

Global Testbeds Positioning

Abstract

In the global efforts to deploy large-scale environments that can host experimental research towards the Future Internet, a number of mainly federative approaches have emerged which are all motivated by the objective to control and manage such a large and diverse environment through a common framework. Although a global agreement is difficult to reach, it is nevertheless worthwhile to analyse the options that exist to enable larger federations and to allow multi-lateral access and utilisation of these precious resources.

Introduction

In the US the GENI initiative currently investigates 5 different approaches for a GENI control framework (CF). The figure below illustrates the relationships of the different control frameworks with existing frameworks, like Emulab, PlanetLab, and OMF, as well as the relationships of European Initiatives like Onelab, PlanetLab Europe and G-Lab to PlanetLab. The figure depicts the current status of the CFs and the technologies used for their implementation.

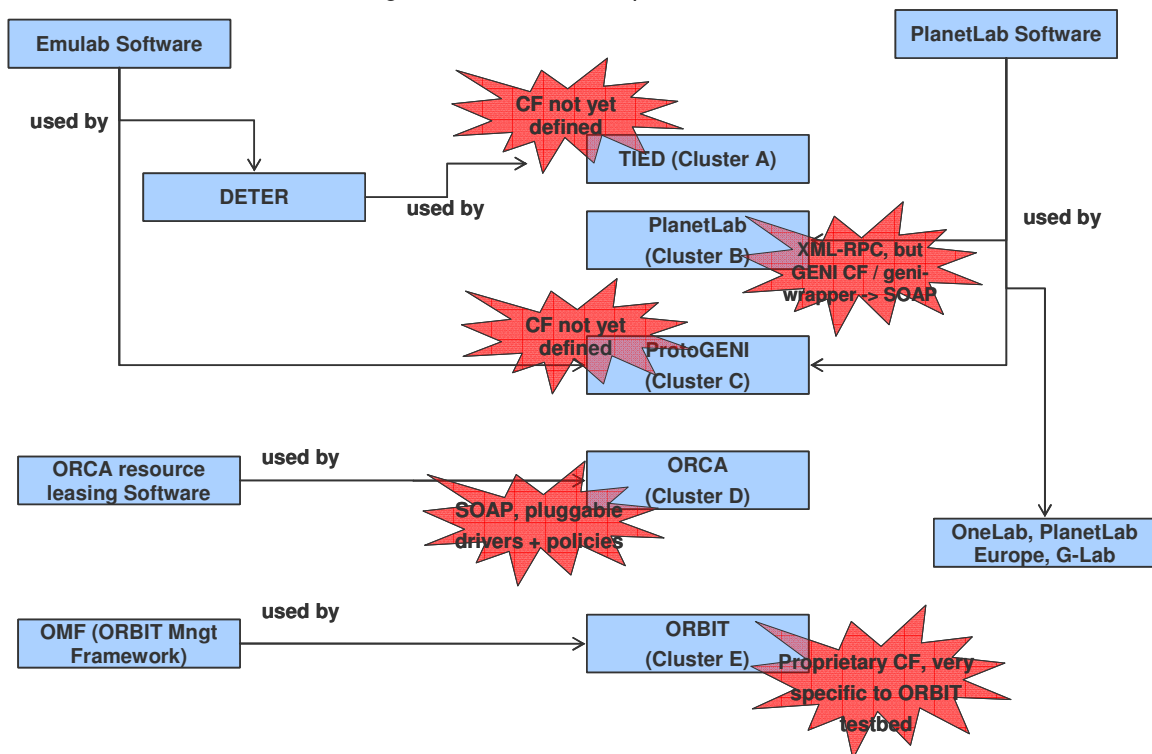


Figure 1 GENI control framework

In Europe the main approaches are centred on 3 projects, namely PII (based on the Panlab framework), Onelab2, and Federica. The remainder of this paper will focus on the work of PII and will compare it to the work of GENI.

