



GÉANT Testbeds Service

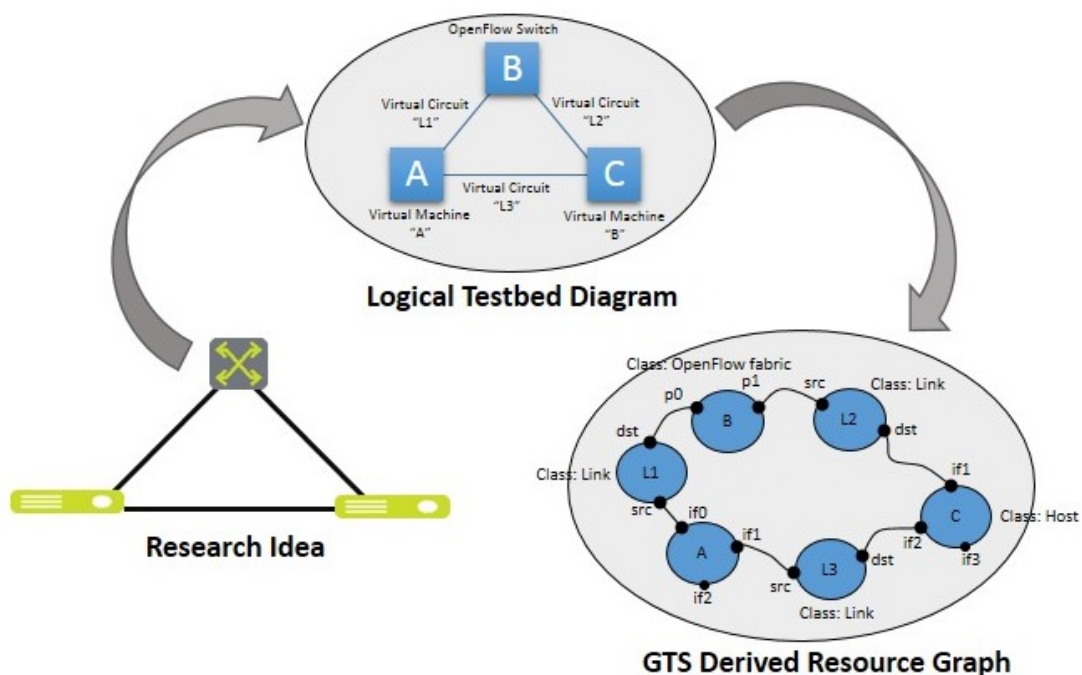
Flexible, Powerful, Tools for Advanced Research



GÉANT Testbeds Service (GTS) v1.0

Overview

The "GEANT Testbeds Service" (GTS) offers user defined experimental networks to the network research community for the purpose of testing novel networking and telecommunications concepts, at scale, and across a geographically realistic European footprint. GTS provides dynamically provisioned network environments consisting of computational servers, data transport circuits, and switching/forwarding elements. These testbeds can be scheduled in advance to coordinate testbed experiments with other non-GTS activities, and testbed components can be selected from any of the GEANT core points of presence that have GTS services available. Each testbed constitutes an isolated and insulated virtual environment that can function autonomously from other testbeds or other production services.



The GTS treats all testbed networks as *graphs*

Figure 1 – GTS network graphs

Testbeds can be defined and orchestrated by the user through a web based Graphical User Interface or via a Command Line Tool within a Linux shell/script environment. The user is

able to access their testbed environment through a standard VPN internet gateway. The testbed itself can communicate with other facilities external to GTS through explicitly defined data flow interfaces incorporated into the testbed data plane itself

The GEANT Testbeds Service is a production capability offered by GEANT. As a production service, the testbed facilities are made available on 24x7 basis. The GTS facilities are continuously monitored by the GEANT Network Operations Center and outages of infrastructure generate immediate alerts to the Level 1 support staff. Level 2 support is available during normal business hours to address unusual or more complex service related problems requiring more expertise. Level 2 support is available outside of normal business hours by special prior arrangement.

GTS version 1.0 is currently deployed in four GEANT core sites (Copenhagen, Bratislava, Ljubljana, and Amsterdam) with three more sites available in 2015-Q1 (Hamburg, Paris, and Milan). These facilities make available approximately 50-100 Virtual Machines, and approximately 20 to 40 virtual Openflow switching fabrics. These components can be linked via 1Gbps virtual circuits provisioned over the GTS core transport infrastructure with approximately 140 Gbps of total backbone capacity.

