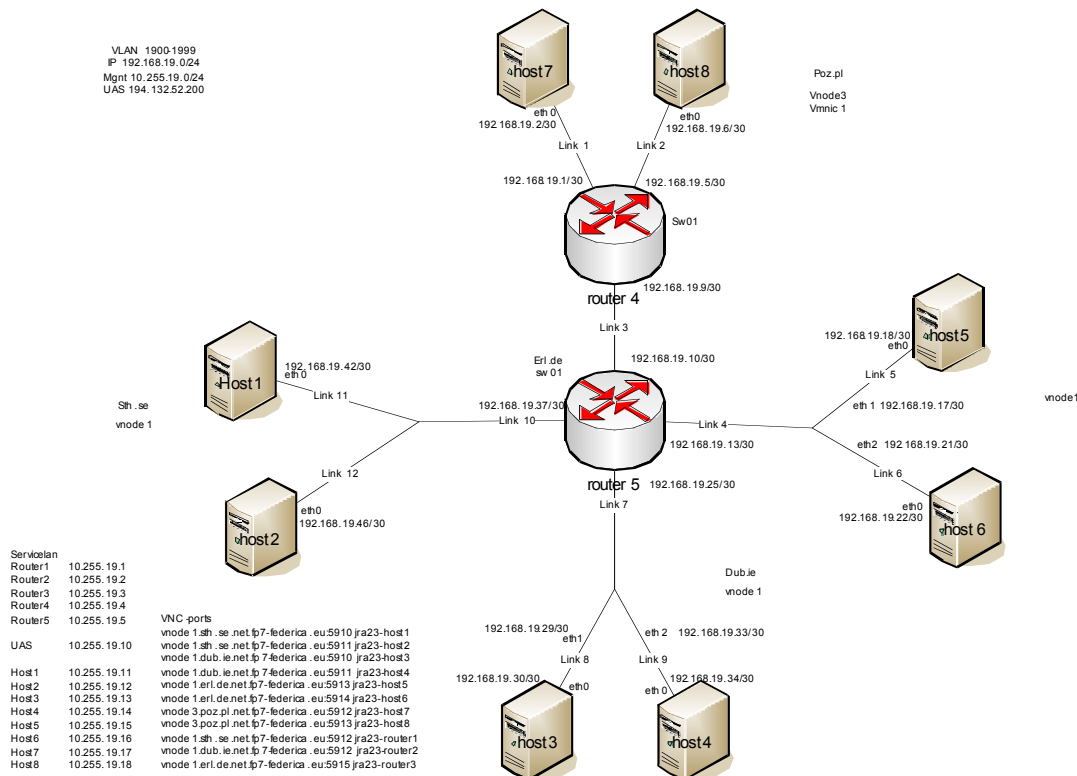


Fig. 2 FEDERICA Slice Jra23

To allow users access to their routers and virtual machines, a User Access Server (UAS) accessible through a public internet address is installed to act as a bridge. A user can access all the routers of this slice and also the Linux machines by opening a Secure Shell (SSH) connection from the UAS to the machine in which it wants to connect. However, the user must first be connected to the UAS and also through SSH. Then he can connect to the corresponding IP addresses, which range from 10.255.19.1 to 10.255.19.18. Furthermore, the user can connect through VNC to the Host machines that act as Test machines (not to MANTICORE servers). To do this, the user can tunnel VNC ports over SSH with a command such as the following:

DJRA2.4: Final prototype testing



➤ `ssh -L5901:vnode1.sth.se.net.fp7-federica.eu:5910 jra23@jra23.net.fp7-federica.eu`

, where:

- 5901 would be the local port used to connect from the localhost to the remote machine
- `Vnode1.sth.se.net.fp7-federica.eu:5910` is the hostname:port of the Host machine where we want to connect.
- `Jra23` is the Username used to connect to the UAS.
- `Jra23.net.fp7-federica.eu` is the hostname of the UAS.

For instance, this tunnel makes the `127.0.0.1:5901` to be `jra23-host1` VNC port. The corresponding VNC ports configured in the machines to allow remote connections are listed in the previous picture.

The range of private addresses used internally for the Data plane of the Slice is `192.168.19.0/30`. All the hosts of the slice can see each other thanks to configuring the corresponding routing protocols on the routers of the slice.

