

## **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 DSA2.3: SA2 Final Report**

## **Version 1.0**

<b>Dissemination Level</b>	Public
<b>Contractual Date of Delivery</b>	31 August 2010
<b>Actual Date of Delivery</b>	7 January 2011
<b>Editor</b>	M.Campanella - GARR
<b>Contributors</b>	M. Potts - Martel B. Rhoads- KTH

## **Reviewers**

## **Abstract**

This deliverable summarizes the experience gained in the SA2 activity of the FEDERICA project. The document focuses on the user support (Task TSA2.1) and reports the procedures and improvements since the description in deliverable DSA2.1 (released March 2010). The detailed description of “support tools”, “user portal” and “policies” are not repeated here. Considerations on achievements and lessons learned in the Support Activity SA2 are listed as a contribution for future improvements of user support.

## Executive summary

This deliverable summarizes the activity and experience learned in the SA2 activity of the FEDERICA project. The document focuses on user support (Task TSA2.1), which has been officially operational from November 2008 to the October 2010, and it is still supporting users in their experiments at the time of writing.

The SA2 activity has been split in the Description of Work, in four Tasks:

TSA2.1: User support during use

TSA2.2: Prioritization, policies AAI charging and accounting, gathering user feedback

TSA2.3: Tools for virtual slicing and network provisioning, e2e monitoring, fault management of virtual resources, AAI, interaction of virtual infrastructures

TSA2.4: User oriented tools for creation and control of virtual infrastructure.

In the initial period of the project, it became evident that Task TSA2.3 and TSA2.4 had to closely cooperate and soon united their effort for producing the toolbench. Part of the TSA2.4 effort was actually devoted to better support monitoring and producing the User Portal in collaboration with JRA1. This has been reported in the previous DSA2.1 deliverable. The TSA2.3 effort for the development of a tool to partially automate and simplify virtual resources creation and management and whole slice creation, proceeded in the last year with substantial support from JRA1. The manual and description of the last version of the tool is reported in DJRA1.3, released at the same time as this deliverable.

The TSA2.2 activity has summarized its findings in deliverable DSA2.2 and those results have not been extended in the last reporting period.

This document focuses therefore on Task TSA2.1 “User Support”. It describes the operational procedures, the various elements and the overall workflow to create and support a slice<sup>1</sup> in the FEDERICA infrastructure.

Slice creation and management is a complex operation, which starts from technical discussion with the users to fully specify the slice blueprints. The NOC, which is combination of SA1 and SA2 effort, then implements the slice, configures the user portal and updates the substrate configuration and monitoring and handles it over to the researcher. The NOC continues to support the users during the experiment, for example when technical issues arise or there is the need to reload a blocked resource with no user connectivity. The NOC gracefully terminates the slice when it is no longer needed.

Part of the configuration and management effort in the various phases of the slice lifetime has been automated, and part has remained manual. This has proven to be the best balance between the complexity of slice management, the available effort and the intrinsic difficulty of virtual resource mapping to satisfy user’s requirement and operational rules, and the number of requests.

The experience has indicated that more automation is needed to scale the facility to serve faster a higher number of users. Partial automation and an Application

---

<sup>1</sup> A slice is defined as a set of virtual resources that has been customised to fit the needs of a specific user, and that is running on top of the physical infrastructure, sharing it with other parallel slices.

Programming Interface are also needed to enable a basic level of federation, to implement a service that creates virtual testbeds without human intervention. Automation is particularly needed to ensure that the all information on slices and substrate is synchronized quickly between all the facility management and control components. This includes, but is not limited to all physical and virtual resources inventory, the topologies and configuration, monitoring and alarms.

## Table of Contents

<b>Executive summary .....</b>	<b>2</b>
<b>1 Introduction .....</b>	<b>5</b>
1.1 Document overview .....	5
<b>2 The FEDERICA infrastructure .....</b>	<b>6</b>
<b>3 The NOC .....</b>	<b>9</b>
3.1 TSA2.1, TSA1.1 and the United Network Operation Centre.....	9
3.2 United NOC objectives and Tasks.....	10
3.2.1 Physical Substrate .....	10
3.2.2 User Support .....	11
3.2.3 Slice Creation and Management .....	12
<b>4 Conclusions .....</b>	<b>14</b>
<b>Appendix A Slice creation .....</b>	<b>16</b>

## 1 Introduction

FEDERICA is a 34-month European Commission co-funded project in the 7th Framework Programme, GÉANT and e-Infrastructure unit. It started in January 2008 and it engineered, implemented and operated as a service an e-Infrastructure for researchers in current and future Internet. This infrastructure is agnostic in respect of the types of protocols, services and applications that may be trialled, whilst allowing disruptive experiments to be undertaken.

The infrastructure is based on computing elements, routers/switches and network circuits in which virtualization technologies are available. The infrastructure main service is the creation of “virtual infrastructures” named “slices”, which contain network topologies and virtual system to be handled to the users. The slices are under full user control and can support innovative and disruptive experiments.

The main objective of the Service Activity SA2 (Operational User support and toolbench development) is the support of the researchers in the use of the FEDERICA slices for the whole duration of the tests.

The user is granted access to the infrastructure after an initial discussion and eventual approval by the User Policy Board. The SA2 effort starts already during the initial user’s dialogue phase with the UPB task and continues until the end of the experiment and the return of the resource in the availability pool. The largest effort in the process lies in the creation of the slices and their maintenance during the experiment.

The SA2 activity had the task to support both the NOC (or operate on its behalf) and the user in the process, by specifying workflows and automating procedures, also by developing a set of tools.

In the initial period of the project, it became evident that task TSA2.3 and TSA2.4 on tool development had to closely cooperate and soon united their effort for producing the toolbench. Part of the TSA2.4 effort was actually devoted to better support monitoring and producing the User Portal in collaboration with JRA1. This has been reported in the previous DSA2.1 deliverable. The TSA2.3 effort for the development of a tool to partially automate and simplify virtual resources creation and management and whole slice creation, proceeded in the last year with substantial support from JRA1. The manual and description of the last version of the tool is reported in DJRA1.3, released at the same time of this deliverable.

In the following we therefore focus on TSA2.1 effort and on user support workflow. The experience learned and the improvements are also considered in the document as contribution for future developments.

### 1.1 Document overview

This document contains information regarding:

- The interaction between the infrastructure support and the user support, in particular how the two tasks have to closely cooperate and may eventually be managed by a single entity (the Network Operation Centre – NOC);

- The description of the steps to create a slice, using examples of provisioned slices;
- The issues and possibility for improvement in the effort.

It is structured as follows:

Section 1, entitled “Introduction”, is a short introduction describing the background, content and structure of this deliverable.