The virtual network environments offered by GTS provide a revolutionary new capability for both network researchers as well as for developers of advanced networked applications. GTS networks can offer powerful self contained networks for simulations and experimentation, or they can be employed as “software defined exchange” fabrics – intelligent service-infused network domains bonding together multiple otherwise isolated research islands into a contiguous user defined service region.

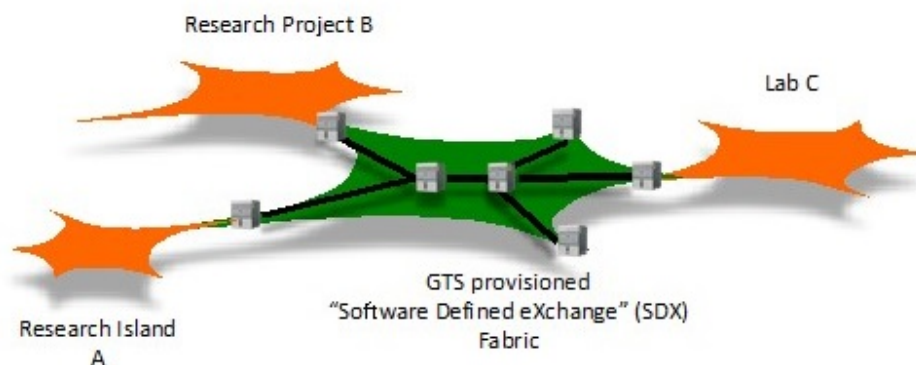


Figure 2 - Software Defined eXchange Fabric

New service features and capabilities, infrastructure, locations, and resource types will be continuously introduced in response to user requirements/

“Testbeds” Purpose and Target User Community

The GTS purpose is to provide a facility to the European network research community for testing novel network architectures quickly and easily. The service will offer computational nodes and switching/forwarding nodes as virtual “resources” that the researcher can reserve and incorporate into their own custom testbed environment. The resources allocated to a testbed will be placed under control of the user researcher/team. GTS provides a generic application program interface (API) that allows the researcher to maintain control of their respective testbed resources and which will allow the service administrative and operations staff to assert management control over the entire service facilities.

The following describes the GEANT Testbeds Service version 1.0.

GTS v1.0 Features

The purpose of a network “testbed” is to provide an environment that has the characteristics of a real world network, yet which is isolated and under sufficient positive control to prevent any unanticipated behavior from causing disruptions to other ongoing production services or that might interfere with other unrelated testbed activities. These testbeds should consists of flexible elements or “resources” that have properties that the researcher can configure to meet their needs.

The target customer base the GEANT Testbeds Service is network research, including applications research related to distributed or networked service environments.

The “service” provides user-defined virtualized network environments that can be applied to explore a wide range of network research topics. Some of the key features are as follows:

1. Testbeds can be described using an easy to understand high level domain specific language – the “DSL”
2. The GTS v1.0 Testbed Resource portfolio includes:
 - 2.1. “Host”s - Linux based virtual machines that provide a highly flexible platform for the deployment of new software based services.
 - 2.2. “Link”s – Point-to-point bidirectional data transport circuits that move Ethernet frames transparently and unmodified end to end. Link resources can be provisioned with up to 1 Gbps transport capacity.
 - 2.3. “OFX”s - Openflow switching fabrics. These are packet forwarding resources that will be available in GTS. The OpenFlow fabric is based on OpenFlow 1.3. Each OFX has its own controller port and a variable number of data plane ports.

2.4. Virtual External Domains – VEDs are virtual resource classes that represent external domains. These resources contain port(s) specifications map to existing external circuits that connect GTS service domain to the external a specific remote domain

3. Testbeds are insulated and isolated from one another. Thus activity within one testbed will not impact other active testbeds.
4. The service allows resources and testbeds to be reserved for use during some window in the future. GTS advanced reservations are mapped to specific infrastructure facilities at reservation time in order to guarantee their availability at the scheduled time.
5. The service provides an API that allows users to develop their own orchestration agents to interact with the GTS core service agent to manage testbeds and resources.
6. The service provides a web based interactive graphical user interface (GUI) to manage a testbed through its entire lifecycle.
7. The service provides a Linux based command line interface (CLI) tool that can be used manually within a shell or within canned scripts to manage testbeds via the service.
8. The service provides classic internet access for users to reach and interact with testbed resources. This allows users to easily stage software or data to/from their testbed resources, or allows those resources to access external servers to download data.
9. The service is a production service. This means that the service is reliable and available with high confidence level 24x7x365. The service infrastructure is monitored, and there is a call center that users can contact to interact with the GTS operations and management personnel.