Taxonomy – Categorisation

In order to compare Panlab and GENI we introduce the following taxonomy for the areas of concern:

- Central coordinating intelligence
- Resource providers

- Resource consumers
- Resource users
- Central office

As illustrated in the figure below each of the initiatives have partitioned the areas of concern similarly, although different terminology is used for communicating the concern. Concerning the resources that may constitute the respective environments, Panlab includes hardware, software and services, which may be virtualized or physical, and is cross-cutting across all layers, namely network, service platforms, and applications. In contrast the GENI resources are mainly networking resources, whereas slices are composed of slivers and can span across several aggregates (term used for a GENI resource provider). A sliver is a part of a slice that is located inside a single aggregate. Virtualization and the so called Slice-based Facility Architecture are the driving forces in the GENI design.

	Central coordinating intelligence	Resource Providers	Resource Consumers	Resource Users	Central Office
PANLAB	TEAGLE	Panlab Partner	Panlab Customer	Panlab End User	Panlab Office
GENI	GENI Clearinghouse / GENI Admin and Opa Org	GENI aggregates	Research Organization / principal	GENI End user	GENI office

PANLAB Resources	Hardware, software/services, virtualized or not, across all layers
GENI Resources	Networking resources, Slices across several aggregates and “slivers”, virtualization is driving force

Figure 2 Taxonomy – Categorisation

High Level Requirements

Further on analysing the high level requirements that drive the Panlab and GENI frameworks we also discover several similarities. These are listed and compared in table 1 below:

Table 1 High Level Requirements

	PANLAB	GENI
TEAGLE / Clearinghouse	<ul style="list-style-type: none"> • Federation of resources, currently NO federation of suites! • Resource Discovery / Registration • Cross-Domain provisioning of resources • Resource Description (Services and Hardware!) • Service/Resource Composition • Reservation of resources 	<ul style="list-style-type: none"> • Federation of aggregates, components and suites • Secure Operation
Panlab Partner / GENI aggregates	<ul style="list-style-type: none"> • Identification • Registration • Resource allocation • Resource description • Protect resources -> security 	<ul style="list-style-type: none"> • Identification • Registration • Resource allocation • Resource description

	PANLAB	GENI
	<ul style="list-style-type: none"> • Easy to join federation • Incentive to join federation • Quality assurance (SLA conformance) 	
Panlab Customer / GENI principal	<ul style="list-style-type: none"> • Identification • Registration • Authentication • SLAs • Quality assurance • Roles / policies • Search • Export Federation Functionality • Book testbeds • Book experts • Need for test tools • Compare results with previous federated experiments • Control virtual customer testbed (VCT) (e.g. start, stop, reset) • Provision VCT (e.g. add, delete, modify) 	<ul style="list-style-type: none"> • Identification • Registration • Authentication • Privileges / roles • Resource and topology discovery • Resource to resource connections • Setup verification • Experiment and sliver control • Forensic and usage data collection • Experiment status events and notifications • Experiment status commands
Panlab / GENI End user	<ul style="list-style-type: none"> • Experiment data collection and management 	<ul style="list-style-type: none"> • Experiment data collection and management
Panlab / GENI resources	<ul style="list-style-type: none"> • Identification • Registration • Description/modeling • PTM adaptor needed • Authorization • Resource assignment (booking) • Root/configuration access (not for all resources!) = component programming • Quality assurance (SLA conformance) 	<ul style="list-style-type: none"> • Identification • Registration • Resource sharing (virtualization or other means) • Description • Authorization and policies • Resource assignment • Component programming • Disconnected operation • Resource to resource connection

The main differences between Panlab and GENI with respect to the target features of the environments that can be identified are largely allocated in the following areas of concern:

- **Exclusive booking of resources:** if supported by the respective resources, Panlab provides exclusive control over these booked resources in a virtual customer testbed (VCT). In contrast in GENI this is not foreseen. This relates to resource management and is a result of virtualization being the main technology in the GENI environment.
- **Federation of suites:** GENI defines itself as a novel suite of infrastructure to support experimental research in network science and engineering. There can be multiple suites which can federate under multiple GENI clearinghouses. Panlab currently defines one level of federation, which is controlled by Teagle. The option that there may exist several instances of Teagle which federate at a higher level has been discussed in Panlab, but has not been analyzed in-depth yet.