Section 2, entitled “The FEDERICA infrastructure”, briefly describes the substrate to provide the needed information and terminology to be used in the following sections.

Section 3, entitled “The NOC”, describes the role and evolution of the infrastructure and user support. It also describes the operational procedures that the NOC put in place to create and control a slice. The section has brief reminders of tools, monitoring system and user portal to better describe the various steps.

Finally, Section 4, entitled “Conclusions”, closes the main part of the document.

A step-by-step procedure for creating a slice using VLAN links without guaranteed capacity is reported in Appendix A.

## 2 The FEDERICA infrastructure

The FEDERICA infrastructure spans 12 countries in Europe and it is composed by 14 Points of Presence, each hosted by a different project partner. Figure 1 below shows its topology.

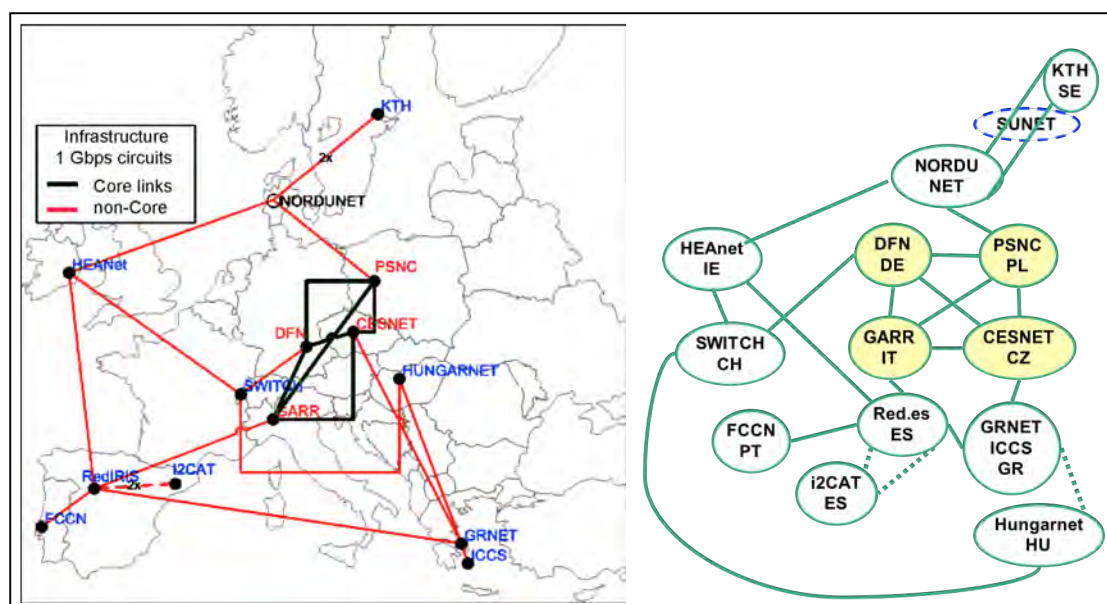
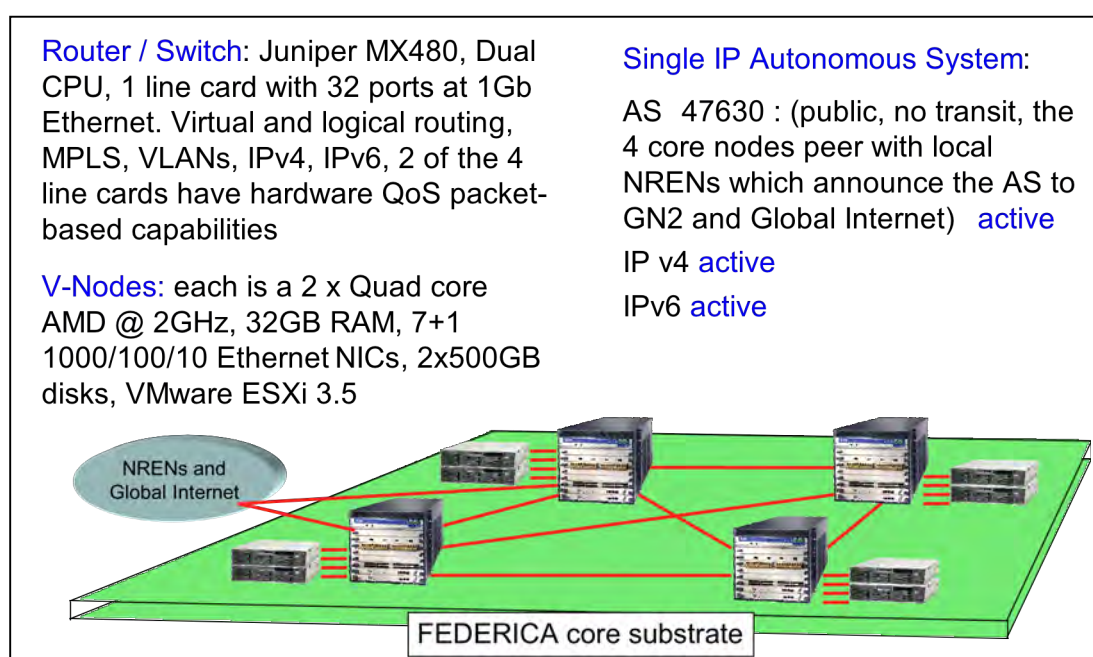


Figure 1: FEDERICA topology; physical, left and logical, right.

The physical substrate comprises 4 core PoPs located in CESNET (Czech Republic), DFN (Germany), GARR (Italy), and PSNC (Poland), and 10 non-core PoPs located in FCCN (Portugal), GRNET (Greece), HEAnet (Ireland), I2CAT (Spain), ICCS (Greece), KTH (Sweden), NIIF (Hungarnet), PSNC (Poland), RedIRIS (Spain), and SWITCH (Switzerland).

Core PoPs are connected between them in a full mesh topology by 1Gbps dedicated GÉANT+ network circuits, terminated directly on FEDERICA equipment, Juniper MX480 routers. An exception to this rule is the FEDERICA DFN PoP, while the GÉANT PoP is located in Frankfurt, the FEDERICA infrastructure is located in Erlangen and four dedicated  $\lambda$  channels are set up by DFN, extending the GÉANT circuits to FEDERICA installation site.

Figure 2 summarizes the information on the equipment in the core and its capabilities, as well as its configuration as part of to Internet. The FEDERICA infrastructure is a public Autonomous System in both IP version 4 and Version 6 which has 4 full BGP peering in the four core PoPs with the local NREN and through them to GÉANT and Internet. The physical infrastructure and all the slices are by default protected by firewall configurations.



