



GÉANT Testbed Service (GTS) User and Resource Guide



Version 2.0

June 2015

Document Code: <GN4-P1-YY-nnn-nnn>

Authors:

© DANTE on behalf of the GN3plus project.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 605243 (GN3plus).

Contents

1	Introduction	4
2	Testbed topology and resources	5
3	How to get started	6
3.1	Access to GTS	6
3.2	Login to the GUI	6
4	Working in the testbed environment	9
4.1	Defining a testbed with DSL	10
4.2	Reserving testbed resources	12
4.3	Activating testbed resources	14
4.4	Setting reservation windows	16
4.5	Deactivating testbed resources	17
4.6	Releasing testbed resources	18
4.7	Querying details	19
4.8	Internet Access Gateway (IAGW)	22
4.8.1	Using IAGW with Windows	22
4.8.2	Using IAGW with Debian / Ubuntu	24
4.9	Example: Working with GTS OpenFlow switches (OFXs)	28
5	Help and Support	36
6	Introduction to Domain Specific Language	37
6.1	DSL in GTS	37
6.2	Quick Start	37
	Appendix I: Resource Guide	42
I.	Composite types	43
II.	Host	44
III.	Link	46
IV.	OFX	47
V.	External Domain	50
VI.	BNF Grammar	52
VII.	Using the Command Line Interface (CLI)	53
VIII.	Examples: One host	54
IX.	Example: Two hosts linked together	55

X.	Example: Triangle between three locations	56
XI.	Example: Three triangles built using DSL code iterations	58
XII.	Example: Two OpenFlow switches, one controller and three hosts	59

1 Introduction

Welcome to the new GÉANT Testbeds Service (GTS)! This new production service gives you the opportunity to setup your own customized experimental platform in order to be able to test novel concepts in networking and telecommunications.

This new service is made available in several releases and will grow further over time. This user guide describes features available in GTS Version 1.1.

The service allows you to set up an isolated customized packet-based testbed environment that is completely isolated from other testbed domains and enables you to simulate your particular experiment over resources in a real network without having to worry about impacting other testbeds or production services. After a short registration process a first time user has access to a pool of virtualized resources that can be programmed and reserved using a Domain Specific Language (DSL) description or selected via a graphical user interface. This reservation of resources also allows a user to specify start and end times during when these resources should be activated and made available for experimentation. New features available in this version include an Internet Access Gateway (IAGW) for experimenters so that they have the possibility to load their own software and also an external interface that can be selected as a resource whenever a testbed needs to be linked to external nodes.

The GTS facility is intended for use by researchers who are part of the GÉANT Community and its associated OpenCall projects [FAR-2014, SZE-2014]. Other researchers will be provided access on a space available basis.

2 Testbed topology and resources

Currently, GTS allows you to select resources from the following geographical locations: Amsterdam, Prague, Ljubljana and Bratislava¹ (see Figure 1). All these locations are sitting on the GÉANT backbone network point of presences. Later on, it is expected that new resource locations will be available outside of GÉANT, in the NRENs or other peer networks' domains.

Available resources in Version 1.1 of GTS include:

- Hosts (VMs with data ports; implemented using OpenStack)
- Virtual circuits (Ethernet pipe with data ports; implemented using Network Service Interface (NSI))
- OFXs (OpenFlow switch instances (OpenFlow Specification 1.3) with data ports)
- External domain.

The detailed description of the resources, their parameters and default attributes can be found in the Appendix I.

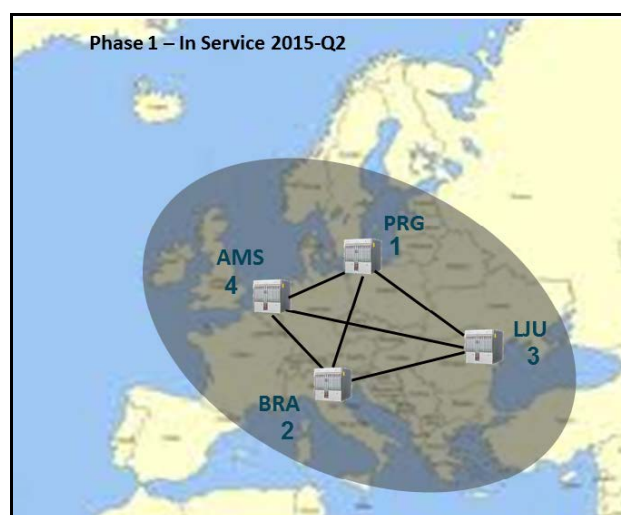


Figure 1: GTS physical infrastructure topology.

¹ Currently three more GÉANT backbone locations are being added in London, Milan and Hamburg.

3 How to get started

This section describes the steps you need to complete until your testbed setup is ready for use.

3.1 Access to GTS

The GTS front-end web Graphical User Interface (GUI) is available at: <http://gts.geant.net>. It is recommended to currently use the Google Chrome browser [GOO-2014] for the GTS GUI.

NOTE: Whenever you get the message “The service is currently undergoing maintenance, may not be fully functional. Check back later!” it practically means that we are upgrading and verifying the software/hardware functions of the production environment. The system may not be fully functional while the notice is displayed.

The GUI is the primary user interface of GTS but there are other means of accessing the testbed services. GTS also makes available a Command Line Interface (CLI) for more experienced users. Both the GUI and the CLI are built on a restful API in front of the GTS core software. In later versions of the software, the API will also be open for application designers and system integrators.

3.2 Login to the GUI

First time users have to register and define VPN credentials (Figure 2); users that are already registered can login with their user names and passwords as shown in Figure 3. During the initial registration process you have to define your username and provide your real name, e-mail address, password, password confirmation and the name of your project. You also have to define a VPN username and password that you will later use to VPN into your project. Once you are logged in you have access to the VPN information by clicking on the blue information icon next to your project's name (Figure 4). Logout of your project using the dropdown menu below your project name (Figure 4). This dropdown menu also gives you access to your user profile and allows updates if needed (Figure 6).

In the current implementation of Version 1.1, new user registrations always end up in new projects. If you wish to add multiple users to your single (already existing) project, please contact the GTS service managers.

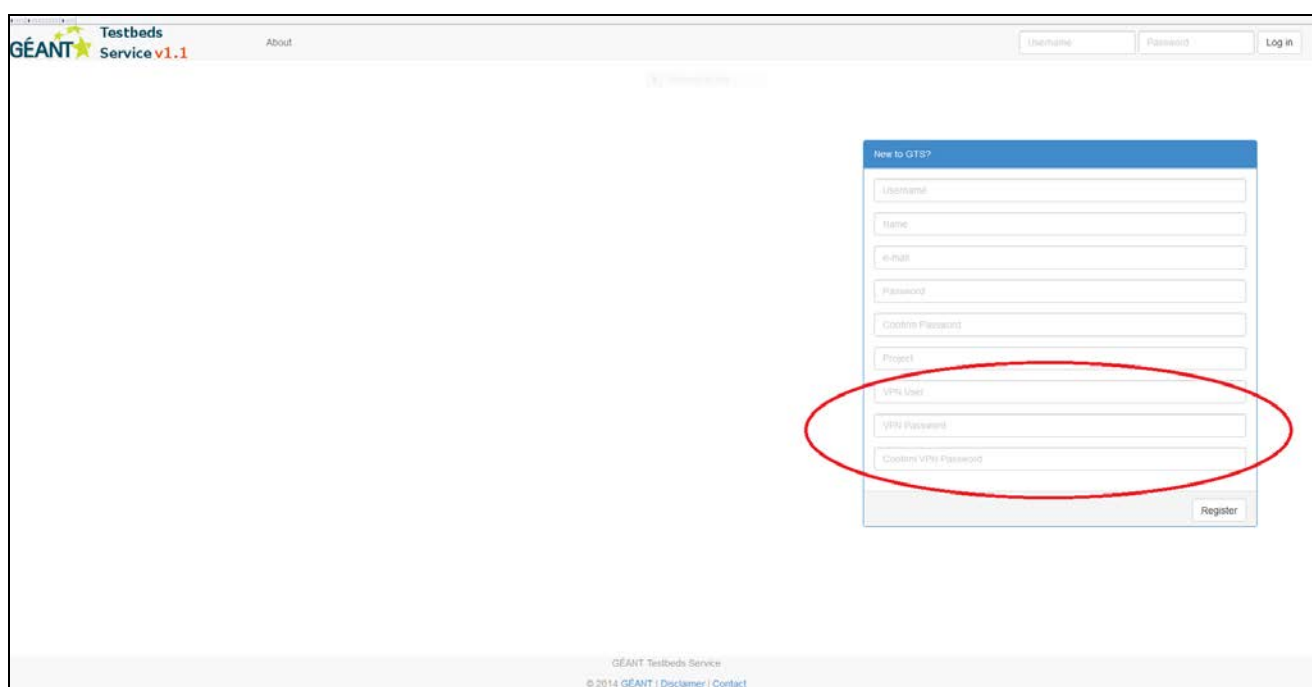


Figure 2: First-time registration

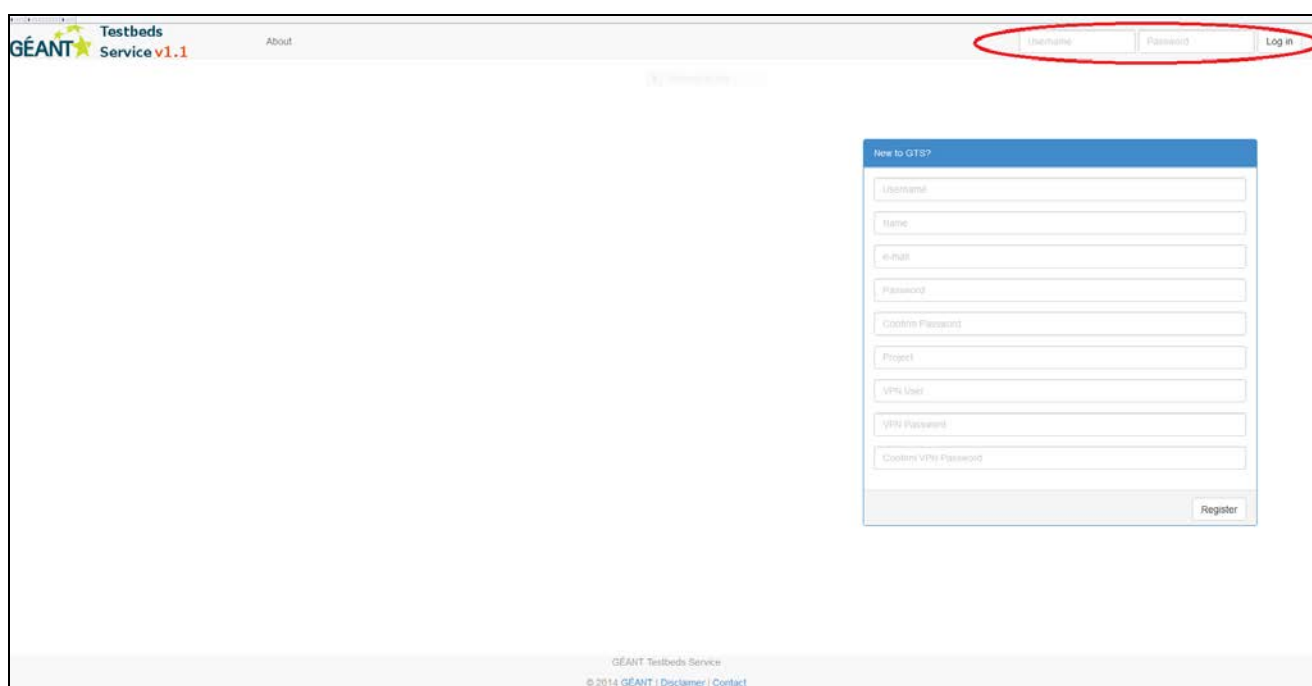


Figure 3: Login as registered user



Figure 4: Project and VPN related information



Figure 5: Log out of the project via the dropdown menu next to the project's name