- **Available / targeted resources:** The type of resources that are made available in Panlab can be of any type as long as they can be described (see below) by the Panlab adopted formalism. As such they can be physical or virtual resources at the network, platform or application level. GENI focuses clearly on network capabilities.
- **SLAs:** Panlab considers the inclusion of a Service Level Agreement (SLA) as part of the contract among the involved stakeholders, i.e. between the Panlab office the testbed providers, and the customers of the services. SLAs are currently not addressed by GENI. There is, however, the notion of setup verification (CF requirement) to ensure the proper instantiation of requested resources.
- **Economics:** Panlab considers an economic model based on revenue. The economic model is deemed necessary in Panlab in order for an economically self-sustainable federation to emerge. Public funding required for Panlab operations shall be reduced to a minimum and reach zero in the long term. In contrast GENI is an infrastructure that can be considered a “public good” inside the GENI community based on resource sharing and fair use principles.
- **Emergency stop:** The GENI environment foresees a central emergency stop “switch”, which is able to suspend all activities in the GENI infrastructure. Such a “switch” is currently not foreseen in Panlab, although technically certainly possible. However, this would involve complex technologies such as anomaly detection and prevention.
- **Component programming:** Within given constraints usually set by usage policies, GENI allows full programmability of the components in the environment. In contrast in Panlab the programmability of a configured testbed, called Virtual Customer Testbed (VCT), is largely taking place before provisioning while during the execution phase only configuration access is granted to the customer of the VCT.
- **Resource description:** Panlab uses a generic, model based description technique, based on DEN-ng, to describe the available resources. In GENI the resource description is currently unclear. The current development suggests that GENI resources will be described according to RSpec which is an XML schema based description.

The figure below summarises the differences in a table.

	PANLAB	GENI
Exclusive booking of resources	✓	✗
Federation of suites	✗ (currently)	✓
Available / targeted resources	any resource / service	mainly network capabilities
SLAs	✓	✗
Economics	revenue based	public good
Emergency stop	✗ (currently)	✓
Component programming	configuration access	root access
Resource description	generic, model based	probably RSpec, not clear yet

Figure 3 Difference in requirements

What is TEAGLE?

For federating different testbeds residing in autonomous administrative domains and provide composite infrastructures using resources across the boundaries of the domains, three major blocks of functionality need to be provided:

- Resource Description,

- Service Exposure, and
- Service Composition.

These functionalities are considered the enablers for the federation, since they represent essential capabilities of the technical infrastructure of the federation. Teagle is a concept in Panlab that subsumes the functionalities listed above.

Teagle federates any type of resources as long as such resources are uniformly described (see next section). In this respect Teagle is resource technology agnostic and only relies on the fact that a given resource can be managed via a Web Services interface.

This is a challenge for some resources such as highly proprietary network components, since a configuration command issued by Teagle to the component has to be mapped onto a resource specific configuration command.

The Panlab architecture is generic enough, and as long as the above conditions are met, is able to support all types of different resources.

Teagle as a broker for heterogeneous resources

Whatever is offered by a testbed towards the Panlab federation must be described according to a common information model in order to allow Teagle to understand the functionality of resources, offer them towards external customers and provide compositions of distributed functionalities. The services need to be described uniformly to allow Teagle as a broker residing between the customer and the resource, to match requests to concrete functionalities.

Resources are exposed as services via a component that sits at the border of a testbed (representing an administrative domain) and which is called Panlab Testbed manager (PTM). The PTM provides the mapping of federation level communication messages to resource specific management and configuration commands. The PTM has a “device-driver”-like interface that allows adaptors for specific testbed resources to plug into the PTM. Resources in turn are controlled via standard interfaces (e.g. SNMP) or proprietary interfaces. In this case the resource adaptors for the each proprietary interface must be implemented, normally by the resource owner.