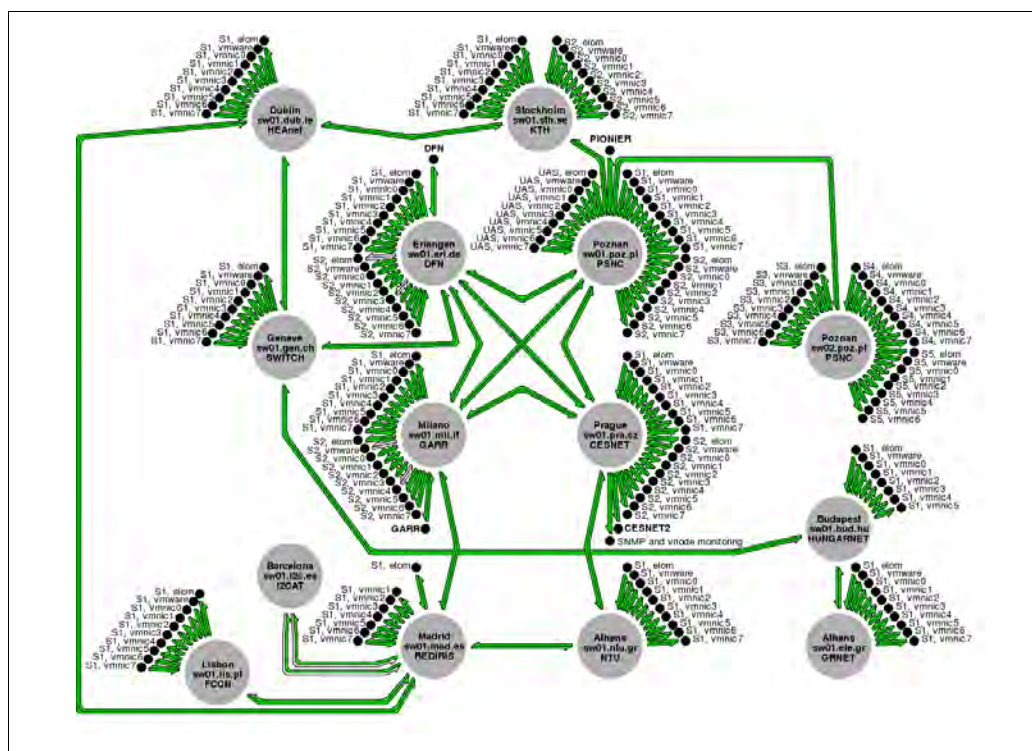
**Figure 2: The core part of the substrate, hardware and software details**

The other key component of the infrastructure is the monitoring tool, which can track the state and utilisation of all resources in the infrastructure and offer statistics and alarms to the NOC and the user. The G3 system from CESNET was used, collecting physical and virtual network state and traffic information on Juniper MX and EX interfaces via SNMP and IGMP protocols. The monitoring system was enhanced to gather information from the virtualisation servers and virtual resources in slices using Nagios and the VMware API available in ESXi. A single flow monitoring probe is implemented in the CESNET PoP to provide per-packet information on selected flows (the slice configuration has to include resource in the CESNET PoP).

Figure 3 shows, as taken from the web graphical interface, the snapshot of the status all circuits, routing/switching equipment and V-Nodes. The graphical interface allows to explode any single element in detailed statistics. The number of elementary entities monitored, as an example a CPU, a network interface, System Memory (either physical or virtual) is more than 53'000. For each entity data is shown in historical format with different time averages. The data is available on request also for the

resources in each slice in the same format. Access to monitoring information is restricted to the owner of the resources, while the NOC has access to all the substrate information and to the slice elements on request. Due to the very large number of entities, which are available for monitoring, the system has been tailored in two aspects using input by the NOC and the users:

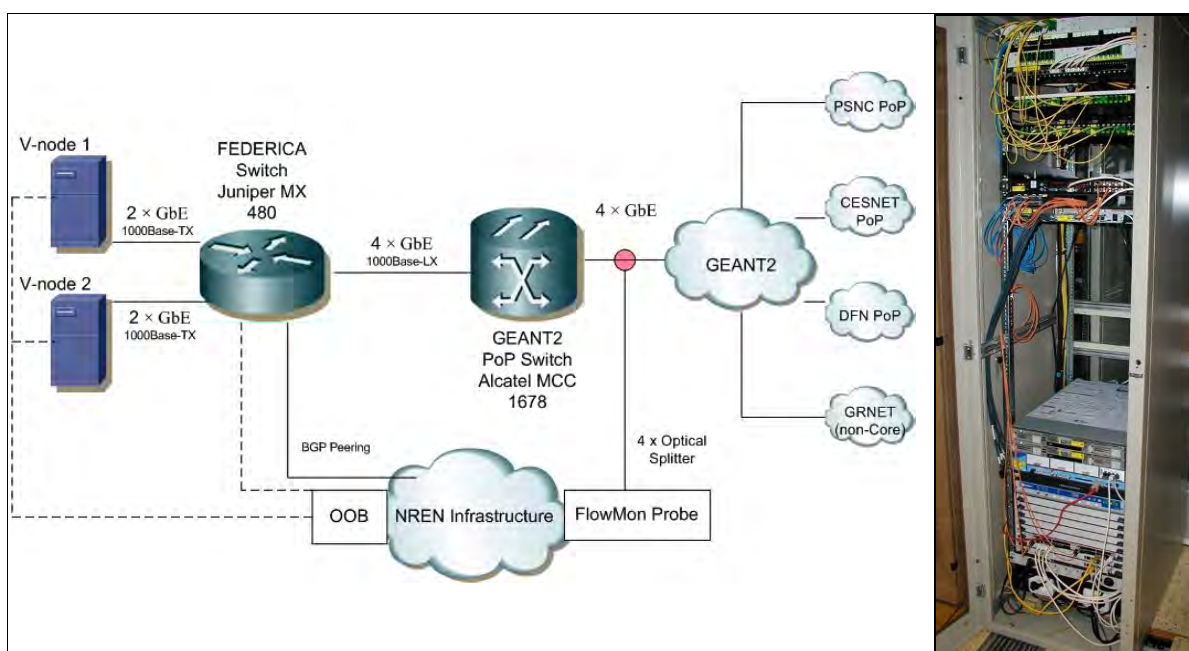
1. The elements to be monitored. The choice has been made to ensure that all critical entities in the substrate are constantly monitored, to ensure that failures are promptly signalled. The resources in each slice are monitored according to user’s requirements. A default set of resource in each slice is monitored for state.
2. The monitoring interval for each resource. Due to the large number of entities to monitor, it is mandatory to avoid peaks of SNMP requests, as they will cause loads in the CPUs of the equipment to be monitored. Usually the intervals are not in phase with each other and the system works in ‘relaxed’ mode, with longer delays between measurements.