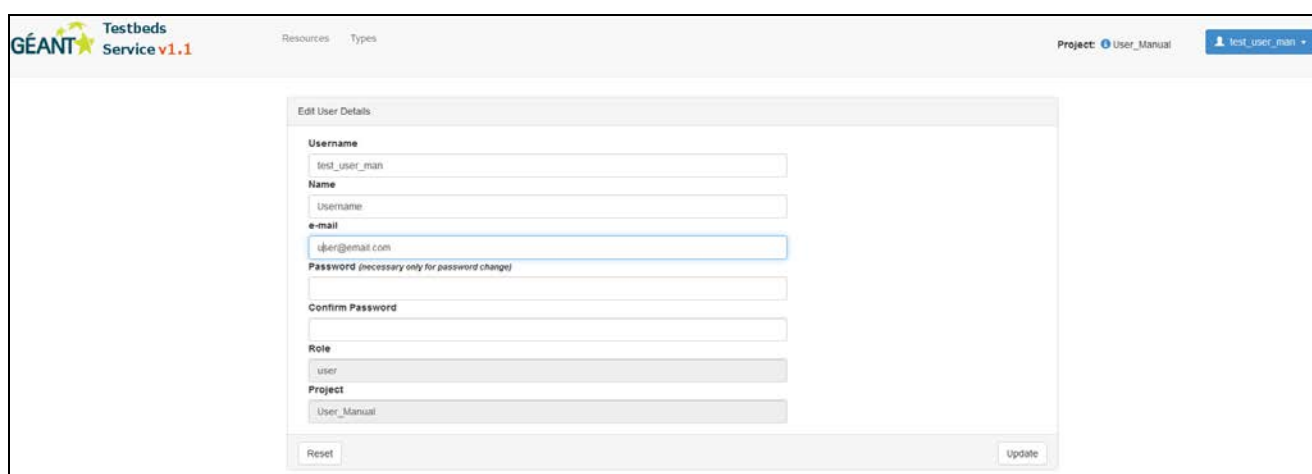


Figure 6: Update or reset user profile

During a login process you may find a disclaimer that indicates that the testbed facility is currently under maintenance and its use may be limited (Figure 7).




Figure 7: Disclaimer

4 Working in the testbed environment

The GTS GUI acts as an agent to the testbed resources and allows a user to build and monitor an isolated domain of interacting resources, i.e. a user's testbed is independent from other users' testbeds and can be operated without any impact on other testbed environments.

After logging in to GTS and naming and creating a new project, the user builds a testbed description for the Testbed Control Agent (TCA). In a next step, the testbed description document is fed to the Resource Manager (RM) which in turn parses the document, allocates resources, and sets up the testbed control plane. The testbed is then activated and the user controls it via the GUI.

For example, the user may decide to define router nodes at first and may then browse and select a set of PC resources to act as the end systems. Each of these resources has data plane ports defined as integral components of the resource. The user must then define the connectivity among these nodal resources by indicating how the various nodal resources are to be interconnected via these network links. At the end the user saves the testbed description to the project's repository and submits the testbed description to the SA2 Service Engine for processing. For some experiments it may be necessary to select virtual resources that are geographically apart or at certain locations. In such use cases users may select a resource's location from the list of available installations.

The SA2 service engine analyzes the testbed description for any obvious errors or omissions, and finding nothing critical, the Service proceeds to locate and reserve the indicated resources. Note: this is an iterative process – the researcher may have only defined a partial set of resources and wishes to reserve these first, where upon the user will then define additional resources to be incorporated into the testbed. Upon successfully reserving all resources, the researcher requests the SA2 Service Engine to instantiate the testbed. As resources are initialized and brought into service, control of those resources are passed to the researcher's own testbed control agent. Instantiating resources is complete, when the selected resources have been activated.

The user can program the desired testbed environments and reserve resources using a Domain Specific Language (DSL) description or select resources via the GUI (for more information on DSL please see Chapter 6). This reservation of resources also allows a user to specify start and end times during when these resources should be activated and made available for experimentation. Only when all requested resources are available the reservation process completes with success.

Reserved resources can then be **activated** by the user and at this time all resources will actually be allocated. At some time during an experiment a user may wish to **deactivate** the resources, but keep the reservation of resources valid for use at a later point in time. Once an experiment is finished and the project is no longer needed, the user should also **release** all resources so that resources can be made available to other projects. Resources can also be **queried** for status information.

4.1 Defining a testbed with DSL

After you have logged in as a registered user you are now able to upload Domain Specific Language (DSL) descriptor files that set up the type of experimental environment that you would like to have in your testbed (for more information on the DSL please see Chapter 6).

You can define your own resource types or classes with DSL code (for example: an OpenFlow switch connected to a controller). To add a DSL file, please click on “Types” (top navigation menu, Figure 8) and then select “Add” (Figure 9) which will open up a window (see Figure 10) where you can browse through your files to upload an already existing DSL file. Click on ‘Submit’ to finish your testbed description.



Figure 8: Select “Types”



Figure 9: Add “Type”

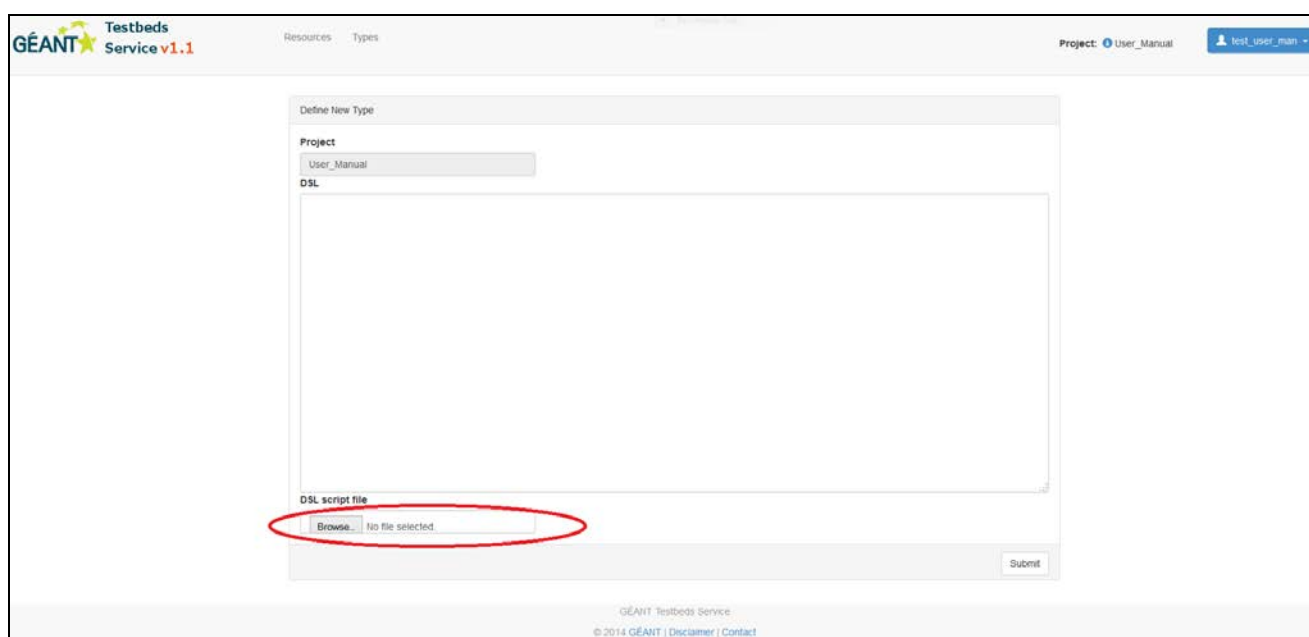


Figure 10: Submitting DSL code via an already existing file

Another option is to configure your testbed by filling in your DSL description using the editor window of the GUI (Figure 11): Just type in your code and click on 'Submit'.

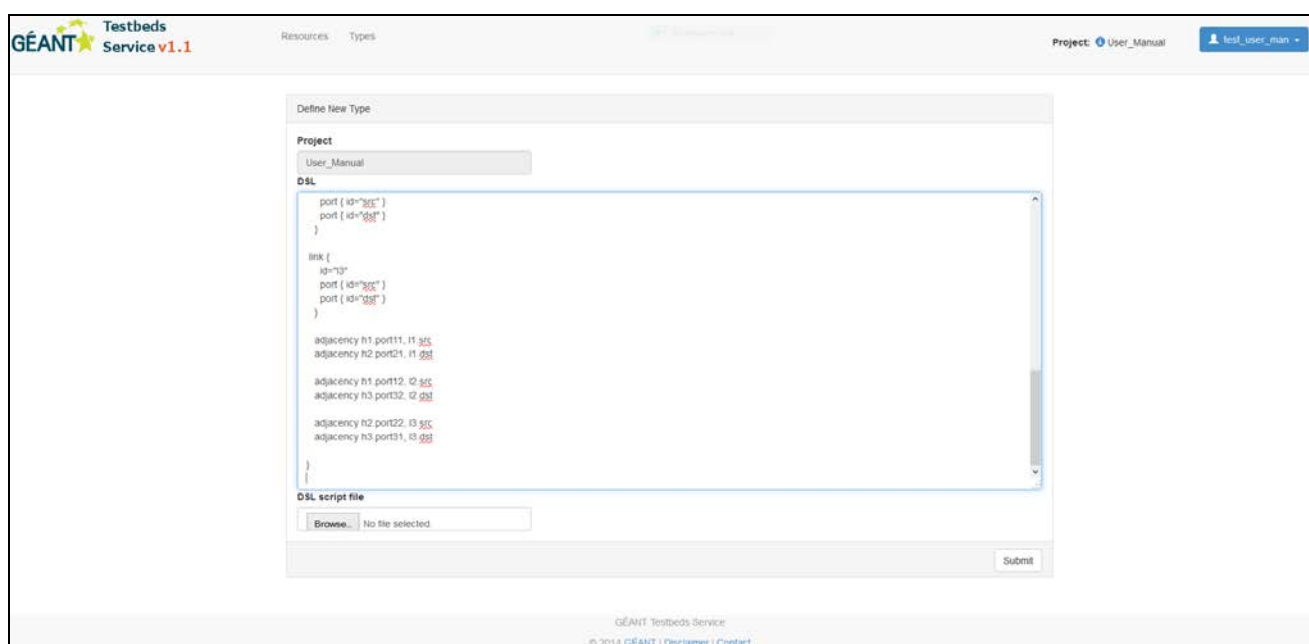



Figure 11: Submitting DSL code directly

4.2 Reserving testbed resources

After submitting your DSL code the system takes you back to the 'Types' screen and you will see the new testbed type you have added to your project (Figure 12)



Types			
Name	Description	Project	Actions
triangle	Triangle between CPH, BRA and LJU	User Manual	 

Figure 12: Added testbed type

When you click on the type's name you have read-only access to your DSL code. For a created testbed type you can choose the following actions:

- You can RESERVE the testbed by clicking on the green button.
- The red button allows you to 'Undefine' or delete the type and the DSL file is dropped.