User Registration and Authorization Policy

The GTS allows new users to register to use the service directly from the login splash page "gts.geant.net." Initially, the service places a very low bar to registration allowing the target research community to easily engage and to begin learning how to use the service. The Service Management team will be responsible for initializing access policy and credentials, and approving more advanced user authorization requests.

GTS supports four user "classes":

- 1) **"Basic"** users. These are users who wish to trial the service to judge its capabilities prior to undertaking large scale work. Such users will be easily registered and will have rights to build and run relatively simple testbeds. These users will be tracked and authorized just as any other user. However, authentication measures will be minimized in order to allow new users to register easily and then quickly become familiar with using the service. Basic users will have substantial limitations on the overall breadth and depth of the testbeds they will be allowed to instantiate, and the testbed data plane will only be allowed very limited external access to other systems or services. This introductory status is intended to provide an easy learning experience for new or prospective users without imposing strict or convoluted security or accounting policy necessary for more advanced, resource intensive, or sensitive projects.
- 2) **"Intermediate"** users. These users will need to provide additional authentication credentials as part of their registration and/or authorization quotas. The ProjectIDs associated with intermediate users will be allowed more privileges in terms of resources of the testbed service. Intermediate user profile allows the user to provision external gateways for the testbed to source or sink traffic across the secure testbed perimeter. Larger testbeds (more resources) will be allowed. It is expected that most network researchers will fall into this category.
- 3) **"Advanced"** users will be allowed customized security profiles such as allowing real traffic to transit a testbed, or to insert custom hardware or software into the resource database. The user policy profiles for advanced users may be highly customized to reflect desired capabilities of the research being performed, the demonstrated expertise of the research team, the specifics of the testbed itself, etc.
- 4) **"Operations"** personnel will have the greatest privileges. Agents with Operations credentials will effectively have full control privileges over resources associated with that function, and will have ability to step in to halt or otherwise manage testbeds that pose problems to the overall continuous availability of the service.

In general, users will be assigned an “allocation” to their account. When this allocation is used up, the user must apply for additional allocations. Each resource will be assigned a “cost” or a rate at which the allocations are consumed over time within the GTS service. This initial accounting process will be reviewed and may be modified as necessary to fit the accounting requirements for GEANT and to meet the community’s needs. Other means of accounting for resource utilization may be proposed and can be considered and/or implemented over the long term.

Once the user is registered, the access to and use of the service is intended to be fully automated – requiring no further human intervention except in error conditions.

There are obvious issues regarding user authorization profiles when we consider the prospect of large, persistent, or multi-domain testbeds. In GTS v1.0, the simple authorization profiles will be applied.

Contact and User Support Processes

The GEANT Testbeds Service website provides the primary interface to prospective users. This website is a first stop/one stop nexus for all aspects of GTS within the context of the GEANT Project. The GTS website provides online training tutorials, information found in this Service Description, and all current documentation – GTS 1.0 Architecture Description, User & Resource Guide, and the Engineering Installation & Configuration Guide. In addition, the GTS provides directions and links to download the GTS software. The GTS website offers a user registration gateway, and trouble ticket system. Phone and email contact information will also be available.

Service Operations and Support

The GEANT Testbeds Service is a production service. This means that the ability to create and use testbeds is expected to be available on a 24 by 7 operational basis. Therefore, the facilities that comprise the service infrastructure will be continuously monitored to insure their availability, the virtual resources assigned to testbeds will be monitored to insure these components of the service are functioning properly, and there will be a means for users to contact the GEANT Operations Center at any time to resolve questions about the operational status of the service or a service instance.

The GEANT Network Operations Center will be tasked with monitoring the service facilities to ensure their continuous operational readiness. Key aspects that will be monitored are:

- **GTS hardware elements.** This includes all GTS “pods” and their constituent hardware components including compute node servers, the OpenFlow switching

- hardware, MX routers, and control plane switch. The hardware elements of the Central Server Facility will also be monitored to include the CSF hardware servers as well as the virtual machines running on those servers.
- **GTS Software agents.** These are processes that run on one of the hardware elements or VMs. These agents include functions such as OpenNSA daemons, OpenStack processes, database agents, shared file services, etc.
 - **User resource instances.** Proper monitoring and evaluation of resource availability depends upon the lifecycle state of a resource and the availability of the underlying infrastructure elements that are used to create the virtual resource instances and support their proper functioning. GTS v1.0 will provide only a very basic monitoring of resources instances to insure they transition properly into the Active state. Later GTS versions will define a comprehensive resource monitoring architecture that considers life cycle state, underlying resource/infrastructure dependencies, and resource specific performance requirements.
 - **The Internet Access Gateway..** The IAGW for each registered GTS Project will be monitored to insure the VPN gateway is operational, and that the key IAGW services (DHCP, NAT, NFS, etc.) are running.
 - **Testbed border I/O activity** Monitoring of external data flow ports is important to prevent testbeds from interfering with other testbeds or other non-GTS services. In GTS v1.0 external data plane interfaces are limited to specific destinations and capacities. The IAGW capacity is likewise constrained to insure adequate yet safe access to/from testbeds.