**Figure 3: FEDERICA infrastructure top-level view by the G3 monitoring system**

In addition to the G3 monitoring system, the collaboration with the GN3 project allowed DFN to use the HADES monitoring system of GN3 to perform on-demand active measurements of packets Round Trip Time and circuit performance on the core Ethernet circuits of the FEDERICA infrastructure. This allows the NOC to access information on the same circuits gathered from two different type of equipment.

Figure 4 depicts, as an example, the CESNET PoP with connection details with the GN3 Alcatel equipment. The HADES monitoring system (based on PerfSONAR) is hosted in the GN3 PoPs



**Figure 4: CESNET PoP Rack and GN3 connection details**

For more details the deliverable DSA1.3 “The FEDERICA Final Infrastructure” provides an in-depth description of the facility.

### 3 The NOC

This section examines the operation of the user support. The NOC objectives, tasks and functionalities are examined, to conclude that a united SA1 and SA2 NOC is more efficient as indeed has been deployed. The NOC operations are then described, with particular details on slice creation.

#### 3.1 TSA2.1, TSA1.1 and the United Network Operation Centre

The main objectives of the Service Activity SA2 of user support and slice creation, imply operations on the substrate and on the whole infrastructure, which largely overlaps with responsibilities from SA1. The Task TSA1.1 has the objectives of infrastructure support, virtual infrastructure factory and monitoring. The SA1 Activity is composed by all partners hosting a PoP and it is led by PSNC. The TSA2.1 partners, led by KTH are more focused on user support, but participated since the beginning to the infrastructure planning and construction. The slice creation procedure and the user support require an extensive knowledge of the substrate and virtual resources and the full control of all resources.

For these reasons, and to optimize the effort since the beginning, the core Network Operation Centre of FEDERICA has been composed by partners in both activities, which worked as a single unit.

The SA2 NOC was initially defined to be composed by KTH, FCCN, NIIF, and NORDUNET, with KTH being the Task leader. After the initial construction phase, in which KTH has been a very active participant, the NOC became a united NOC composed by PSNC, KTH, and FCCN with effort from Red.es contributing in the authorization and authentication.

The effort has been coherent with a constant increase of the user support effort with time, in parallel with minimal substrate maintenance.

## **3.2 United NOC objectives and Tasks**

The NOC tasks are physical substrate maintenance and evolution and user support, which includes slice management.

### **3.2.1 Physical Substrate**

The physical substrate support leverages the experience of the partners in networking and systems and follows a set of standard procedures that the NRENs have in place in their network.

The basic FEDERICA substrate is configured a single network domain based on a mesh of 1 Gbps circuits. Its IP addressing and routing schema are standard in IPv4 and IPv6; it is a single non-transit AS domain using BGP. The GÉANT NOC is responsible for the optical and SDH transport layers from which the circuits are created. Routing and switching equipment are configured for providing VLANs, MPLS LSPs, standard routing for the management networks connecting all equipments.

The V-Nodes are very stable and reliable from the hardware point of view and can run unattended for long periods of time. There has been only one issue with the virtualization software (see below).

Standard procedures (not detailed here) for scheduled maintenance for configuration changes, Software and/or Hardware upgrades, fault identification and solutions, and monitoring statistics and alarms (based on the CESNET G3 system) have been deployed.

The substrate maintenance has been rather straightforward.

The issues encountered are considered part of normal operation and very difficult to prevent. There have been:

- One or two fibre cuts per year in the GÉANT optical footprint
- Ethernet circuit malfunctions, needing transport equipment reset
- In some V-Nodes, the virtualization software degraded performance with time and needed to be rebooted after a couple of months.

### 3.2.2 User Support

The user's support is composed by two main areas:

1. Users' consultation and advice. This happens before the user receives a slice and continues for the whole duration of the experiment. The initial phase is led by the User Policy Board (UPB) and the NOC as technical consultant. The NOC has also the knowledge of the network status and can provide tentative timescales.
2. Slice creation and management. This task is the more complex and has also to be performed in as short time as possible (or according to user's request). The task needs configuration in the substrate and may extend up to configuration in the virtual system images.

The initial user's consultation and advice is performed based on the documents submitted to the UPB and proceeds usually via the NOC mailing list (`federica-noc@fp7-federica.eu`) and the user's Portal (described in the previous SA2 deliverable), which is then the main information access tool. Its role starts with supplying preliminary information to the potential users and continues with providing guidance on the account and slice creation and deletion procedures and implementing the primary interface between the end-users and the NOC throughout the life of the slice (e.g. supporting change and failure management).

The portal functionalities and usage are shown diagrammatically in Figure 5.

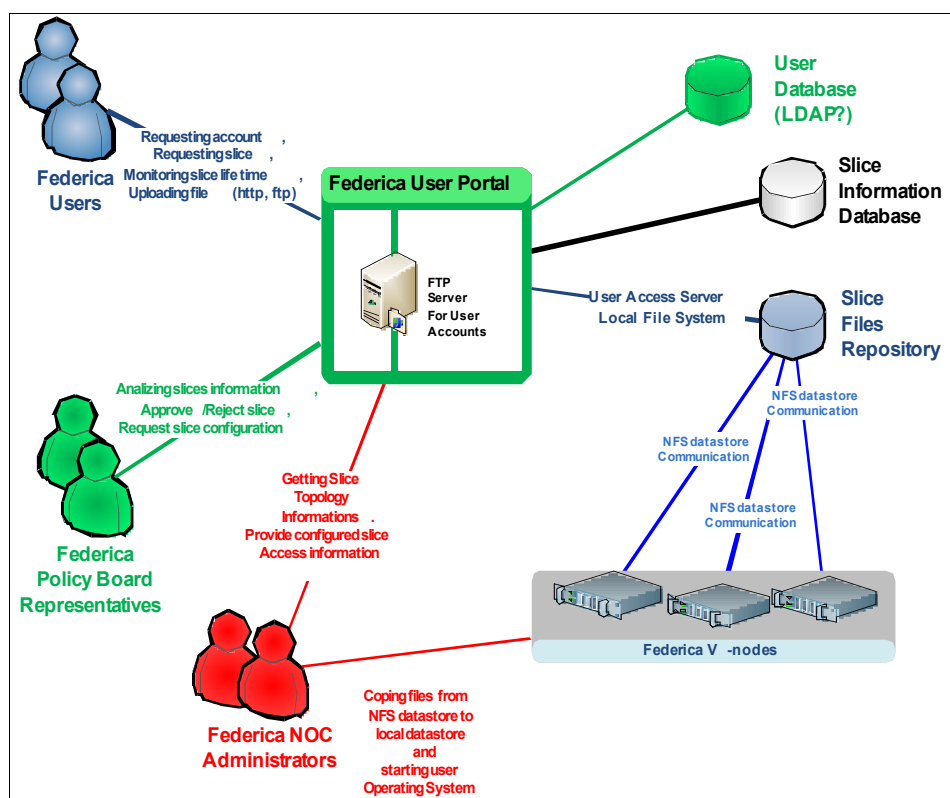


Figure 5 : The users' portal architecture

### 3.2.3 Slice Creation and Management

The slice creation process relies on the information provided by the users and on the current state of the substrate. The goal is not just to create what the user has requested, but also to optimize the use of the substrate, ensure a level of reproducibility to computing resources, even if the user did not made an explicit request and avoid any influence on other slides. In case reproducibility guarantees are requested for some or all the resources, appropriate configurations will be applied in the router/switches and in the V-Nodes. The effectiveness and the procedures to configure a resource for reproducibility is the subject of a separate internal note.