After you selected the green button to reserve your testbed you will see the following information on your screen that lists your new resource type (Figure 13):

Resources




Provider ID	ID	Status	Type	Actions
+ 407	163180	RESERVED	Testbed 6	  

Figure 13: Reserved new resource type

You can expand the listed information to show you more details on the individual resource components that were reserved as part of your new type by clicking on the '+' sign next to the Provider ID on the left (Figure 14):










Resources						
Provider ID	ID	Status	Type	Actions		
407	163180	RESERVED	Testbed 6			
Host-1423840498510	h3	RESERVED	Host			
Host-1423840498755	h2	RESERVED	Host			
Host-1423840498911	h1	RESERVED	Host			
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	RESERVED	Link			
OpenNsaLink-TA-T628c26a930e9	l2	RESERVED	Link			
OpenNsaLink-TA-Tb265a30714fd	l1	RESERVED	Link			

Figure 14: Listing of reserved resources

The reservation process is complete when your resources show the status 'RESERVED'.

Note: In case you specify a certain location for a resource and your request cannot be granted because not enough VMs are available at that time at that location then your request may go through if you resubmit your DSL without the specification of a location (if that is not a requirement for your experiment in some way).

4.3 Activating testbed resources

Before *reserved* resources can be used, they must be *activated*. Click on the green button on the right to start the activation process (Figure 15).









Resources					
Provider ID	ID	Status	Type	Actions	
407	163180	RESERVED	Testbed 6		
Host-1423840498510	h3	RESERVED	Host		
Host-1423840498755	h2	RESERVED	Host		
Host-1423840498911	h1	RESERVED	Host		
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	RESERVED	Link		
OpenNsaLink-TA-T628c26a930e9	l2	RESERVED	Link		
OpenNsaLink-TA-Tb265a30714fd	l1	RESERVED	Link		

Figure 15: Activate a resource

The activation of resources prompts the user to indicate a start time and end time. The activation can happen in two ways:

1. If you specified your reservation time window (start time and end time) in your DSL, the resources will automatically go from reserved state to active when the time window starts (Figure 16). During that defined time window (i.e. when your reservation is valid) you can actually deactivate and activate your resources as you wish.
2. If you did not specify a reservation window in your DSL, your reservation is valid starting from the current time until practically forever. This will likely change in later versions of GTS! In this case you can activate and deactivate your resources whenever you want.









Testbeds					
Select Reservation to activate					
Reservation Id: 357 Start time: 13-Feb-2015 15:14:00 GMT End time: 31-Jan-2016 22:59:59 GMT					
Provider ID	ID	Status	Type	Actions	
407	163180	RESERVED	Testbed 6		
Host-1423840498510	h3	RESERVED	Host		
Host-1423840498755	h2	RESERVED	Host		
Host-1423840498911	h1	RESERVED	Host		
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	RESERVED	Link		
OpenNsaLink-TA-T628c26a930e9	l2	RESERVED	Link		
OpenNsaLink-TA-Tb265a30714fd	l1	RESERVED	Link		

Figure 16: Reservation start and end time

You can follow the status of your resources while it is changing from 'RESERVED' to 'ACTIVE' (Figure 17):









Resources						
Provider ID	ID	Status	Type	Actions		
407	163180	ACTIVATING	Testbed 6			
Host-1423840498510	h3	ACTIVE	Host			
Host-1423840498755	h2	RESERVED	Host			
Host-1423840498911	h1	RESERVED	Host			
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	RESERVED	Link			
OpenNsaLink-TA-T628c26a930e9	l2	RESERVED	Link			
OpenNsaLink-TA-Tb265a30714fd	l1	RESERVED	Link			

Figure 17: Activating resources










Resources						
Provider ID	ID	Status	Type	Actions		
407	163180	ACTIVE	Testbed 6			
Host-1423840498510	h3	ACTIVE	Host			
Host-1423840498755	h2	ACTIVE	Host			
Host-1423840498911	h1	ACTIVE	Host			
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	ACTIVE	Link			
OpenNsaLink-TA-T628c26a930e9	l2	ACTIVE	Link			
OpenNsaLink-TA-Tb265a30714fd	l1	ACTIVE	Link			

Figure 18: All resources are activated

For activated resources the status turns to 'ACTIVE' (Figure 18); please allow a few minutes for processing as all resources are created dynamically and not just picked from an already existing pool.

4.4 Setting reservation windows

This version does not allow the setting of specific start- and end times yet for activation and deactivation processes. Simply select the provided default settings by clicking on them (Figure 19) (**Please note**: Closing the window is not enough to choose the default values!):



Figure 19: Selecting default start- and end-times

4.5 Deactivating testbed resources

At some time during an experiment a user may wish to deactivate the resources, but keep the reservation of resources valid for use at a later point in time. In other words: Deactivated resources are not released and returned to other users until the experimenter actually *releases* the resources. Deactivated resources can be returned to an active state by simply activating them again at any time during the experiment.

To deactivate a resource, simply click on the yellow button on the right (Figure 20):









Resources					
Provider ID	ID	Status	Type	Actions	
407	163180	ACTIVE	Testbed 6		
Host-1423840498510	h3	ACTIVE	Host		
Host-1423840498755	h2	ACTIVE	Host		
Host-1423840498911	h1	ACTIVE	Host		
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	ACTIVE	Link		
OpenNsaLink-TA-T628c26a930e9	l2	ACTIVE	Link		
OpenNsaLink-TA-Tb265a30714fd	l1	ACTIVE	Link		

Figure 20: Deactivation of a resource

The start- and end-time reservation window will pop up again where you can specify the time when your resources are to be deactivated (in this version only available with default setting); select the (default) setting by clicking on the line with the start and end times (closing the window is not enough!) and wait for the status of the resources to change to 'RESERVED' again (Figure 21). But that does not mean that they are released (please see section 4.6 for more details); instead they remain reserved for the user to be available for the next activation process.









Testbeds					
Select Reservation to deactivate					
Reservation ID: 357 Start time: 13-Feb-2015 15:14:00 GMT End time: 31-Jan-2026 22:59:59 GMT					
Provider ID	ID	Status	Type	Actions	
407	163180	ACTIVE	Testbed 6		
Host-1423840498510	h3	ACTIVE	Host		
Host-1423840498755	h2	ACTIVE	Host		
Host-1423840498911	h1	ACTIVE	Host		
OpenNsaLink-TA-Tdca9f9cb2ddd	l3	ACTIVE	Link		
OpenNsaLink-TA-T628c26a930e9	l2	ACTIVE	Link		
OpenNsaLink-TA-Tb265a30714fd	l1	ACTIVE	Link		

Figure 21: Selecting an end-time

4.6 Releasing testbed resources

Resources that are no longer of interest should be *released* so that these resources can be made available to other projects and users.




Resources					
Provider ID	ID	Status	Type	Actions	
+ 2491	8963620	RESERVED	Testbed 6		 

Figure 22: Releasing resources

To release the testbed resource click on the red button on the right as marked in Figure 22.

4.7 Querying details

The individual resources that were reserved for your testbed become visible when you click on the '+' sign to the left of the Provider ID of your resource (Figure 23); each resource can then also be accessed by clicking on its link (Figure 24):





Resources					
Provider ID	ID	Status	Type	Actions	
 2491	8963620	RESERVED	Testbed 6		 

Figure 23: Obtaining details via Provider ID











Resources					
Provider ID	ID	Status	Type	Actions	
 2491	8963620	RESERVED	Testbed 6		 
Host-1425297436046	h3	RESERVED	Host		
Host-1425297436919	h2	RESERVED	Host		
Host-1425297437784	h1	RESERVED	Host		
OpenNsaLink-TA-T863cc624f840	l3	RESERVED	Link		
OpenNsaLink-TA-T6ee5ff17d8b1	l2	RESERVED	Link		
OpenNsaLink-TA-T115cf06baffe	l1	RESERVED	Link		

Figure 24: List of individual testbed resources

By clicking on a resource you can also obtain more information on ports, locations as well as reservation start and end times (Figure 25):

Resource: Host						
Provider id	Id	Status	Description	Ports	Location	Console
Host-1425297436046	h3	RESERVED		port32 port31	AMS	http://62.40.126.10:6080/vnc_auto.html?token=ae65b422-d799-4753-853b-9adc19851d69

Reservations		
Id	Start Time	End Time
2553	Mon, 02 Mar 2015 11:57:00 GMT	Thu, 31 Jan 2526 22:59:59 GMT

Figure 25: List of individual testbed resources

To obtain a tree structure of your testbed resource and to see port adjacencies or adjacencies for external ports, click on the Provider ID itself (Figure 26 and Figure 27):




Resources					
Provider ID	ID	Status	Type	Actions	
+ 2491	8963620	RESERVED	Testbed 6	  	

Figure 26: Access to tree structure

Tree Structure	
<div> <div>▼ 2491</div> <ul style="list-style-type: none"> Host-1425297436046 Host-1425297436919 Host-1425297437784 OpenNsaLink-TA-T863cc624f840 OpenNsaLink-TA-T6ee5ff17d8b1 OpenNsaLink-TA-T115cf06baffe </div>	
Port Adjacencies for Nested Resources	
Nested Resources	Adjacent Ports
Host-1425297436046	Host-1425297436046:port32 ⇌ OpenNsaLink-TA-T6ee5ff17d8b1:dst Host-1425297436046:port31 ⇌ OpenNsaLink-TA-T863cc624f840:dst
Host-1425297436919	Host-1425297436919:port22 ⇌ OpenNsaLink-TA-T863cc624f840:src Host-1425297436919:port21 ⇌ OpenNsaLink-TA-T115cf06baffe:dst
Host-1425297437784	Host-1425297437784:port12 ⇌ OpenNsaLink-TA-T6ee5ff17d8b1:src Host-1425297437784:port11 ⇌ OpenNsaLink-TA-T115cf06baffe:src
OpenNsaLink-TA-T863cc624f840	OpenNsaLink-TA-T863cc624f840:src ⇌ Host-1425297436919:port22 OpenNsaLink-TA-T863cc624f840:dst ⇌ Host-1425297436046:port31
OpenNsaLink-TA-T6ee5ff17d8b1	OpenNsaLink-TA-T6ee5ff17d8b1:src ⇌ Host-1425297437784:port12 OpenNsaLink-TA-T6ee5ff17d8b1:dst ⇌ Host-1425297436046:port32
OpenNsaLink-TA-T115cf06baffe	OpenNsaLink-TA-T115cf06baffe:src ⇌ Host-1425297437784:port11 OpenNsaLink-TA-T115cf06baffe:dst ⇌ Host-1425297436919:port21
Adjacencies for External Ports	
No external ports for this testbed	

Figure 27: Tree structure and port adjacencies

The blue button in the section 'Actions' in your resource list offers you a visualization with Google Earth Plugin (Figure 28 and Figure 29):




Resources					
Provider ID	ID	Status	Type	Actions	
+ 2491	8963620	RESERVED	Testbed 6	  	

Figure 28: Access to visualization



Figure 29: Visualization with Google Earth Plugin²

² Google and the Google log are registered trademarks of Google Inc., used with permission.