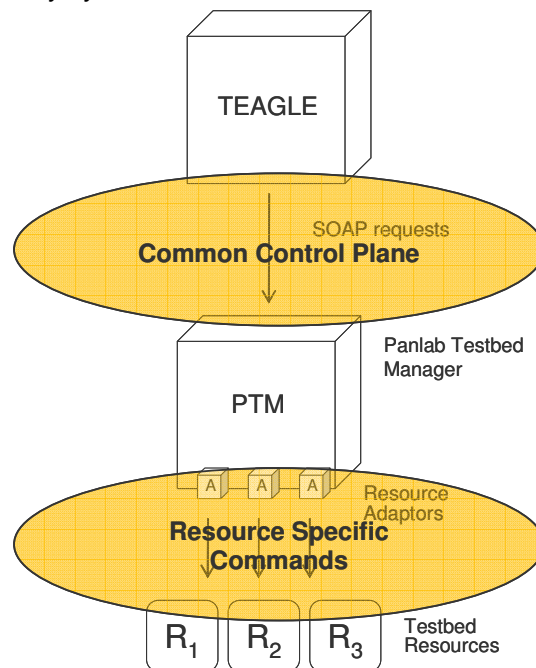


Figure 4 Panlab Testbed Manager

Teagle offers service composition functionality that allows the orchestration of (atomic) building blocks offered by different testbeds. In order to do so, it requires a solid description of what is available. A repository holds descriptions of the available resources in the federation and instructions on how to invoke them.

Teagle provides a frontend portal that allows a customer to browse the available federation resources and help him to express the desired functionality in a Virtual Customer Testbed (VCT) description. Once the customer has provided the VCT description, a composition engine provides a provisioning workflow that defines how to provision the requested environment as a composition of building blocks from different testbeds. The field of service orchestration is still subject to extensive research, however

Panlab endeavours using current state of the art technologies in this area. First implementation stages are foreseen to rely on a number of manual processes while fully automated service composition remains the grand vision to be achieved in later implementation stages.

How can Teagle broker Federica Resources?

Federica uses the concept of slices for provisioning compositions of virtual computing and communication resources. The customer interacts with a Federica network operations center (NOC) and can request a slice. To ease the choice and configuration option for the customer, Federica is planning to present a choice of so called services templates. A service template is a configuration of a slice that is known to work.

A Federica slice is generated when requested by the customer and does not exist unless requested and subsequently generated by the NOC.

Certainly the resource description of a Federica service template must be available in the Panlab repository of service offerings. When this resource description is available, Teagle can broker Federica resources, concretely slices based on service templates. When requested by the Panlab customer, Teagle can then issue a request for a slice to the Domain Manager at the Federica NOC. This Domain Manager maps to the PTM (in Panlab nomenclature) and must be implemented. Once Federica offers an API to programmatically request of a slice the implementation of the Domain Manager should not be too difficult. What is unclear currently is what formalism Federica will use to describe resources.