According to user’s requirement the networking technology used to create the topology may differ from VLANS to MPLS, with routing or without. There isn’t therefore a predefined procedure, and each slice has to be engineered separately. As default, VLANs and routing-instance in Juniper are used.

A step-by-step procedure for creating a slice using VLAN links without guaranteed capacity is reported in Appendix A. Each slice will have a VLAN just for management purpose, in addition to the user’s topology.

Figures 6 and 7 show for the VMSR slice respectively the user topology (as requested) and the NOC topology information, which includes the physical placement of the resources in the substrate.

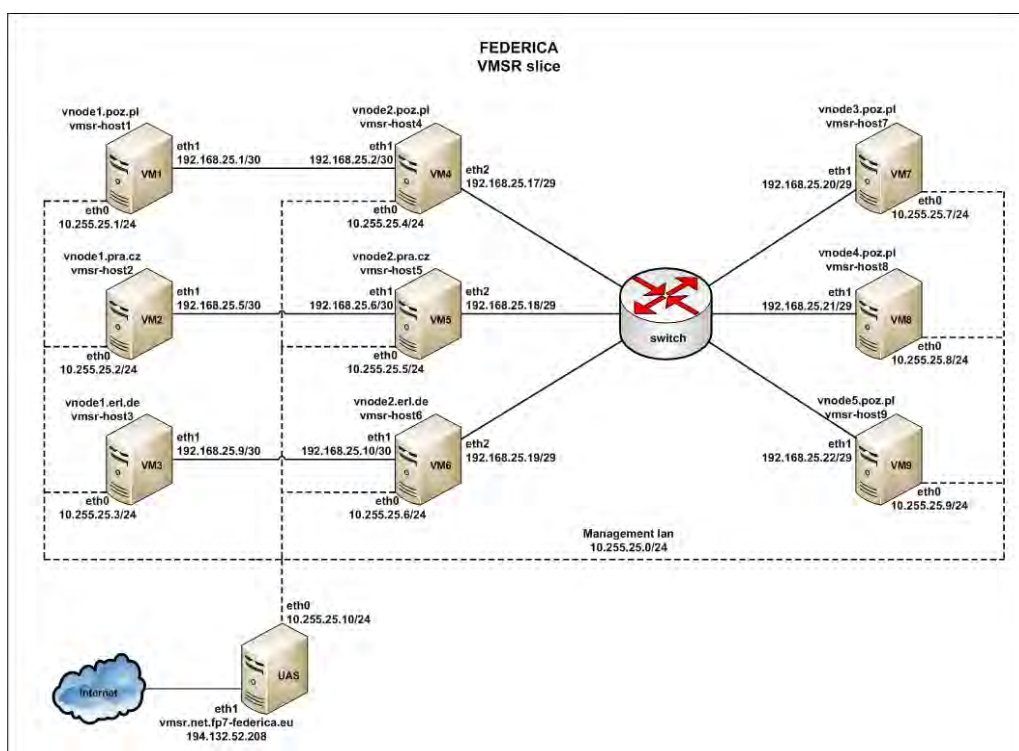


Figure 6: VMSR - User's view

As it can be noticed a single link in the user topology, may actually imply many physical circuits and the transit by more than one switch/router.

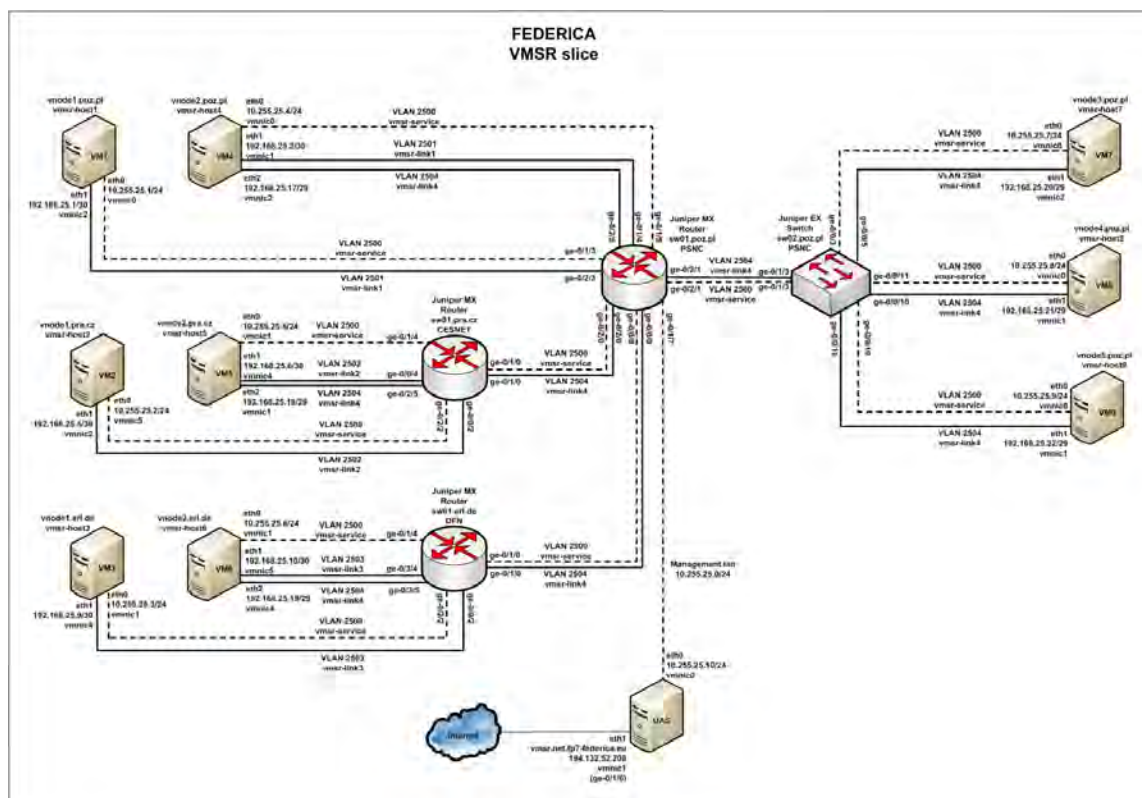


Figure 7: VMSR - NOC full view

At the time of writing the slice implementation is using a manual database based on textual information in the Wiki. This approach has been effective up to present because of the relatively low number of slices and change rate in the information. The use of the tool produced by the project or any other management application will introduce a database function, which needs to be engineered for the FEDERICA infrastructure, where virtualization introduces complex relations between entities.