4.8 Internet Access Gateway (IAGW)

In order to be able to upload software to your testbed environment there is an Internet Access Gateway (IAGW) you can reach using the VPN information that you defined during your initial user registration process. The VPN is managed by a VM set up at the GTS Central Server Facility and remains in place as long as at least one of your hosts is active.

4.8.1 Using IAGW with Windows

For systems using versions of Windows set up a new VPN network connection via the Control Panel and its Network and Internet options.

To set up the VPN use the IP address shown when you click on the blue 'i' icon next to your project name (Figure 30 and Figure 31) and also the VPN user name and password you defined during the initial project registration process.



Figure 30: VPN information

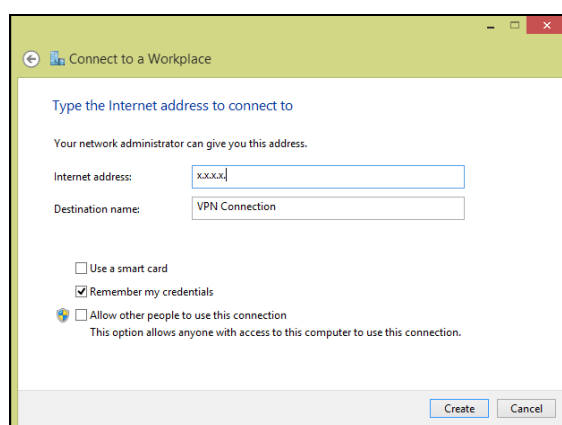


Figure 31: VPN IP address

Once you are connected via VPN, use a file transfer program like WinSCP or PuTTY to transfer software to a testbed host. Before you use WinSCP or PuTTY access the console of the host that you would like to transfer to by clicking on the appropriate host (Figure 32) which will then provide a console link (Figure 33).

Resources						
Provider ID	ID	Status	Type	Actions		
2491	8963620	ACTIVE	Testbed 6			
Host-1425297436046	h3	ACTIVE	Host			
Host-1425297436919	h2	ACTIVE	Host			
Host-1425297437784	h1	ACTIVE	Host			
OpenNsaLink-TA-T863cc624f840	l3	ACTIVE	Link			
OpenNsaLink-TA-T6ee5ff17d8b1	l2	ACTIVE	Link			
OpenNsaLink-TA-T115cf06baffe	l1	ACTIVE	Link			

Figure 32: Accessing the console of a host

Resource: Host						
Provider id	Id	Status	Description	Ports	Location	Console
Host-1425297436919	h2	ACTIVE		port22 port21	AMS	http://62.40.126.10:6080/vnc_auto.html?token=85b2f30a-3612-429a-8688-5bb5be0d83fb

Reservations		
Id	Start Time	End Time
2554	Mon, 02 Mar 2015 11:57:00 GMT	Thu, 31 Jan 2526 22:59:59 GMT

Figure 33: Console link

Click on the console link; you will then be prompted for user name and password (gts, gts). Look for the IP address provided on the screen for `eth0` (Figure 34); this is the IP address you should use for your file transfer program along with user name and password (gts, gts). You are then ready to transfer your software to your host.

```

Ubuntu 14.04.1 LTS sa2host tty1
sa2host login: gts
Password:
Last login: Fri Mar  6 10:11:08 CET 2015 on tty1
Welcome to Ubuntu 14.04.1 LTS (GNU/Linux 3.13.0-37-generic x86_64)

* Documentation:  https://help.ubuntu.com/

System information as of Fri Mar  6 10:11:08 CET 2015

System load:  0.0          Processes:      68
Usage of /:   16.3% of 6.76GB Users logged in: 0
Memory usage: 1%          IP address for eth0: 172.16.0.11
Swap usage:   0%

Graph this data and manage this system at:
https://landscape.canonical.com/

gts@sa2host:~$ _

```

Figure 34: IP address for `eth0`

4.8.2 Using IAGW with Debian / Ubuntu

For a VPN connection to GTS using Debian / Ubuntu proceed with the following steps:

- Get into root environment
\$sudo su
- Switch into directory 'peers'
#cd /etc/ppp/peers
- Setup up a VPN configuration file 'projectname' and adjust the values marked in red (Figure 35):

```
### Build VPN connection to GTS VPN server
pty "pptp <VPN IP> --nolaunchpppd --nobuffer --timeout 10"

### GTS VPN - user name
name <VPN user name>

### This command reconnects to GTS VPN server after a disconnect
persist

### MTU value should be standard value.
mtu 1400

### used for scripts in ip-up / ip-down
ipparam <YOURPROJECTNAME>

### Terminate after n consecutive failed connection attempts.
### A value of 0 means no limit. The default value is 10.
maxfail 0

### IMPORTANT: Do not use this option for the GTS VPN connection.
### Otherwise all packets are routed over this connection.
### Please comment it out
#defaultroute

### Some useful settings.
remotename <YOURPROJECTNAME>
lock
noauth
nobsdcomp
nodeflate

#Require encrypt
require-mppe-128
```

Figure 35: Adjust VPN configuration file

To find out the correct IP and user name that you should set in this VPN configuration file click on the blue 'i' icon next to your project name (Figure 36) and use the values you set during your initial project registration and then save the file.



Figure 36: VPN information

- In the next step adjust the file `chap-secrets` indicating your VPN user name, your project name and VPN password that you set during your initial project registration (Figure 37):

```
# Secrets for authentication using CHAP
# client      server      secret      IP addresses
# IMPORTANT: <VPN PASSWORD> MUST be set in double quotes
<VPN USER> <YOURPROJECTNAME> "<VPN PASSWORD>" *
#end
```

Figure 37: Adjust file `chap-secrets`

- Then switch to the directory `ip-up.d`
`#cd ip-up.d/`
- and set the project route with an executable file `yourprojectname.route` (Figure 38)

```
#!/bin/sh
# IMPORTANT: Use here the network IP of interface eth0 of your
# GTS VM (NOT the VPN IP!)
if [ "${PPP_IPPARAM}" = "<YOURPROJECTNAME>" ]; then
  /sbin/route add -net <YOUR NETWORK IP>/24 dev <PPP INTERFACE NAME>
fi
```

Figure 38: Set project route with executable file

and make sure you use the network IP address of `eth0` of your host. To find this IP address click on the appropriate host (Figure 39) which will then provide a console link (Figure 40).

Resources					
Provider ID	ID	Status	Type	Actions	
2491	8963620	ACTIVE	Testbed		
Host-1425297436046	h3	ACTIVE	Host		
Host-1425297436919	h2	ACTIVE	Host		
Host-1425297437784	h1	ACTIVE	Host		
OpenNsaLink-TA-T863cc624f840	l3	ACTIVE	Link		
OpenNsaLink-TA-T6ee5ff17d8b1	l2	ACTIVE	Link		
OpenNsaLink-TA-T115cf06baffe	l1	ACTIVE	Link		

Figure 39: Accessing the console of a host

Resource: Host						
Provider id	Id	Status	Description	Ports	Location	Console
Host-1425297436919	h2	ACTIVE		port22 port21	AMS	http://62.40.126.10:6080/vnc_auto.html?token=85b2f30a-3612-429a-8688-5bb5be0d83fb

Reservations		
Id	Start Time	End Time
2554	Mon, 02 Mar 2015 11:57:00 GMT	Thu, 31 Jan 2526 22:59:59 GMT

Figure 40: Console link

Click on the console link; you will then be prompted for user name and password (gts, gts). Look for the IP address provided on the screen for `eth0` (Figure 41); this is the IP address you should set as **<YOUR NETWORK IP>** in the file `'yourprojectname.route'` as shown in Figure 38.

```
gts@sa2host:~$ ifconfig
eth0      Link encap:Ethernet  HWaddr fa:16:3e:01:6d:30
          inet addr:172.16.0.10  Bcast:172.16.0.255  Mask:255.255.255.0
          inet6 addr: fe80::f01b:3eff:fe01:6d30/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:13284 errors:0 dropped:0 overruns:0 frame:0
          TX packets:10753 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:68437809 (68.4 MB)  TX bytes:1001988 (1.0 MB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:25 errors:0 dropped:0 overruns:0 frame:0
          TX packets:25 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1904 (1.9 KB)  TX bytes:1904 (1.9 KB)

gts@sa2host:~$
```

Figure 41: IP address for `eth0`

In the file `'yourprojectname.route'` you also define the **<PPP INTERFACE NAME>** (see Figure 38) as `pppx` where 'x' would be '0' if you have no other VPNs defined (or else would be incremented to the appropriate number, but please **note**: only count VPNs set up with `ppp`, not VPNs set up with OpenVPN).

- This file `'yourprojectname.route'` should then be made into an executable file
`#chmod 700 yourprojectname.route`
- To connect with the VPN use
`#pon projectname`
- To disconnect the VPN use
`#poff projectname`

- **To test:** Type
#ifconfig
and look for your pppx with address 172.16.0.1 (Figure 42):

```
ppp0      Link encap:Point-to-Point Protocol
          inet addr:172.16.0.1  P-t-P:172.16.0.253  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1396  Metric:1
          RX packets:14584 errors:0 dropped:0 overruns:0 frame:0
          TX packets:13996 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:3
          RX bytes:2568663 (2.4 MiB)  TX bytes:966234 (943.5 KiB)
```

Figure 42: check for pppx

- and then check if you can ping your host, for example
#ping -c 5 172.16.0.10
- if your ping does not work, check with
#ip route show
if you have the following entry:
172.16.0.0/24 dev ppp0 scope link
- **To upload software:** ssh into your host with user name and password (gts,gts)
#ssh gts@172.16.0.10 and upload files with
#scp filename gts@172.16.0.10:
- **Please note:** Currently GTS does not support automatic security updates on VMs. It is therefore recommended to switch into root environment on the host with password (gts)
\$sudo su
and update / upgrade the list of available packages with
#aptitude update
#aptitude upgrade

4.9 Example: Working with GTS OpenFlow switches (OFXs)

The following steps guide you through a DSL example³ that sets up a simple OpenFlow environment with two OpenFlow switches (OFXs)⁴, a controller and three Virtual Machines (VMs). Each of the two OpenFlow switches is at a different location (in Ljubljana (LJU) and in Bratislava (BRA)) and they are both connected:

Step 1: Define the VM that will be the controller

```
TwoOFXwithController {
  description = "Two OFX with one controller VM. The OFXs are at
different locations and connected."
  id = "TwoOFXwithController"

  host {
    id="controller"
    port { id="eth1" }
    port { id="eth2" }
  }
}
```

A host called “controller” is defined as controller; the controller has two ports, because it will later have to connect to each of the two OpenFlow switches in LJU and BRA.

Step 2: Define a host (VM) in Bratislava

```
host {
  id="dataBRAVM1"
  location="BRA"
  port { id="eth1" }
}
```

The host has one port for data transfer.

³ The complete DSL code is available in the Appendix, Section VII

⁴ In GTS OpenFlow hardware switches are separated into instances referred to as OFX switch fabrics.