2.2 I2Cat PoP

As is explained in the previous chapters, the other infrastructure used in order to build the complete scenario is the i2Cat PoP of the FEDERICA infrastructure.

The i2Cat PoP is composed of one Juniper MX 32000 Switch (called `sw01.i2c.es`) and two Sun Fire X2200 M2 Servers (`vnode1.i2c.es` and `vnode2.i2c.es`). The switch is connected to each Server through 8 Gigabit Ethernet interfaces (Fig. 3).

In order to perform the tests, one Virtual Machine called TestX with Linux Ubuntu 10.04 Desktop located at the Vnode1 was created.

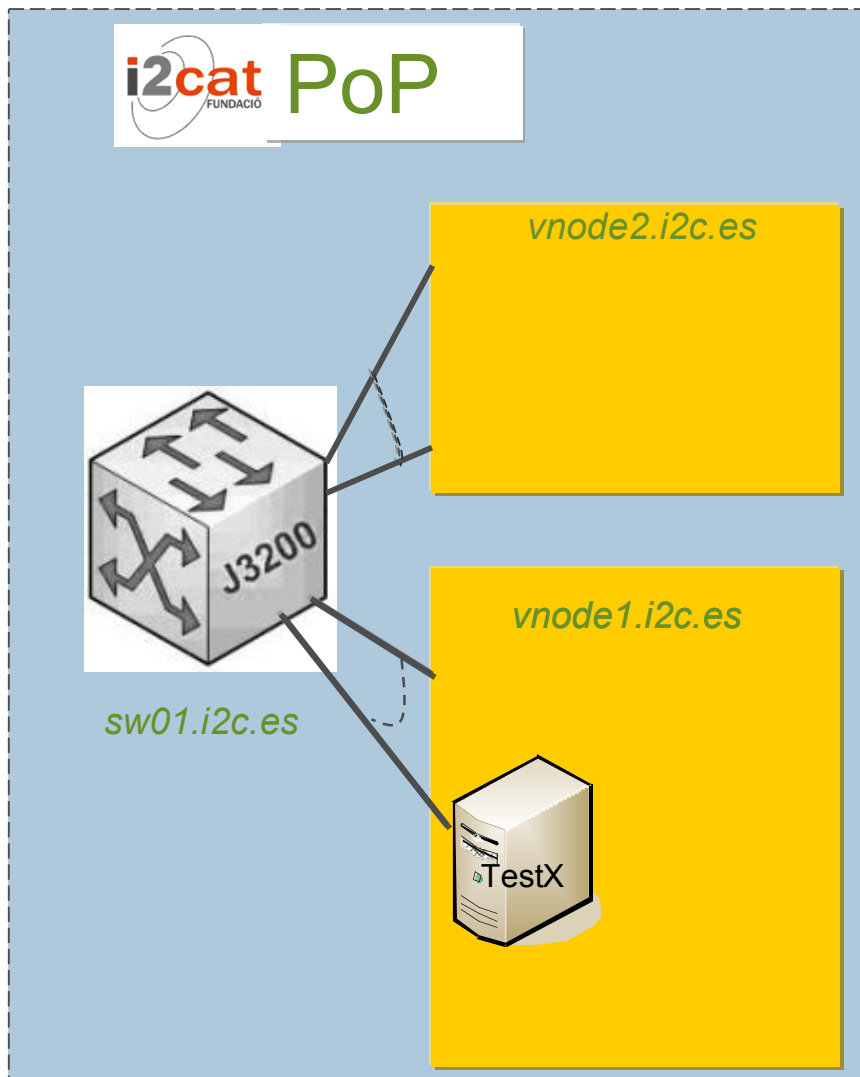


Fig. 3 i2Cat PoP

2.3 Interconnection between the two facilities

As the next major requirement, these two infrastructures had to be physically connected. In principle, this should not be a problem as shown in Fig. 4, the FEDERICA substrate is already physically interconnected.