These relations are at present realized also by an accurate naming scheme. The database needs also to be shared between monitoring, slice creation function, user portal and the API to automate the FEDERICA service. The current experience is sufficient to engineer and implement the database functionality.

All slice configurations have a backup on a storage server in case of failures. The backup of user configuration, as well as checkpointing of the virtual hosts are not yet implemented and are a responsibility of the user.

## 4 Conclusions

The SA2 effort and user support provided has been able to implement and operate the FEDERICA facility to serve about 15 different users' groups, each with different requirements and specification for the slice. The virtualization capabilities of the facility have performed according to plans.

The experience suggests some improvements and guidelines for the evolution of the NOC effort:

- A distributed NOC, made by partners in various countries, works well, but efficiency can be gained by collapsing the effort in one or two sites.
- The time to devote to support the user is larger than expected. The service is targeted at researchers which trial complex technologies and every slice poses usually specific and innovative challenges. The NOC and the UBP should closely collaborate and receive adequate effort.
- The tool has demonstrated to be feasible to automate part of the provisioning system. An integration or communication between the databases of the used application (monitoring, provisioning, security) is needed to allow scaling in number of slices and in the complexity of each.
- Mapping the virtual resource to the physical substrate is a complex operation. It is still done manually, because of the constraints (avoid overbooking, ensure reproducibility where needed) and variability of the substrate allocation. In the case that reproducibility is not requested, the rules can be relaxed for network links and for virtual hosts, although it is preferred to maintain reproducibility where possible. The procedure and policy needs to be improved according to users' experience and requirements.
- The standard network management procedures of the NRENs work well in the FEDERICA environment. The NOC needs to be capable of handling not just networking equipment, but also computing elements.
- To ensure scalability and increase in efficiency, the toolbench needs to be further improved to provide a greater amount of automation. It is envisaged that the more than one application will need to be used at the same time (e.g. one for computing equipment and a second for networking equipment), which will be maintained in sync by the management main application.



## **Appendix A**    **Slice Creation**

### **A.1**    **Scope of the Appendix**

This appendix outlines the procedures for manual slice creation and is based on the experience in creating test slices in the FEDERICA substrate. This appendix focuses mainly on the network creation and not the creation of the virtual hosts, i.e. the procedures for creating the network for a slice.

The creation of virtual hosts follows the procedures for the specific virtualization application (VMware at present). The system host images will be loaded by the NOC or by the users.

The use of the tool developed automates some, but not all, the steps below. Until the tool is the only application in use to manage the substrate and the slices, the database of the tool and the manual database need to be kept coherent manually.

### **A.2**    **Get a slice description from the customer**

Document D5 – “Resource description” is the starting point for the user’s network topology. In case it’s not available, in any case a similar diagram should be produced with the help of the user. It will streamline the slice creation and serve as a base to evolve from.

To get the initial data, the FEDERICA User Portal provides a structured way for customers to provide data needed for a slice. Example of data that can be entered into the Portal:

- \* Date when the slice should be created
- \* Date when the slice should be removed
- \* Slice description (scope, responsible, administrative data)
- \* Slice topology description - drawing, and type and characteristic of each resource
- \* Files (system images), which should be uploaded to V-Nodes (if requested).

### **A.3**    **Assign a name for the slice**

Since almost all configuration data is name based, the name of the slice must be unique and should reflect the customer/project using it, if possible. In the case of test slices we have e.g. “meter” for the measurement slice, “Cesnet” for the slice that CESNET needed to generate different loads on the network. Both are examples of slice’s name that reflects either a project name or a customer name.

### **A.4**    **Create a RT-case for the slice**

The RT-case in the trouble ticket system will serve as the tracker for the slice during implementation and operation. It is also possible to link cases that affect the slice or requests made as ‘troubles’ However RT is not good to have comprehensive information in a structured manner, therefore is not enough and the complete information on the slice needs to be elsewhere.

Trouble ticket system is located at <http://rt.net.fp7-federica.eu/rt/>

### **A.5 Create a Wiki page for the slice**

The Wiki page serves as a structured informational page with details about the slice for the NOC. This includes information such as contact information, network diagrams and slice layout, i.e. the information that is more or less static for a slice lifetime. Both RT and Wiki should have a link to each other to make it easy to cross reference information.

### **A.6 Assign VLAN range**

To assign unique VLAN ranges check current allocation in the Wiki page:

<https://wiki.fp7-federica.eu/bin/view/SA2/VLANAssignment> .

### **A.7 Assign Service LAN IP range**

The service LAN is the LAN that initially connects all resources to allow the user to reach them. As with VLAN's it is preferable to have unique address ranges for the service LANs so that it is easier to isolate what traffic belongs where, even if the VLANs are in the private range. The address ranges are maintained in a Wiki page:

<https://wiki.fp7-federica.eu/bin/view/SA2/ServiceLanIPAssignment>

The issue to consider when assigning a unique IP-range for the service LAN is that it might conflict with user's specification for its own ranges. It has to be agreed with the user and the information be passed when handling the slice.

### **A.8 Assign slice IP range**

Analogous to the service LAN there is a need to assign an IP range for the slice itself. This might clash with a requirement from the customer and hence might overlap with other assigned ranges. The default for assignment, when any address can be used is to use 192.168.xx.0/24 where xx corresponds to the VLAN range xxyy.

### **A.9 Assign substrate links for to use by the slice**

As a function of the requirement of the presence of a logical router or a FEDERICA managed router in the slice, the layout of links can be done rather freely. The link layout is tightly coupled with the virtual host layout.

### **A.10 Assign virtual hosts to V-Nodes**

Once the basic information needed for the network is available, it is the turn to lay the slice out on the substrate. The process is not trivial, as it needs to optimize the use of the substrate, whilst complying with users' requirements. Considerations have to be done for circuit usage and V-Node usage.

It is desirable to try to keep core link usage limited in order to ensure that the slice behaves as predictable as possible, even if reproducibility has not been explicitly requested. This process has to be done on a slice-by-slice case. If a slice requests logical routers it constraints the virtual node layout, since the logical router must be in one of the core nodes. Overbooking of network resources has also to be avoided.