Step 3: Define OpenFlow switch in BRA

The following DSL code adds an OpenFlow switch in BRA with IP address 10.10.10.1. The controller port for the switch is on port "pControl" (mode="CONTROL"). RYU [RYU-2014] by default launches the controller instance in port 6633, but in this example the port '9966' was specified:

```
ofx {
  id="ofxBRA"
  location="BRA"
  fabricIPv4="10.10.10.1"
  fabricSubnetMaskv4="255.255.255.0"
  controllerPort="9966"
  controllerIPv4="10.10.10.10"
  port { id="p1" }
  port { id="p2" }
  port { id="pLJU" }
  port {
    id="pControl"
    mode="CONTROL"
  }
}
```

Step 4: Connect components in BRA

As a next step, the host "dataBRAVM1" in Bratislava must be connected to the OpenFlow switch and the Switch must be connected to the controller. This is achieved by

- defining the required links ("l1BRA" for the data plane between host and OpenFlow switch and "l1CBRA" for the control plane between the controller and the OpenFlow Switch)
- and** by defining the adjacencies, which actually link the objects:

```
link {
  id="l1BRA"
  port { id="src" }
  port { id="dst" }
}

link {
  id="l1CBRA"
  port { id="src" }
  port { id="dst" }
}

adjacency dataBRAVM1.eth1, l1BRA.src
adjacency ofxBRA.p1, l1BRA.dst

adjacency controller.eth1, l1CBRA.src
adjacency ofxBRA.pControl, l1CBRA.dst
```

Step 5: Add a second host in Bratislava

Now add a second host in BRA by adding the following DSL code:

```
host {
  id="dataBRAVM2"
  location="BRA"
  port { id="eth1" }
}

link {
  id="l2BRA"
  port { id="src" }
  port { id="dst" }
}

adjacency dataBRAVM2.eth1, l2BRA.src
adjacency ofxBRA.p2, l2BRA.dst
```

Step 6: Connect another OpenFlow switch in Ljubljana

In this example the second OpenFlow switch is located in Ljubljana (LJU) with IP address 10.10.11.1. The controller port for the switch is

on port "pControl" (mode="CONTROL").

```
ofx {
  id="ofxLJU"
  location="LJU"
  fabricIPv4="10.10.11.1"
  fabricSubnetMaskv4="255.255.255.0"
  controllerPort="9966"
  controllerIPv4="10.10.11.10"
  port { id="p1" }
  port { id="pBRA" }
  port {
    id="pControl"
    mode="CONTROL"
  }
}
```

Step 7: Add a host to the OpenFlow switch in LJU, connect host to OpenFlow switch, connect OpenFlow switch to controller, and connect the two OpenFlow switches

```

host {
    id="dataLJUVM1"
    location="LJU"
    port { id="eth1" }
}

link {
    id="l1LJU"
    port { id="src" }
    port { id="dst" }
}

link {
    id="loFXBRAOFXLJU"
    port { id="src" }
    port { id="dst" }
}

link {
    id="l1CLJU"
    port { id="src" }
    port { id="dst" }
}

adjacency controller.eth2, l1CLJU.src
adjacency ofxLJU.pControl, l1CLJU.dst

adjacency dataLJUVM1.eth1, l1LJU.src
adjacency ofxLJU.p1, l1LJU.dst

adjacency ofxBRA.pLJU, loFXBRAOFXLJU.src
adjacency ofxLJU.pBRA, loFXBRAOFXLJU.dst

```

Step 8: Configuring the controller interfaces

Once you have reserved your testbed, activate the resources and make sure that all resources are marked as 'ACTIVE'.

- find your controller in the list of resources
- click on the controller in order to get a link to its console (as you would do to get access to the console of any VM; see Figure 43, Figure 44, compare also with Figure 40)

Resources						
Provider ID	ID	Status	Type	Actions		
2499	t1	RESERVED	Testbed			
Host-1425660497378	controller	RESERVED	Host			
Host-1425660498295	h1	RESERVED	Host			
OFX-af387cc6-cd7d-4dd4-a4fc-c151386f3a4d	ofx1	RESERVED	OFX			
OpenNsaLink-TA-Tb1206186cc2e	lc	RESERVED	Link			
OpenNsaLink-TA-T582c87e151d7	l1	RESERVED	Link			

Figure 43: Access to controller VM

Resource: Host						
Provider id	Id	Status	Description	Ports	Location	Console
Host-1425660497378	controller	ACTIVE		eth1	LJU	http://62.40.126.10:6080/vnc_auto.html?token=0f156eb1-2ffd-4f51-94a8-51cb0c401f71

Reservations		
Id	Start Time	End Time
2563	Fri, 06 Mar 2015 16:48:00 GMT	Thu, 31 Jan 2526 22:59:59 GMT

Figure 44: Link to controller console

- at the console you will be prompted for a user name and password (gts, gts)
- at the moment it is still necessary to manually configure the interfaces of the controller (in later versions of GTS it will be possible to simply set this via the GUI); this is because GTS currently does not support port mapping.
- **please note:** 192.168.x.x and eth0 should NOT be used, as they are currently reserved for internal processes!
- configure eth1 interface of controller for BRA (as already defined in the DSL example):

```
$sudo su
#ifconfig eth1 up 10.10.10.10 netmask 255.255.255.0
and check for incoming ARP packets with
#tcpdump -i eth1
```

You will see the incoming echo request from the OFX like this:

```
#16:49:58.702223 IP 10.10.10.1.59938 > 10.10.10.10.6633: Flags [S], seq 688712374,
```


win 65535, options [mss 1460,nop,wscale 3,sackOK,TS val 4260253826 ecr 0], length 0
i.e. an OFX with IP address 10.10.10.1 requests IP 10.10.10.10.

- open the interfaces configuration file with your editor
/etc/network/interfaces
- add the following entries to the file:
auto eth1
iface eth1 inet static
address 10.10.10.10
netmask 255.255.255.0
network 10.10.10.0
- stop interface with
#ifdown eth1
- start interface with
#ifup eth1
- check if you can ping the OFX in BRA with IP address 10.10.10.1
#ping -c 10.10.10.1
- repeat the same process for eth2 interface of controller for LJU (as already defined in the DSL example):
\$sudo su
#ifconfig eth2 up 10.10.11.1 netmask 255.255.255.0

Step 9: Example: Installing RYU as OpenFlow Controller

GTS OpenFlow switches support OpenFlow specification 1.3. Therefore we suggest to use RYU as OpenFlow controller in GTS, because it fully supports OpenFlow specification 1.3. To install the RYU OpenFlow controller the following steps are required:

- get root access with
\$sudo su
- install dependencies
#aptitude install python-setuptools python-pip python-dev libxslt-dev
libxml2-dev zlib1g-dev wget
- get the installation script from RYU
#wget https://raw.githubusercontent.com/sdn2s-tw/ryuInstallHelper/master/ryuInstallHelper.sh
- make the script executable
#chmod +x ryuInstallHelper.sh
- execute the installer
./ryuInstallHelper.sh

- this should install RYU applications and libraries. Follow the output messages during the installation process as in some cases errors may occur where missing packages have to be installed and the installer has to be run again
- you can find all the examples and test scripts under
`#/usr/local/lib/python2.7/dist-packages/ryu`
- to see if the controller works, you can change to the directory with applications
`#cd /usr/local/lib/python2.7/dist-packages/ryu/app/`
and run a simple application, e.g. `simple_switch_13.py`
- RYU uses port "6633" as default ; if you have specified another port in your DSL (such as `controllerPort="9966"` in the example above) you must also adjust the port number with
`#ryu-manager simple_switch_13.py -ofp_tcp-listen-port 9966`
- **please note:** GTS uses OpenFlow version 1.3 switches, therefore only apps for OpenFlow version 1.3 will work
- **please note:** In this version of GTS it is strongly discouraged to change the `datapathID` (Figure 45) that identifies your OpenFlow switch fabric as this identifier has to be absolutely unique!!

GÉANT Testbed Service (GTS)
User and Resource Guide
Document Code:

5 Help and Support

For help and support please contact

- **First Line Contact:** GÉANT Operations Centre
mailto: support@gts.geant.net; phone: +44 1223 866140
- **Second Line Escalation:** SA2-Operations Team
mailto: gn3plus-sa2-operations@geant.net.

The first line of contact would have to be informed if the functionality of the GTS service seems to have been compromised or some equipment seems to be broken, whereas the second line of support would be approached in case of problems with DSL or configuration of the testbed. You also find this support information by clicking on 'About' on the GTS home page (Figure 46). The support page also provides a link where you can open a ticket in case you encounter a problem with the system (Figure 47).



Figure 46: Access to support

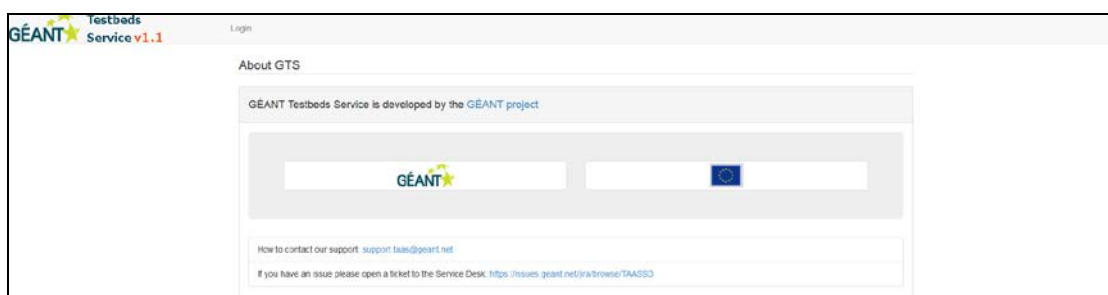


Figure 47: Important links for support

Additional information and training videos are available at <http://services.geant.net/gts/>.

6 Introduction to Domain Specific Language

6.1 DSL in GTS

GTS uses its own Domain Specific Language (DSL) to describe topologies. It is used by users to specify their requests, but also by GTS service to represent the allocated topology that conforms to users' requirements. DSL not only allows defining requests, but also offers Groovy [GRO-2014] programming language syntax that eases up the process of specifying scalable, complex topologies. In this section only DSL syntax is described. Readers should get familiar with Groovy programming language grammar if they want to use programming language features. For basic request definition knowing only DSL syntax is enough. The section is divided into two parts: Quick start and DSL grammar definition. The first offers introduction with examples where after five minutes users will know all concepts and should be able to define basic requests. The latter provides information about all resource types, their parameters and BNF (Backus–Naur Form) grammar of the DSL.

6.2 Quick Start

For every virtual resource there is a resource class or type, called a 'template' that defines all possible attributes of that resource. The process of creating a resource instance from this class template is referred to as 'instantiation'. In a testbed, a researcher can instantiate several instances of a resource and distinguish them by assigning each resource instance a unique name [DEL-D61].

Each resource class has

- a class name / an identifier of that class (not instance!) (name cannot start with a digit)
- external ports (where each port has a port name and a set of attributes for those ports)
- class attributes
- internal resources
- adjacencies among internal resources

- and a set of resource control primitives defined for that resource class.

Resources can be either atomic or composite. Atomic resources denote virtual resources that will be allocated for the user. Those that are composite are logical containers of other resources. This feature allows to define scalable, complex topologies, because composite types form reusable building blocks can be combined together.