Federica Substrate
Core network only

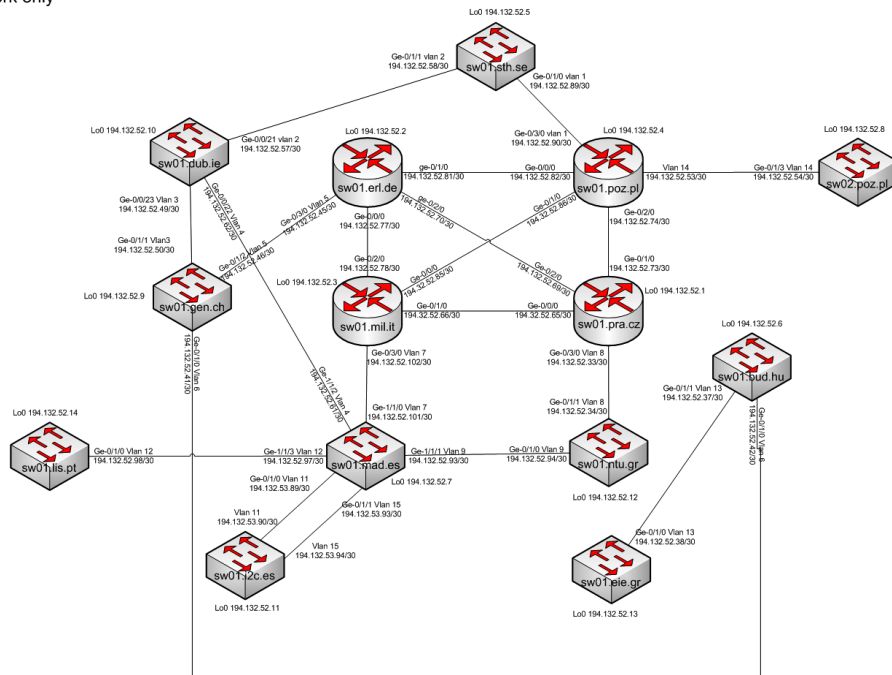


Fig. 4 FEDERICA substrate

federica-noc@fp7.federica.eu

The way to achieve logical connectivity between the i2Cat PoP and the slice is to setup an unfiltered link between the i2Cat switch in Spain (entry point of the FEDERICA i2Cat PoP) and one of the physical routers of the Jra23 slice (sw01.i2c.es as entry point of the slice). According to the figure shown in the Jra23 slice description, this router can be located in Poznan (sw01.poz.pl).

2.4 Complete scenario

Once the two infrastructures had been interconnected, only a few issues remained unsolved. The first was the internal connection between the *sw01.poz.pl* router and the logical router created inside it (router4) and assigned to the Jra23 slice. This was accomplished by a logical tunnel which was configured in the physical router to connect the root router with the logical router.

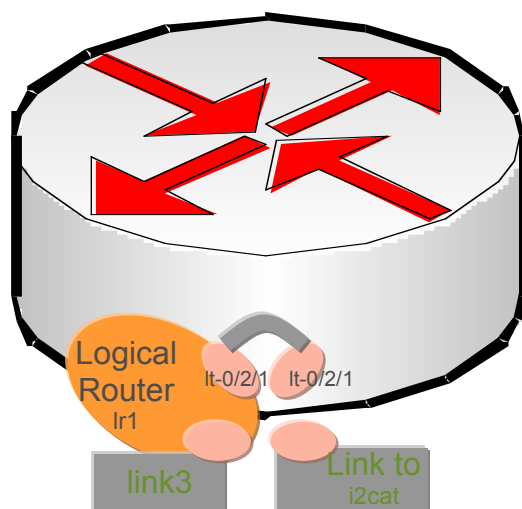


Fig. 5 Details of the sw01.poz.pl router with its internal Logical router 5 interconnected by a logical interface

To finalize the scenario construction, the last step was to deploy MANTICORE software in one of the VMs which had Linux OS installed and also to deploy the SMS Child on the other VM which had the Eclipse Platform installed.. The SMS Child Machine had to be connected to the MANTICORE machine, and this last machine connected with both routers (logical routers) of the slice using the Management Plane IP Addressing scheme (different from the Data Plane Addressing) in order to avoid any collision problem. Only the Jra23 Slice side contained routers therefore, only one MANTICORE instance was deployed to control the routers in the slice, and none for the i2Cat PoP. There was no problem in performing these tests because a connectivity path between the two sides interacting only with one of them could be activated or deactivated. However, we wanted to extend this use case to the real world where each side was owned by different providers. Then in this case, each side would have to deploy an independent instance of MANTICORE and also an independent instance of the SMS Child.