### A.11 Create the links/switches/routers

Each link that will pass through a router at least once, and in order to create a link that is transparent, implies the creation of some type of routing instance or logical router on the Juniper equipment. The creation of the links and of instances in the switches/routers are coupled together and need to be configured at the same time coherently.

The simplest form of device is the switch, a routing-instance of type bridge, to bridge together interfaces needed to make a transparent link. This is the basic building block for slice links, and will be a part of most, if not all links.

On a Juniper router the configuration might look something like the following:

```

routing-instances {
  meter {
    instance-type virtual-switch;
    bridge-domains {
      link3 {
        domain-type bridge;
        interface ge-0/0/0.1103;
        interface ge-0/3/3.1103;
      }
    }
  }
}

```

where the name of the routing-instance is the assigned slice name and bridge-domains contains the names of the configured links. The configured sub-interfaces have the same number as the VLAN-id they have for easy identification.

A router can be configured in two ways, either as a routing-instance or a logical-system. A logical-system is needed if the customer needs to be able to login into the router to do configurations. A routing-instance uses fewer resources on the main router and should be preferred when the user just needs a router, but does not request to be able to configure it personally.

An example of a logical-system is the following:

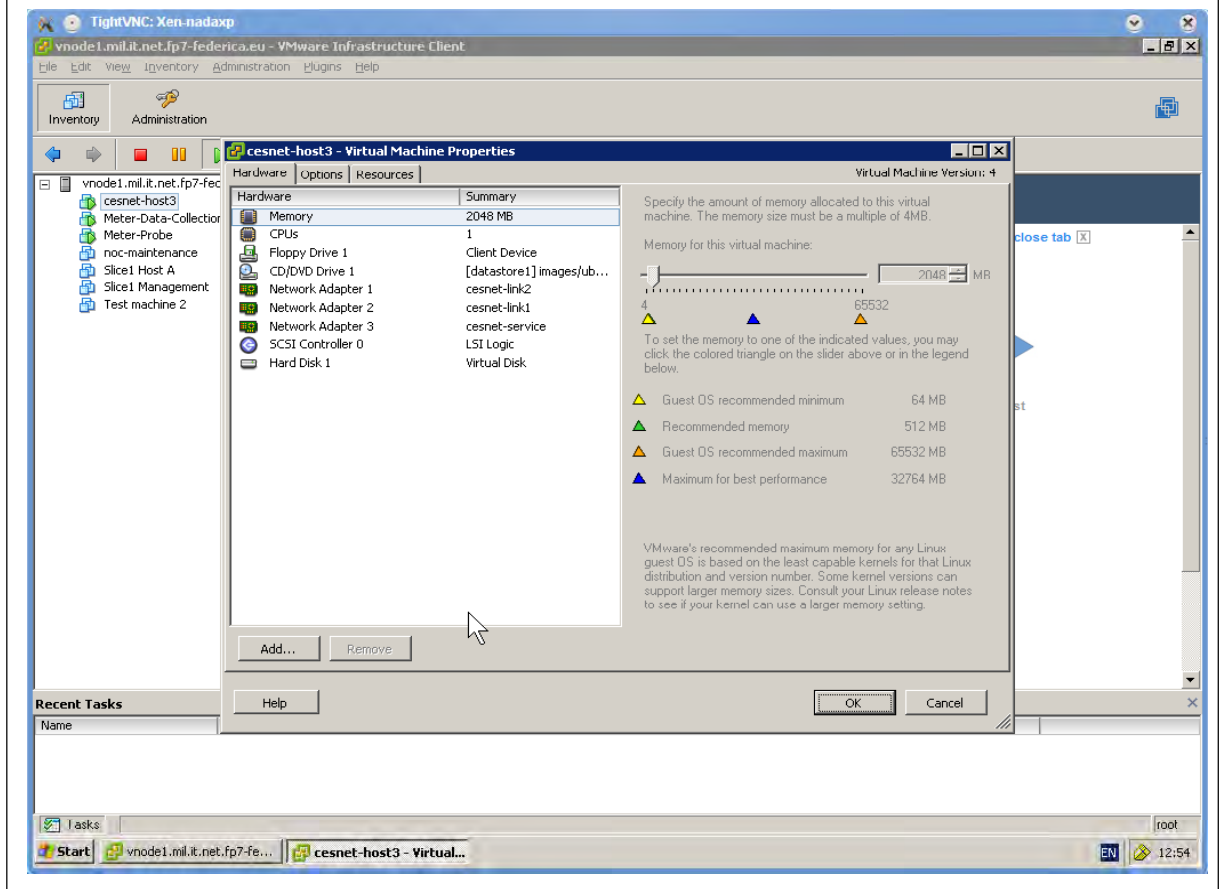
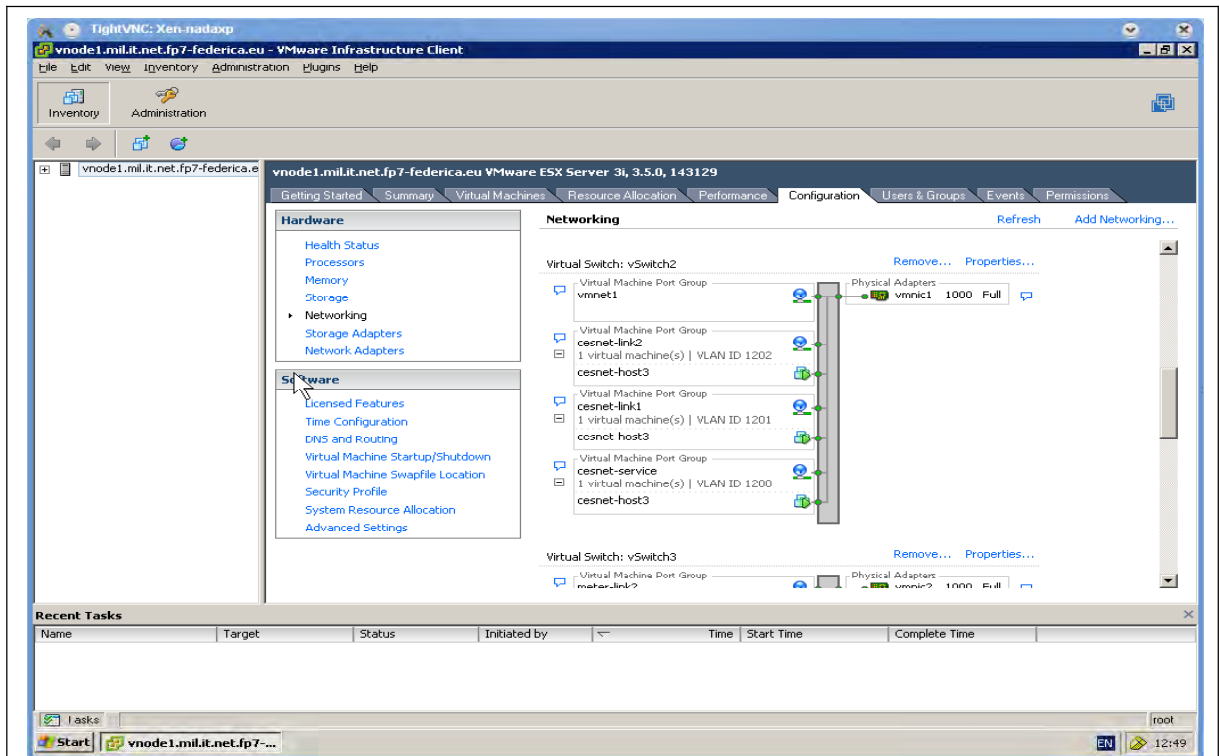
```

logical-systems {
  meter-router1 {
    interfaces {
      ge-0/2/0 {
        unit 1103 {
          description meter-link3-probe2;
          vlan-id 1103;
          family inet {
            address 192.168.2.1/24;
          }
        }
      }
    }
  }
  ge-0/2/3 {
    unit 1101 {