Before going into formal BNF grammar definition let's start with a simple example to get familiar with the DSL syntax. The code snippet below is a definition of a type containing one host. The type called *HostWithTwoPorts* is a composite that contains one host. The host has two interfaces called ports *p1* and *p2*. Each resource instance definition consists of a resource type name and a set of parameters and ports specified in curly braces. The parent-child relation is specified by defining the resource instance within the definition of the composite resource. In this example the host has *HostWithTwoPorts* composite type as a parent.

```
type HostWithTwoPorts {  
    description = "Host that has two ports."  
  
    host {  
        id = "host"  
        port { id = "p1" }  
        port { id = "p2" }  
    }  
}
```

Let's imagine a more complex topology where we have three hosts connected to each other forming a triangle. We will reuse *HostWithTwoPorts* type to specify three hosts. Each resource can be connected with another resource via a transport resource called *link*. Links have exactly two ports called *src* and *dst* denoting source and sink. We need three hosts and three links.

```

type Triangle {
  description = "Triangle using PRG, LJU and BRA hosts."

  HostWithTwoPorts {
    id = "h1"
    host.location = "prg"
  }

  HostWithTwoPorts {
    id = "h2"
    host.location = "lju"
  }

  HostWithTwoPorts {
    id = "h3"
    host.location = "bra"
  }

  link { id = "l1" }
  link { id = "l2" }
  link { id = "l3" }
  ...
}

```

The id parameter in each resource is a user defined identifier that is unique only within the given DSL. It does not have to be a globally unique identifier. It helps users specifying properties and adjacencies. We used ids to define locations of internal nodes. The other benefit is when debugging what went wrong during instantiation of the request. Having hundreds of resources with meaningful identifiers greatly helps to find the cause of the problem in the request. The last step that we need in order to define the triangle are adjacencies between resources. We are using composites to connect to link resources.

```

type HostWithTwoPorts {
    description = "Host that has two ports."

    host {
        id = "host"
        port { id = "p1" }
        port { id = "p2" }
    }

    port { id = "p1" }
    port { id = "p2" }

    adjacency p1, host.p1
    adjacency p2, host.p2
}

```

In order to expose internal details, ie. ports of the underlying host, logical ports must be defined in the composite and connected to the host's ports. Then these logical ports will be used to connect to link ports. Having *HostWithTwoPorts* revisited we can define triangle adjacencies.

```

type Triangle {
    ...
    adjacency h1.p1, l1.src
    adjacency h2.p1, l1.dst
    adjacency h1.p2, l2.src
    adjacency h3.p2, l2.dst
    adjacency h2.p2, l3.src
    adjacency h3.p1, l3.dst
}

```

We specify that host *h1* port *p1* is connected to source of link *l1*. Link *l1* sink is connected to host *h2* port *p1* and so on. As you see we use user specified identifiers to help us denote specific resource instances.

Groovy programming language features come in handy when there is a need to specify complex topologies with possibly hundreds of connections. It can be done by defining procedures for doing so. Even the Triangle example can benefit from using programming. This way we do not have to define an additional type called *HostWithTwoPorts*. Take a look at the example below. We use a loop three times to define three hosts and three links. To ease up iterating over resources we defined two lists hosts and links. The second loop is used to connect all resources together.


```

type Triangle {
  description = "Triangle using Groovy language iterators to define adjacencies."
  id = "t1"

  def hosts = []
  def links = []

  3.times { idx ->
    def h1 = host {
      id = "host$idx"
      port { id = "p1" }
      port { id = "p2" }
    }
    hosts << h1

    def l1 = link {
      id = "link$idx"
      port { id = "src" }
      port { id = "dst" }
    }
    links << l1

    adjacency h1.p1, l1.src
  }

  3.times { idx -> adjacency hosts[(idx + 1) % 3].p2, link[idx].dst }
}

```

Appendix I: Resource Guide

Each resource type has a different set of properties that can be specified or read by users. This section provides a list of resource types and possible properties with description of their field types and possible ranges.

I. Composite types

Aggregate logical resources are specified by this type.

List of parameters:

Parameter	Type	Description	Example
description	String <read, write, optional>	Characterizes the type for the user.	Host with two ports.
id	String <read, write, optional>	User defined identifier. If not specified it is autogenerated by the service.	h1
providerId	String <read-only>	Globally unique identifier of the resource instance. Resource instances have this value set by the service.	nsi20010495
status	RESERVED, ACTIVATING, ACTIVE, DEACTIVATING, RELEASING <read-only>	Current state of the resource.	ACTIVE
ports	Map<id, port> <read only>	Map of ports defined by the user.	port { id = "p1" } port { id = "p2" }

Example:

```
type HostWithTwoPorts {
    description = "Host that has two ports."

    host {
        id = "host"
        port { id = "p1" }
        port { id = "p2" }
    }
}
```

II. Host

A virtual machine is depicted by this type. Each VM has the following components preinstalled:

- Ubuntu 12.04.4 LTS (GNU/Linux 3.8.0-29-generic x86_64)

List of parameters:

Parameter	Type	Description	Example
location	PRG, BRA, LJU, AMS <read, write, optional>	location of the resource	PRG
id	String <read, write, optional>	User defined identifier. If not specified it is autogenerated by the service.	h1
providerId	String <read-only>	Globally unique identifier of the resource instance. Resource instances have this value set by the service.	nsi20010495
status	RESERVED, ACTIVATING, ACTIVE, DEACTIVATING, RELEASING <read-only>	Current state of the resource.	ACTIVE
ports	Map<id, port> <read only>	Map of ports defined by the user.	port { id = "p1" } port { id = "p2" }

Default attributes		Comments
Host { }	cpuSpeed == 2.3	The minimum clock speed of the CPU in GHz.
	image = "<string>"	The ISO image of the operating system installed on the host by default. The current operation system is "Ubuntu 14.04"
	disk == 20	Disk space coming with the host* defined in GB. <i>* Later on persistent storage (NFS) will also be available.</i>
	consoleAccess == "<public URL>"	Public URL of the console access to the host.
port { }	Framing == "<Eth>"	Only Ethernet (802.1Q) framing is available.

Example:

```
host {
  id = "host"
  location = "prg"
  port { id = "p1" }
}
```

III. Link

Represents a virtual circuit between resources. Always has exactly two ports for source and sink.

List of parameters:

Parameter	Type	Description	Example
id	String <read, write, optional>	User defined identifier. If not specified it is autogenerated by the service.	h1
providerId	String <read-only>	Globally unique identifier of the resource instance. Resource instances have this value set by the service.	nsi20010495
status	RESERVED, ACTIVATING, ACTIVE, DEACTIVATING, RELEASING <read-only>	Current state of the resource.	ACTIVE
ports	Map<id, port> <read only>	Map of exactly two ports source and sink	port { id = "src" } port { id = "dst" }

	Default attributes	Comments
Link { }	capacity = {1...10000}	Capacity of the link defined in Mbps. Default value is 10 Mbps
port { }	Framing = "<Eth>"	Only Ethernet (802.1Q) framing is available
	lineRate = {1, 10}	Nominal line rate of the port defined in Gbps. Default value is 1 Gbps.
	directionality = {uniDir, biDir}	Directionality of the port. Default value is bidirectional: "biDir"

Example:

```
link {
  id = "host"
}
```

IV. OFX

OpenFlow switches are represented by this type. Please note that one of ofx ports must be denoted with mode = "CONTROL" to inform the service that this port is going to be a control port.

GTS currently uses Hewlett Packard switches of the HP 5900 Series which support OpenFlow specification 1.3:

HP Comware Software, Version 7.1.035, Release 2208P01
 Copyright (c) 2010-2013 Hewlett-Packard Development Company, L.P.
 HP 5900AF-48G-4XG-2QSFP+

Boot image: flash:/5900_5920-cmw710-boot-r2208p01.bin
 Boot image version: 7.1.035P05, Release 2208P01
 System image: flash:/5900_5920-cmw710-system-r2208p01.bin
 System image version: 7.1.035, Release 2208P01

List of parameters:

Parameter	Type	Description	Example
location	PRG, BRA, LJU, AMS <read, write, optional>	location of the resource	PRG
id	String <read, write, optional>	User defined identifier. If not specified it is autogenerated by the service.	h1
providerId	String <read-only>	Globally unique identifier of the resource instance. Resource instances have this value set by the service.	nsi20010495
status	RESERVED, ACTIVATING, ACTIVE, DEACTIVATING, RELEASING <read-only>	Current state of the resource.	ACTIVE
ports	Map<id, port> <read only>	Map of ports defined by the user.	port { id = "p1" } port { id = "p2" }
controllerIPv4	IPv4 address	IPv4 Address of the controller node	192.168.1.1

controllerPort	Number	TCP port of the controller node	6969
fabricIPv4	IPv4 address	IPv4 address of the fabric switch	192.168.1.2
fabricSubnetMaskv4	IP mask	IP mask of fabric subnet	255.255.255.0

Default attributes		Comments
OpenFlowFabric { }	ctrlProtocol = { "tcp", "ssl" }	Protocol used by the controller to communicate with the fabric. Default protocol is "tcp".
port { }	Framing = "Eth"	Only Ethernet (802.1Q) framing is available.
	lineRate = { 1, 10 }	Nominal line rate of the port defined in Gbps. Default value is 1 Gbps.
	directionality = { uniDir, biDir }	Directionality of the port.

Example:

```

type OneOFX {
    description = "Two hosts connected to Openflow switch."

    host {
        id="h1"
        location="PRG"
        port { id="port1" }
    }

    host {
        id="h2"
        location="PRG"
        port { id="port1" }
    }

    link {
        id="l1"
    }

    link {
        id="l2"
    }

    ofx {
        id="ofx1"
        location="PRG"
        fabricIPv4="10.10.100.1"
        controllerPort="9966"
        fabricSubnetMaskv4="255.255.255.0"
        controllerIPv4="10.10.100.100"
        port { id="port1"}
        port { id="port2"}
        port {
            id="port3"
            mode="CONTROL"
        }
    }
    adjacency h1.port1, l1.src
    adjacency ofx1.port3, l1.dst
    adjacency h2.port1, l2.src
    adjacency ofx1.port2, l2.dst
}

```

V. External Domain

The External Domain resource represents an endpoint in some facility that is outside the GTS service area. This facility is reachable by a preprovisioned virtual circuit that has been established outside of the scope of a normal GTS testbed life cycle by the GTS service management team.

External domains are represented with the following syntax:

```
ExternalDomain {
    id = "ex1"
    location = "lab1"
    port {id = "ep1"}
}
```

Users may specify the following properties of the external domain:

- *id* – alphanumeric string, case sensitive; The ID is a user chosen name of the External Domain resource (unique within the scope of the Testbed DSL) that can be used within Adjacency specifications to reference a particular ED object/port pair, or for other debugging purposes during the life cycle of the testbed;
- *location* – choice of “AMS”, “BRA”, “LJU”, “PRG”, “LON”, “HAM”, “MIL”; case insensitive. The GTS pop location defines where the External Domain virtual circuit appears on the edge of the GTS domain.
- *Port id* - choice of “ep1”, or “ep2”. The port construct specifies a port ID that is preconfigured within GTS (by the service management team) to be the GTS endpoint for an External Domain virtual circuit. The port id must be either “ep1” or “ep2”. These port ids, at the associated location, are configured to represent different External Domains at different times. The user and the GTS service management team must coordinate to arrange for an external Domain facility to be attached to the GTS domain, and to identify the specific External Domain <location>/<portid> pair which will terminate the virtual circuit going to the particular External Domain. For convenience, the user is able to use the External Domain identifier and the port ID to reference the External domain in adjacency specifications.

To reiterate: It is the “location”/“port id” pair that maps to a particular External Domain at any particular time. And within the testbed DSL description, an External Domain resource instance can be specified and referenced much like a Host resource. Below you can find an example with a host “h1” which is defined to be in Prague that is connected to an external domain called “ProtoGENI-SL” which meets GTS in London and presents a port id “ep1” that other testbed resources can connect to. The host *h1* port *eth1* is connected via link *l1* to external domain *ProtoGENI-SL* port *ep1*.