In Fig. 6, the complete scenario used to perform the validation tests is presented.

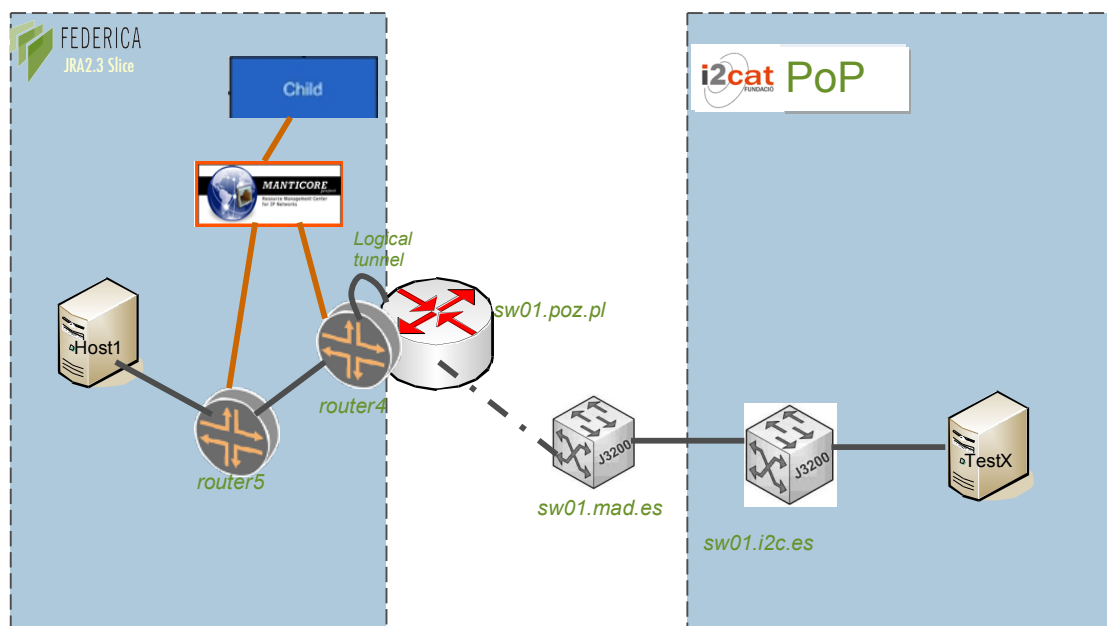


Fig. 6 Complete validation scenario

3 Validation process and test results

Once the scenario has been established the validation tests can start. In order to show the interoperability between the implemented IPsphere components and the FEDERICA infrastructure, we consider the activation of a previously filtered path between the two sides. The activation sequence starts with an invocation of an activation message from the SMS Parent simulator to the SMS Child. Prior to this, the SMS Child has to have published an Element Template containing the data of the Activation Service that MANTICORE needs to execute. Once the SMS Child receives the activation message from the Parent simulator, it invokes MANTICORE in order to activate the corresponding un-filtering service. Finally, MANTICORE modifies the necessary filters on the router interfaces involved in the path to provision the connectivity. To demonstrate the activation of the path, a Video Transmission has been streamed from Host1 to Host X. The complete sequence of the validation test is as follows:

1. Default policer and filter were defined and applied in the *router4* input interface in order to limit the BW for the transmissions through the defined path (Fig. 7).

```

    }
  }
}
firewall {
  policer limitBW {
    if-exceeding {
      bandwidth-limit 5m;
      burst-size-limit 1m;
    }
    then discard;
  }
}
[edit]

```

Fig. 7 Default policer applied in the router4 input interface

- In that moment it was verified that the traffic sent (Video streaming) from Test1 to TestX was experiencing packet losses. In Fig. 8, the reception of the video in the target Host (TestX) at the i2Cat PoP can be seen. Also the fact that the Input bit rate was under 5Mb/s (this was the BW limit configured in the previous step) can be observed.