```



The example screen dump below shows the configuration for the “Cesnet” slice in the V-Node1 in Milan, Italy. By tying the interface to a virtual host the links will be complete. The second screen dump shows the configuration for a virtual node and what network interfaces are defined and where they are connected. By using a consistent naming scheme it is clear where the interfaces are connected. In this case, eth0 is connected to Cesnet link2 and eth1 to link1 and eth2 to the slice management network.

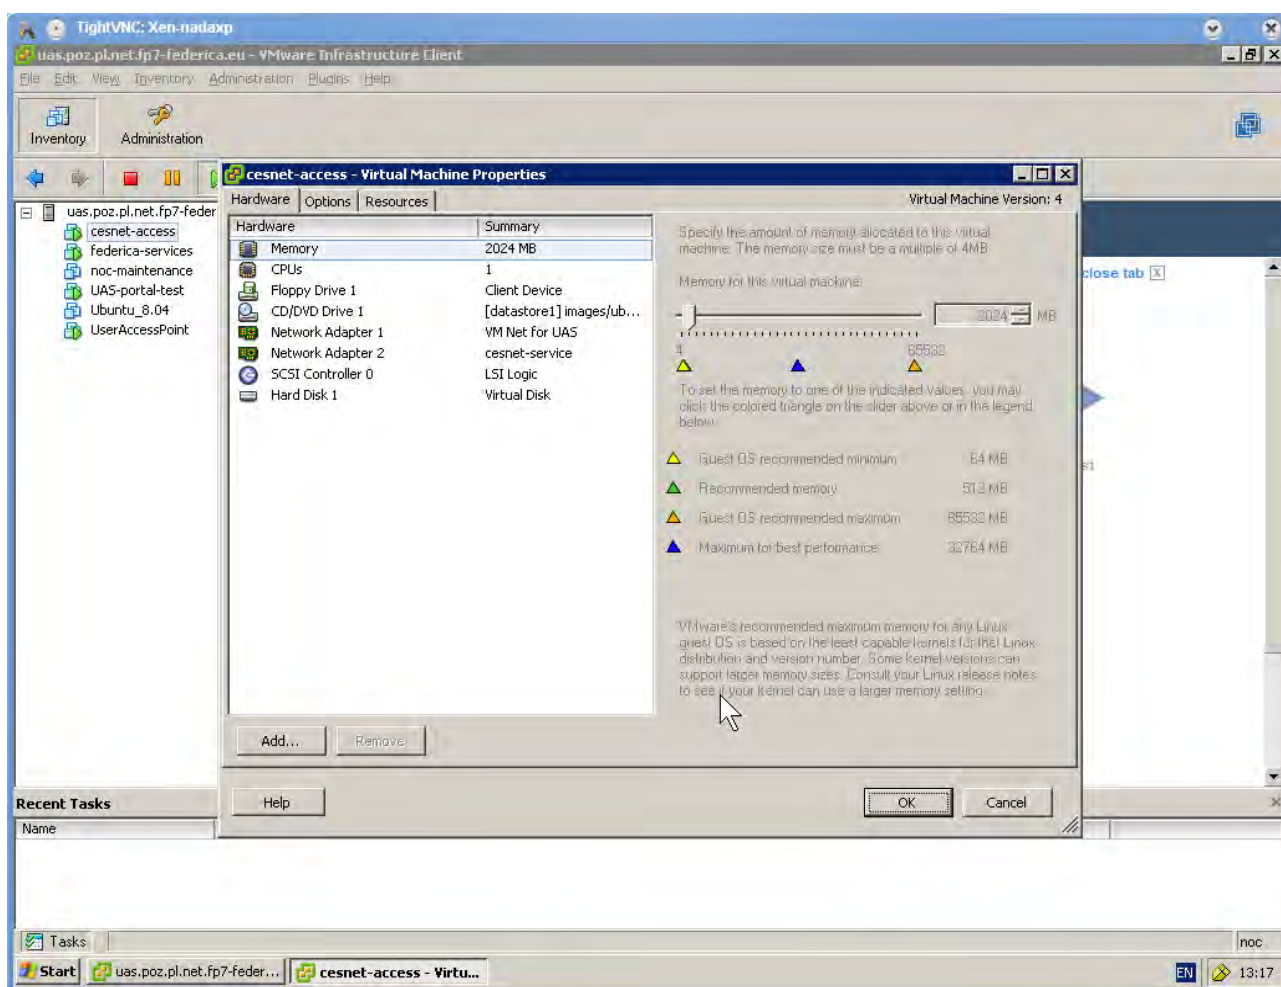


### A.13 Create the slice management network

In the similar way as creating links for a slice, the service LAN for the slice needs to be created. The difference is that the service LAN is just one broadcast domain, and thus only consists of switched routing-instances and all the links share the same VLAN tag, the first VLAN in the assigned range for the slice. It originates from the User Access Server in FEDERICA, `uas.poz.pl.net.fp7-federica.eu`.

The virtual host for the customer has 2 interfaces, one with a public address, and one interface towards the slice management network for the slice.

Since this is a single broadcast domain, it must be built as a star in order to avoid loops. The screen dump below shows the network configuration for the access server for slice “Cesnet”.



The slice creation and management is a complex operation, which starts from technical discussion with the users to fully specify the slice blueprints. The NOC, which is combination of SA1 and SA2 effort, then implements the slice, configures the user portal and updates the substrate configuration and monitoring and handles it to the researcher. The NOC continues supports the users during the experiment, as an example when technical issues arise or there is the need to reload a blocked resource with no user connectivity. The NOC gracefully terminates the slice when it is no longer needed.