Example:

```
ExternalExample {
  host {
    id = "h1"
    port {id = "eth1"}
    port {id = "eth2"}
  }

  link {
    id = "l1"
    port {id = "src"}
    port {id = "dst"}
  }

  ExternalDomain {
    id = "ProtoGENI-SL"
    location = "LON"
    port {id = "ep1"}
  }

  adjacency h1.eth1, l1.src
  adjacency ProtoGENI-SL.ep1, l1.dst
}
```

For more information please contact gn3plus-sa2-operations@geant.net.

VI. BNF Grammar

`<property> ::= <key>=<value>`

`<properties> ::= <property>(;|\n)<properties>`

`<port> ::= port { <properties> }`

`<atomic> ::= <atomic-type> { <properties> <ports> }`

`<atomic-type> ::= host | link | ofx`

`<composite> ::= type <name> { <properties> <resources> <ports> <adjacencies> }`

`<resources> ::= <resource> <resources>`

`<resource> ::= <composite> | <atomic>`

`<adjacencies> ::= <adjacency><adjacencies>`

`<adjacency> ::= adjacency <port-id>, <port-id>`

VII. Using the Command Line Interface (CLI)

[activate | deactivate | release] reservation <reservationId> <user> <password>
to activate, deactivate or release a reservation

[add | remove] project <projectName> <user> <password>
to add or remove a project

add user <userName> <password> <role> <projectName> <user> <password>
to add or remove a user to/from a project

remove user <userName> <user> <password>
to remove a user

define <pathToDslFile> <user> <password>
to define a new type from DSL file

query [resources | types | reservations | projects | infrastructures] <user> <password>
to get information on resources, types, reservations, projects or infrastructures

[query | undefine | reserve] type <typeName> <projectName> <user> <password>
to get type definition or to remove or reserve a testbed instance for a given type

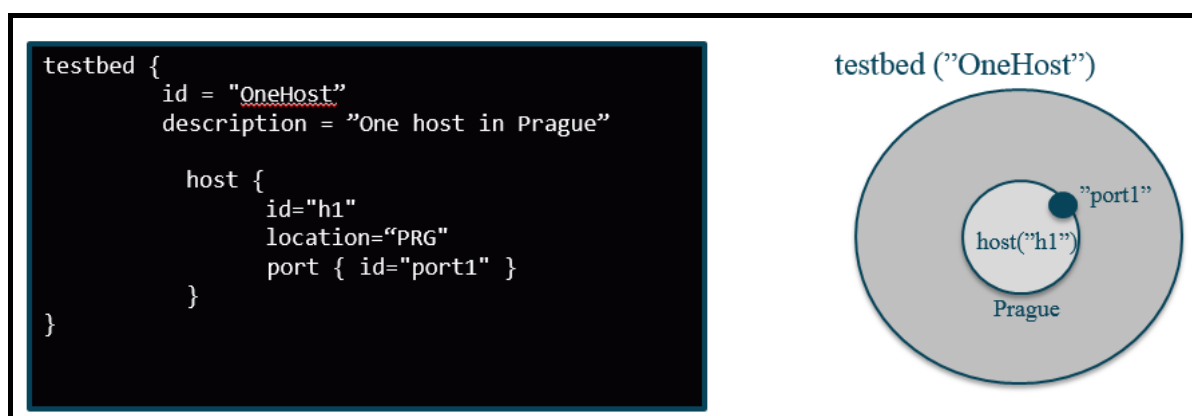
query <ldapQuery> <user> <password>
to get ldap information

VIII. Examples: One host

The following example shows how one host can be described with DSL:



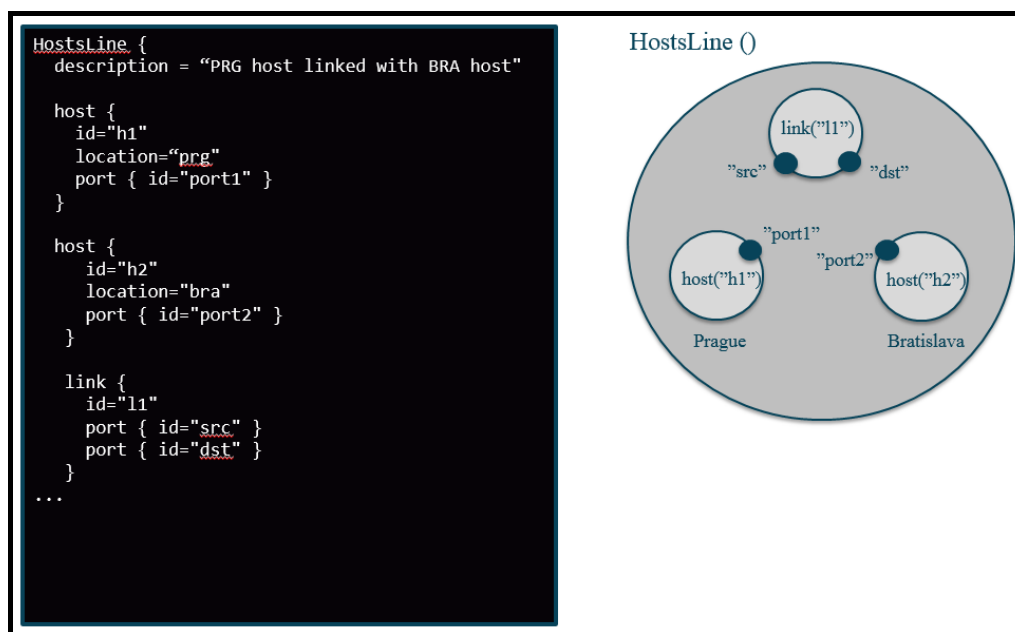
If you need to have one host at a specific location for your experiment then you must specify this in your DSL code:



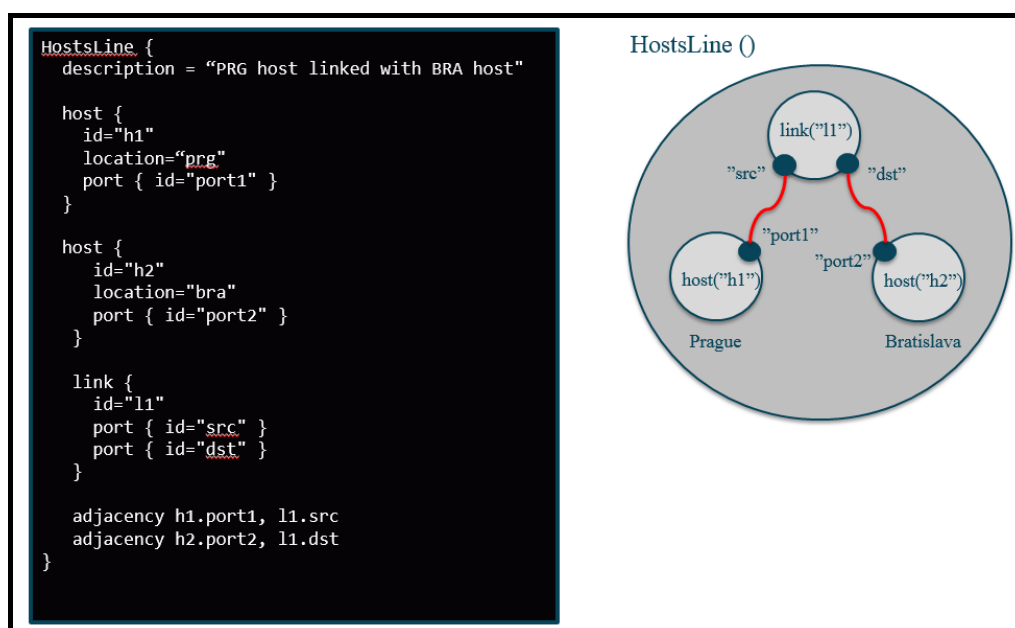
IX. Example: Two hosts linked together

This example shows in two steps how you can define two hosts and a link and connect them to each other:

Step 1: Describe both hosts and the link:



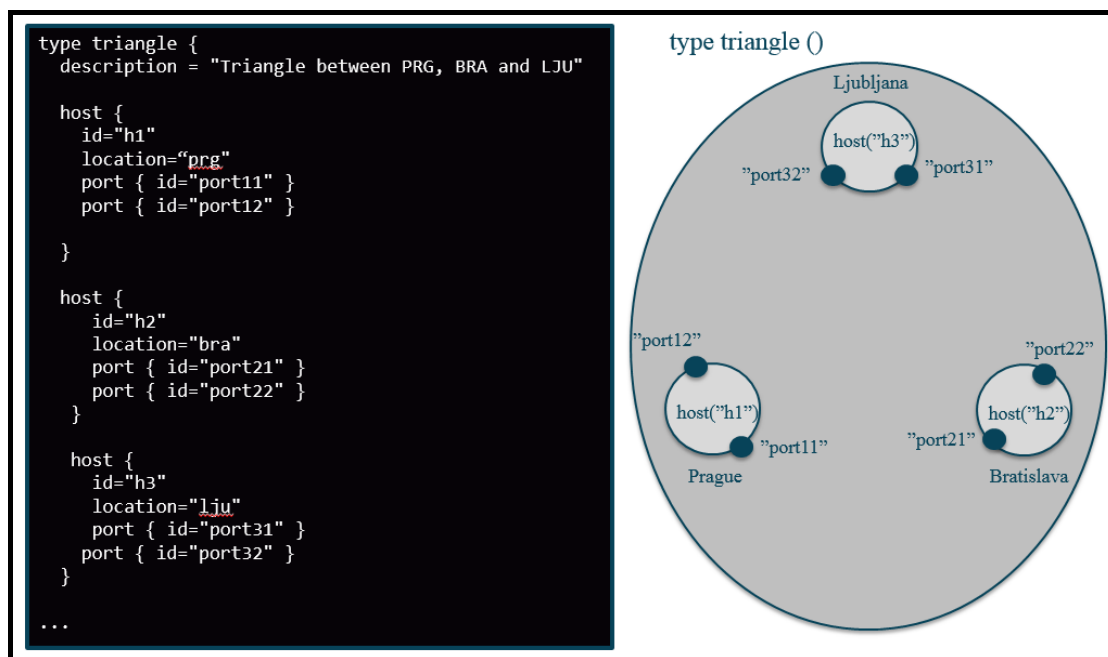
Step 2: Describe the adjacencies to actually connect the entities:



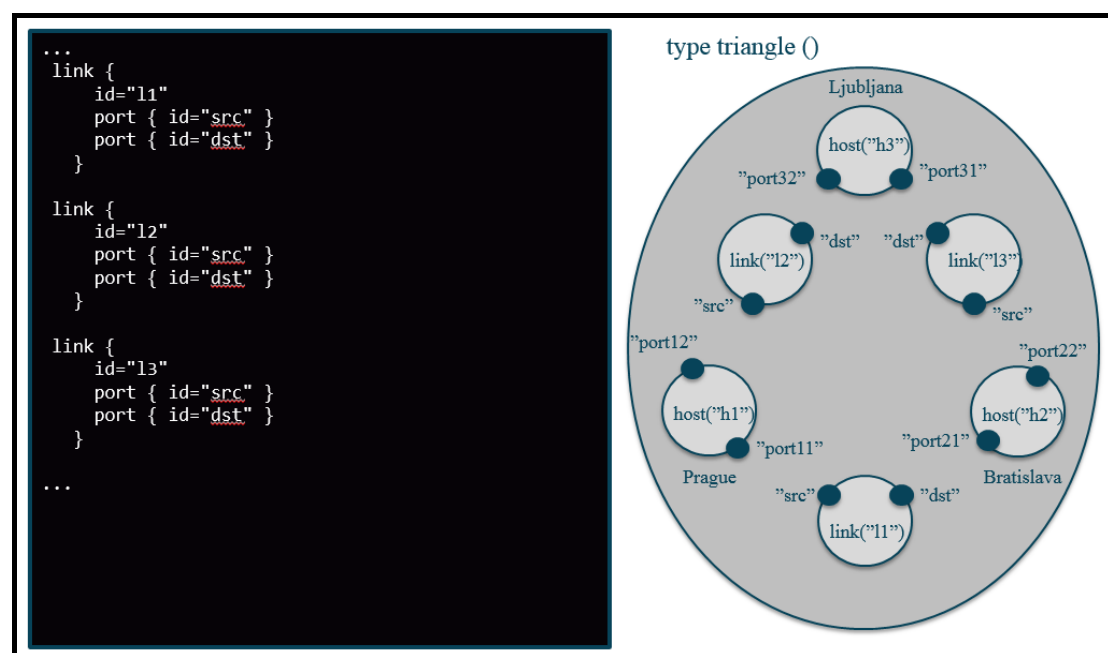
X. Example: Triangle between three locations

For a triangle to be set up between three different locations you need to first describe the three hosts and specify their locations in your DSL code:

Step 1: Describe hosts and locations:



Step 2: Describe the three links between them:



Step 3: Describe the adjacencies to connect them:


```
...
link {
  id="l1"
  port { id="src" }
  port { id="dst" }
}

link {
  id="l2"
  port { id="src" }
  port { id="dst" }
}

link {
  id="l3"
  port { id="src" }
  port { id="dst" }
}

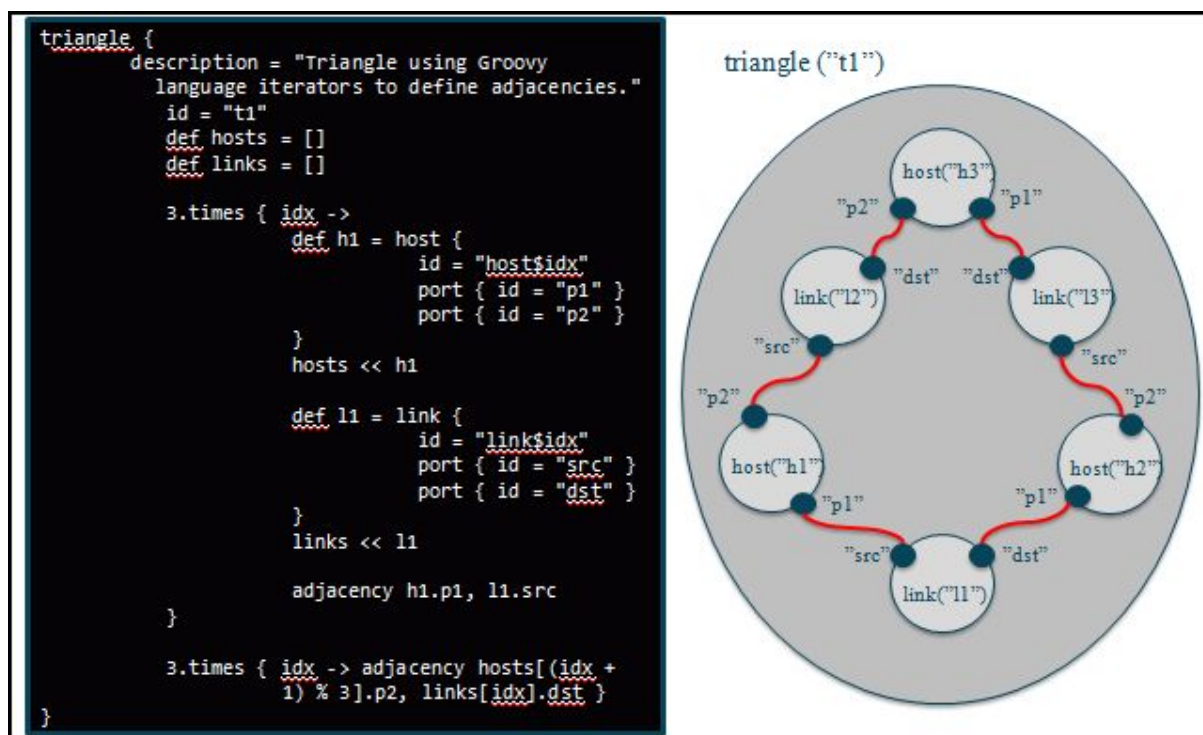
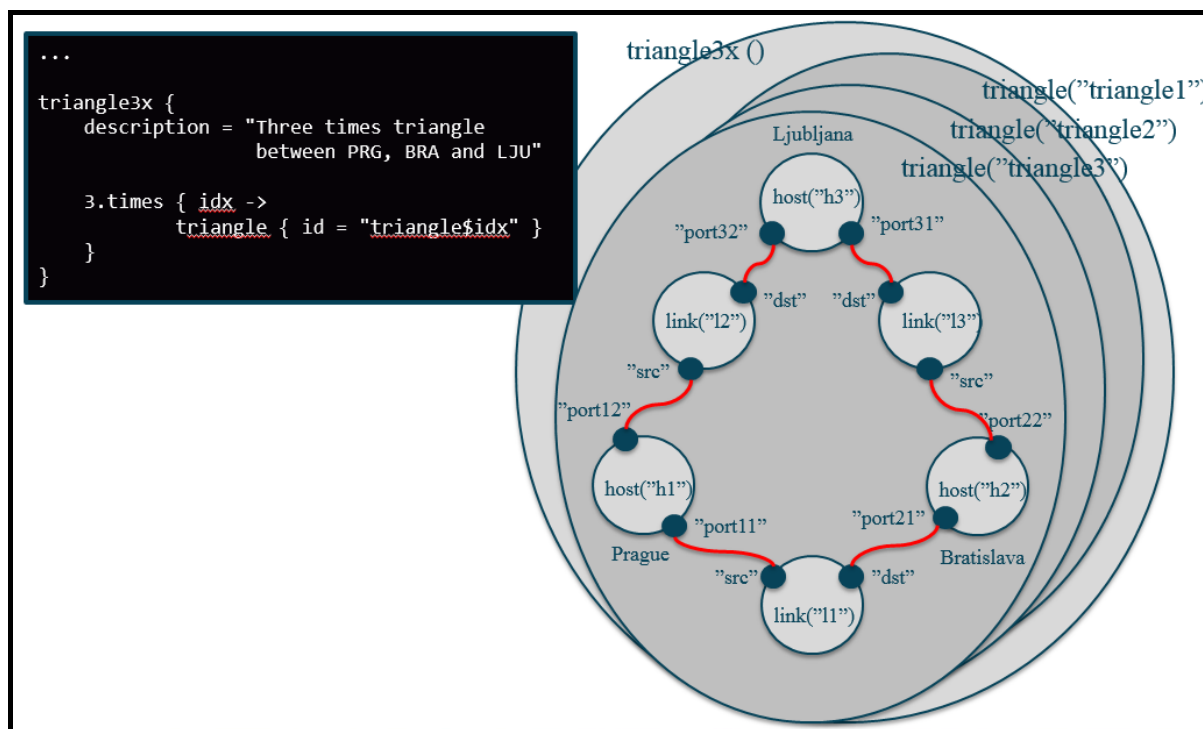
adjacency h1.port11, l1.src
adjacency h2.port21, l1.dst
adjacency h1.port12, l2.src
adjacency h3.port32, l2.dst
adjacency h2.port22, l3.src
adjacency h3.port31, l3.dst
}
```

type triangle ()



XI. Example: Three triangles built using DSL code iterations

The following describes three triangles being setup between Prague, Bratislava and Ljubljana:



XII. Example: Two OpenFlow switches, one controller and three hosts

```
TwoOFXwithController {
  description = "Two OFX with one controller VM. The OFXs are at different locations and
connected."
  id = "TwoOFXwithController"

  host {
    id="controller"
    port { id="eth1" }
    port { id="eth2" }
  }

  host {
    id="dataBRAVM1"
    location="BRA"
    port { id="eth1" }
  }

  ofx {
    id="ofxBRA"
    location="BRA"
    fabricIPv4="10.10.10.1"
    fabricSubnetMaskv4="255.255.255.0"
    controllerPort="9966"
    controllerIPv4="10.10.10.10"
    port { id="p1" }
    port { id="p2" }
    port { id="pLJU" }
    port {
      id="pControl"
      mode="CONTROL"
    }
  }

  link {
    id="l1BRA"
    port { id="src" }
    port { id="dst" }
  }

  link {
    id="l1CBRA"
    port { id="src" }
    port { id="dst" }
  }

  adjacency dataBRAVM1.eth1, l1BRA.src
  adjacency ofxBRA.p1, l1BRA.dst

  adjacency controller.eth1, l1CBRA.src
  adjacency ofxBRA.pControl, l1CBRA.dst

  host {
    id="dataBRAVM2"
    location="BRA"
    port { id="eth1" }
  }
}
```

```

link {
    id="l2BRA"
    port { id="src" }
    port { id="dst" }
}

adjacency dataBRAVM2.eth1, l2BRA.src
adjacency ofxBRA.p2, l2BRA.dst

ofx {
    id="ofxLJU"
    location="LJU"
    fabricIPv4="10.10.11.1"
    fabricSubnetMaskv4="255.255.255.0"
    controllerPort="9966"
    controllerIPv4="10.10.11.10"
    port { id="p1" }
    port { id="pBRA" }
    port {
        id="pControl"
        mode="CONTROL"
    }
}

host {
    id="dataLJUVM1"
    location="LJU"
    port { id="eth1" }
}

link {
    id="l1LJU"
    port { id="src" }
    port { id="dst" }
}

link {
    id="lOFXBRAOFXLJU"
    port { id="src" }
    port { id="dst" }
}

link {
    id="l1CLJU"
    port { id="src" }
    port { id="dst" }
}

adjacency controller.eth2, l1CLJU.src
adjacency ofxLJU.pControl, l1CLJU.dst

adjacency dataLJUVM1.eth1, l1LJU.src
adjacency ofxLJU.p1, l1LJU.dst

adjacency ofxBRA.pLJU, lOFXBRAOFXLJU.src
adjacency ofxLJU.pBRA, lOFXBRAOFXLJU.dst
}

```

References

- DEL-D61 Deliverable D6.1 (DS2.3.1): Architecture Description: Dynamic Virtualised Packet Testbeds Service
- FAR-2014 F. Farina, P. Szegedi, J. Sobieski, GÉANT World Testbed Facility - Federated and distributed Testbeds as a Service facility of GÉANT, published at FIDC 2014 in Karlskrona, Sweden, 2014
- GOO-2014 Google Chrome Browser, <https://www.google.com/chrome/browser/>
- GRO-2014 GROOVY, <http://groovy.codehaus.org/>
- RYU-2014 RYU Project Team, RYU SDN Framework, <http://osrg.github.io/ryu-book/en/Ryubook.pdf>, available on September 29, 2014
- SZE-2014 P. Szegedi, User's Training 1 – The Basics