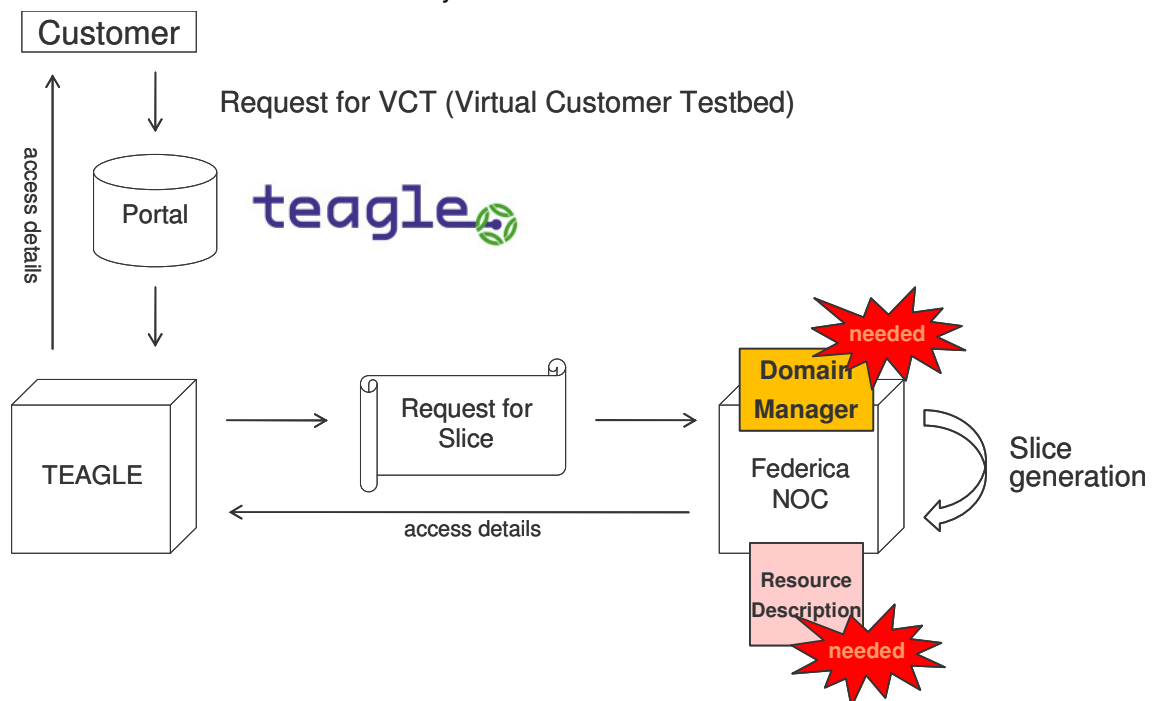


Figure 5 Example Federica

How can Teagle broker PlanetLab Resources?

PlanetLab defines its resources based on RSpec which is an XML schema based description. For inclusion of PlanetLab resources into the Teagle repository a mapping of RSpec descriptions to the DEN-ng information model must be developed. Once this mapping would be available PlanetLab resources can be added to the Teagle repository and brokered as any other resource. RSpec is currently under investigation as the formalism for describing GENI resources.

A Panlab wrapper could be developed that maps Teagle commands to the necessary commands to manage PlanetLab slivers and slices. In principle this should be possible, since a prototype implementation of the GENI interfaces to PlanetLab functionality is now available. The module, called **geniwrapper**, is available at <http://svn.planet-lab.org/>. Although currently PII has no implementation experience in developing a Panlabwrapper for PlanetLab, this should be possible following the same implementation principles as the implementers followed for the **geniwrapper**.

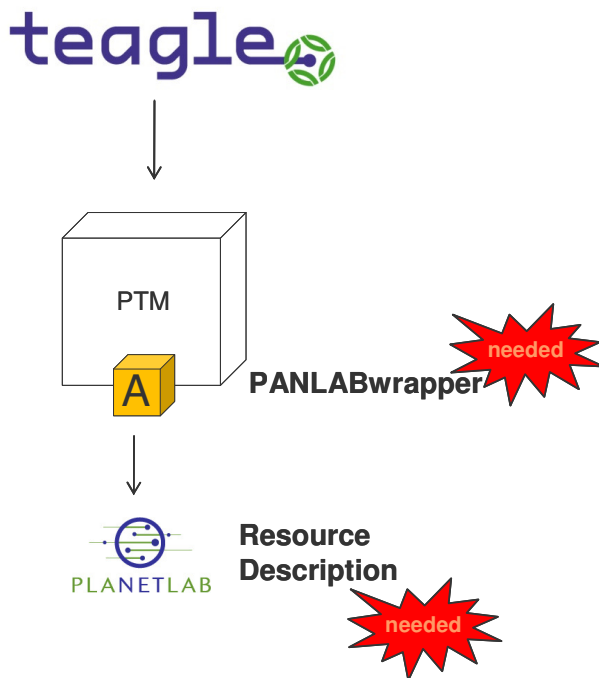


Figure 6 Example PlanetLab

Teagle as a member of the GENI clearinghouse federation

GENI uses the concept of a clearinghouse with which a researcher (customer) and all resources are registered and have established identities. The clearinghouse is also responsible to grant authorisation credentials to the researcher allowing her to create slices.

GENI foresees a growth path to international, semi-private, and commercial “GENIs”, by allowing other clearinghouses to federate with the NSF GENI clearinghouse. In such a scenario Teagle could be acting as a federated GENI clearinghouse being able to broker GENI resources, but also able to offer resources to GENI experiments.

The case in which Teagle brokers GENI resources is very similar to the two previous examples (Federica and PlanetLab). The case that Panlab is offering resources towards GENI has not been examined yet. In most cases this implies that a researcher in GENI should interrogate Teagle and configure a VCT as any other Panlab customer. Since GENI is largely based on virtualisation techniques, it would imply that a “breakout” component is implemented that bridges GENI to Panlab.

The GENI control framework architecture is currently under development, likewise the Panlab control framework, so that there exists a good opportunity to align the future developments.