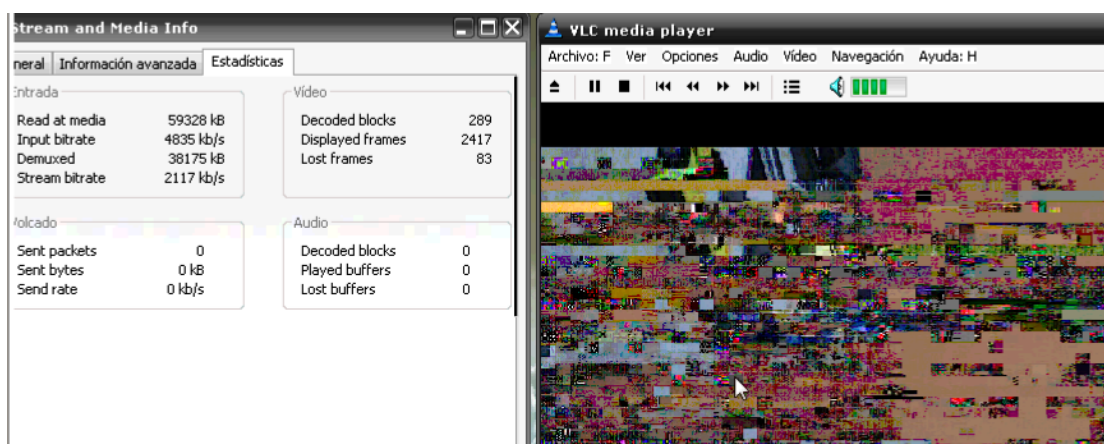


Fig. 8 Video reception before unfiltering the path (there is the default BW limitation)

- An *ElementTemplate* that mapped the *policer_modification* Service on MANTICORE was published from the SMS Child to the Publishing Registry.
- In order to start the activation of the path, the SMS Parent simulator (*Test_SMSChild* module) requested the activation of its corresponding element (published before) to the SMS Child.

5. When the SMS Child received the Parent request, it sent the element *activationRequest* in order to request from MANTICORE the activation of the path passing through the routers of the jra23 slice.
6. MANTICORE which in turn processes the Child request, configured the *router4* input interface modifying the policer (increasing the BW) applied to the interface. In Fig. 9, the moment can be seen when MANTICORE is configuring the router by means of the *NetConf* protocol.

```

#: <session-id>33675</session-id>
#: </hello>
#: ]]>]]>
14:30:11,373 INFO engine.NetConfProtocolHandler [ServiceThread-64,send:151] <rpc>
<edit-config>
<target>
<candidate/>
</target>
</config>

#: <rpc>
<edit-config>
<target>
<candidate/>
</target>
#: <rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:junos="http://xml.juniper.net/
INFO engine.NetConfProtocolHandler [ServiceThread-64,send:156] <configuration>
<firewall>
<policer>
<name>limitBW</name>
<if-exceeding>
<bandwidth-limit>30000000</bandwidth-limit>
<burst-size-limit>500000</burst-size-limit>
</if-exceeding>
<then>
<discard/>
</then>
</policer>
</firewall>
</configuration>

#: <configuration>
<firewall>
<policer>
<name>limitBW</name>
<if-exceeding>
<bandwidth-limit>30000000</bandwidth-limit>
<burst-size-limit>500000</burst-size-limit>
</if-exceeding>
<then>
<discard/>
</then>
</policer>
</firewall>

```

Fig. 9 NetConf stream when MANTICORE is configuring the router

7. Finally, when the policer has been configured to increase the BW through the interface, the input bit rate in the target Host is increased causing the Video Quality on the streaming reception to be improved (Fig. 10).