In addition to real time monitoring, the GTS service logs all significant events that occur in the functioning of the system. These logs can be used for forensic purposes or may be processed for accounting or periodic service reporting purposes.

GTS Support Training.

The Network Operations Center personnel are trained on operational capabilities and problem solving. User training on how to use the Testbed service will be offered to the research community.

GTS Documentation

The GTS provides a number of documents describing various aspects of the service. These can be found by following the documentation links at the GTS page "gts.geant.net". The key documents are described below:

GTS v1.0 Architecture Guide. This is a detailed description of how the GTS service functions. This covers internal functional modules, data structures, and protocols.

GTS v1.0 Engineering, Installation, and Configuration Guide. This document is intended for network engineers and planners. It provide design and engineering requirements, hardware specifications and tradeoffs, and installation and configuration of the GTS open source software

GTS v1.0 User Guide and Resource Catalog. This document describes how users are able to access and manipulate Testbeds within the service. This describes the GUI and CLI interfaces, the API library, and a description of the current resources available to users. The Resource guide provides usage tips as well as a complete description of attributes and other issues associated with resource classes.

Glossary

There has been much work within the R&E community pursuant to virtual networks and services. And often, different terms are used referring to very similar concepts, or the same terms are used for different concepts. This glossary will try to define the use of a number of terms in the context of the GEANT Testbeds Service.

Resource – An allocatable service unit. Resources are created by virtualizing infrastructure components. For instance, a PC server may be virtualized to provide “Virtual Machine” resources – each such virtual machine instance acts as a standalone PC server.

Resource Class – A description of a set of resources that all have the same set of attributes. Each resource instance of this class may have different values for its attributes, but they all share the same attribute set. A class describes the common characteristics of a group of resources.

Resource Instance – An example of a particular resource class description. I.e. a particular resource.

Infrastructure – The underlying hardware, software, or services from which the virtualization process will create virtualized allocatable resources. For GTS, the infrastructure consists of the base facilities underlying the virtualization layer such as PC servers, OpenStack agents, manually provisioned circuits, or specialized hardware such as OpenFlow switches or production routers. For example, infrastructure such as a PC server could be virtualized to provide a bare metal “x86 cpu” Resource. Similarly, a server along with a hypervisor software could be configured to provide set of “Virtual Machine” Resources.

Network – An arrangement of switching/forwarding nodes interconnected by a set of data transport links. The switching and forwarding elements process data flows arriving from the data transport links. Such processing may include conventional IP forwarding, Ethernet switching, or some more complex application oriented processing of the data. Processing of inbound flows may generate outbound flows destined for other nodes. Transport links simply transport information transparently from the outbound port of one forwarding element to an inbound port of another switching/forwarding element. In this model, end systems are also considered nodes in the network, though they simply process the data flows in a different manner than conventional forwarding or switching elements.

Virtual Network - a network (see above) where the transport links and forwarding elements are instantiated from virtualized resources rather than directly upon specific physical hardware infrastructure. The virtual network is described as a graph (nodes and

links) consisting of characteristics of the virtual elements rather than the hardware on which it is may be realized. This allows the [virtual] resource allocation process to analyze the network specification and efficiently map each component of the virtual network to appropriate infrastructure according to the stated requirements of the virtual network description. Such virtualization allows hardware resources to be shared in a manner that is [largely] transparent to the virtual network elements.

Slice – A virtual network. The term slice evolved from network research activities such as Planet Lab where experiments would acquire resources from various distributed facilities. The intercommunicating agents/resources so allocated were referred to as a “slice” of the overall network facilities.

Service – A capability that performs some function on behalf of and at the request of another user, or client, agent. The GTS “Service” therefore creates Testbeds on behalf of and at the request of a user.

Testbed – In the context of the GTS, a “Testbed” is a network specifically created to test and evaluate an experimental concept that involves the transport of data between or across a set of distributed agents. A testbed is considered to be “breakable” – i.e. experiments are able to run on the testbed that may cause a component to fail...And this is ok on a testbed.