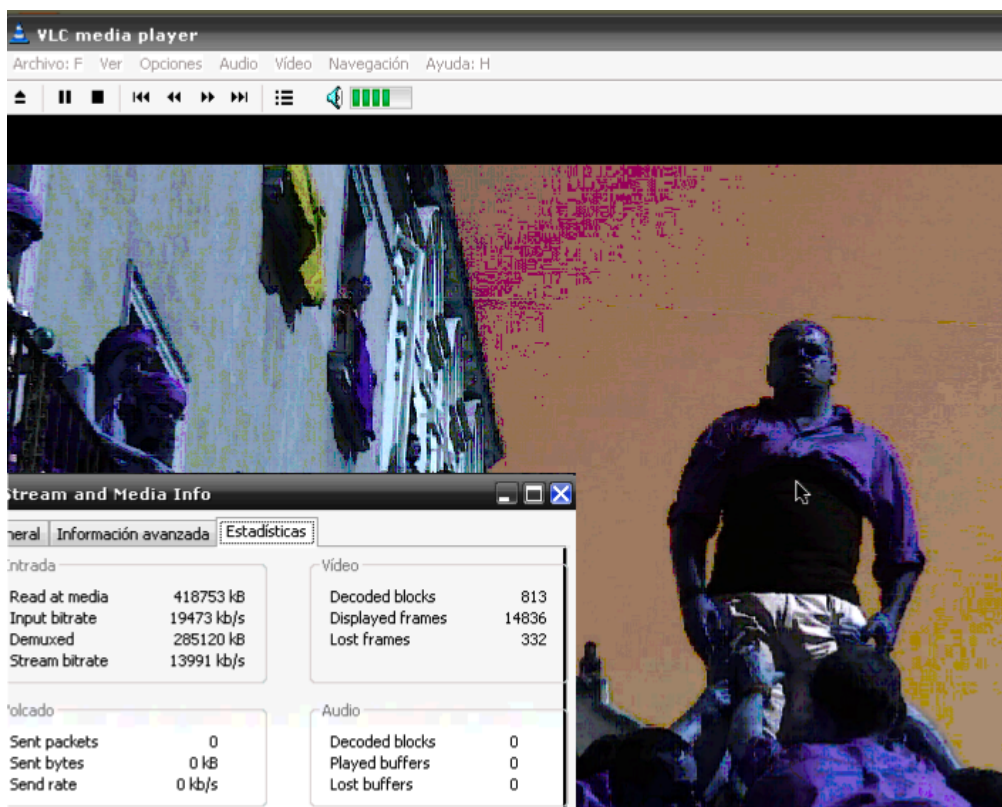


Fig. 10 After unfiltering the path (increasing BW limit) the video reception is enhanced

4 Conclusions and further work

As it was initially stated, the main objective of DJRA2.2 interoperability prototype is to provide the FEDERICA researchers (owning any FEDERICA slice) the mechanisms by which to interconnect FEDERICA slices with each other or with other similar e-Infrastructures (such as GENI). This is accomplished by the establishment of interconnections between slices and other domains.

Prior to this deliverable, an interoperability prototype which contained some of the main components of the IPsphere Framework, was developed and tested in certain scenarios that did not include the FEDERICA domain. In this final deliverable, the focus was on the validation of this prototype in order to establish paths over the FEDERICA infrastructure. To do so, a complex scenario composed of a FEDERICA slice interconnected through the FEDERICA substrate to the i2Cat Point of Presence was built.

After the creation of the complete scenario, several minor tests were performed over it and the major one has been explained in this deliverable concerning the steps followed for demonstrating the expected behaviour of the prototype when it is used over the FEDERICA facility. This final test covered the activation of a pre-established

but filtered path between the two sides of the scenario (from one domain to the other), by means of the SMS Child and MANTICORE interaction. Once this path has been activated a video streaming (that couldn't be transmitted very well before the activation of the service path) has been performed with positive results.

With this demonstration, we can conclude that the prototype, designed and developed following the bases of the IPsphere Forum specifications, has a real and positive application over experimental facilities like FEDERICA. It was also demonstrated that the IPsphere Framework is capable of interoperating with other Management Tools or Frameworks located in the layer between itself and the physical infrastructure.

In the future, this prototype may be completed with the implementation of additional modules of the IPsphere Framework (e.g., SMS Parent) and also with an enhanced SMS Child capable of interoperating with Management Systems other than MANTICORE. This may allow the interconnection of FEDERICA and other experimental infrastructures (managed by their own Management Planes) and follow the latest IPsphere Framework specifications however, this is beyond the scope of the current FEDERICA project.

References

- [Man] Eduard Grasa, Xavier Hesselbach, Sergi Figuerola, Victor Reijs, Dave Wilson, Jean Marc Uzé, Lars Fishcer, Tomás de Miguel, “The MANTICORE Project: Providing users with a Logical IP Network Service”, TERENA Networking Conference 2008
- [Ips] TM Forum, IPSphere Framework: <http://www.tmforum.org/ipsphere>
- [Jun] Juniper Networks SRC Series Session and Resource Control:
<http://www.juniper.net/us/en/products-services/software/ipc/src-series/>