G E N I

Global Environment for Network Innovations

GENI Structure Overview

Document ID: GENI-SE-SY-STRO-01.1

June 15, 2009

Prepared by:
The GENI Project Office
BBN Technologies
10 Moulton Street
Cambridge, MA 02138 USA

Issued under NSF Cooperative Agreement CNS-0737890

TABLE OF CONTENTS

1 Document Scope

This section describes this document's purpose, its context within the overall GENI document tree, the set of related documents, and this document's revision history.

1.1 Purpose of this Document

This document provides an overview of the GENI suite structure.

It is a DRAFT, to be used for discussion..

1.2 Context for this Document

Figure 1-1. below shows the context for this document within GENI's overall document tree.

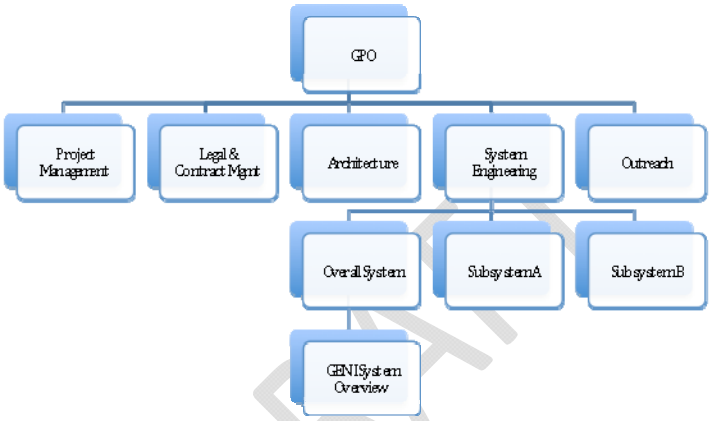

Figure 1-1. This Document within the GENI Document Tree.

1.3 Related Documents

The following documents of exact date listed are related to this document, and provide background information, requirements, etc., that are important for this document.

1.3.1 National Science Foundation (NSF) Documents

Document ID	Document Title and Issue Date
N / A	

1.3.2 GENI Documents

Document ID	Document Title and Issue Date
GENI-SE-SY-RQ-01.9	GENI System Requirements (DRAFT) January 16, 2009 http://groups.geni.net/geni/wiki/SysReqDoc
GENI-SE-SY-SO-02.0	GENI System Overview September 29, 2008 http://www.geni.net/docs/GENISysOvrvw092908.pdf
GENI-INF-PRO-S1-OV-1.12	GENI Spiral 1 Overview September 29, 2009 http://www.geni.net/docs/GENIS1Ovrvw092908.pdf
GENI-SE-CF-RQ-01.3	GENI Control Framework Requirements (DRAFT) January 13, 2009 http://groups.geni.net/geni/wiki/GeniControlFrameworkRequirements
GENI-SE-CF-PLGO-01.2	PlanetLab GENI Control Framework Overview (DRAFT) January 14, 2009 http://groups.geni.net/geni/wiki/PlanetLabGeniControlFrameworkOverview
GENI-SE-CF-PRGO-01.4	ProtoGENI Control Framework Overview (DRAFT) March 25, 2009 http://groups.geni.net/geni/wiki/ProtoGeniControlFrameworkOverview
GENI-SE-CF-ORGO-01.2	ORCA GENI Control Framework Overview (DRAFT) January 14, 2009 http://groups.geni.net/geni/wiki/OrcaGeniControlFrameworkOverview
GENI-SE-SY-TS-UC-LC-01.1	Lifecycle of a GENI Experiment (DRAFT) March 4, 2009 http://groups.geni.net/geni/wiki/ExperimentLifecycleDocument

1.3.3 Standards Documents

Document ID	Document Title and Issue Date
N / A	

1.3.4 Other Documents

Document ID	Document Title and Issue Date

1.4 Document Revision History

The following table provides the revision history for this document, summarizing the date at which it was revised, who revised it, and a brief summary of the changes. This list is maintained in reverse chronological order so the newest revision comes first in the list.

Revision	Date	Revised By	Summary of Changes
-01.1	6/15/09	H. Mussman	Early draft, for discussion within GPO.
-01.2			
-01.3			
-01.4			

2 Introduction

The Global Environment for Network Innovations (GENI) is a suite of experimental network research infrastructure now being planned and prototyped.

A GENI suite will have multiple entities and actors, arranged in a structure that is determined by agreements among the entities that are made by certain actors, the "authorities".

It must also be noted that multiple GENI suites may exist that are somehow related or "federated" with each other. This will create a more complicated GENI multi-suite structure.

The expected range of GENI structures needs to be supported by the GENI architecture, and particularly by the control framework.

Certain functions within GENI, such as end-user opt-in, will also drive the need for particular agreements among the entities and actors.

3 GENI Entities and Actors

3.1 Aggregates that Provide Resources

A GENI suite has these entities that provide resources for experiments within the GENI suite:

- Multiple aggregates (or aggregate components), which are physical entities that provide resources to authorized users for experiments.

Associated with each aggregate are these actors:

- A management authority that is responsible for the management of the aggregate, and which can delegate selected functions to other actors. (Should this be renamed "aggregate authority"?)
- An aggregate administrator is an actor who has been delegated the responsibility for assigning the resources provided by that aggregate.
- An aggregate operator is an actor who has been delegated the responsibility for maintaining the operational readiness of that aggregate.
- The management authority can enter into agreements for the aggregate.
- An aggregate may often be provided by a research institution, e.g., a research group at a university; then, that research group becomes the "management authority", perhaps in combination with other university officials and/or groups.

3.2 Slices that Utilize Resources

A GENI suite has these entities that provide utilize resources for experiments within the GENI suite:

- Multiple slices (abstract entities) that "contain" resources to support experiments (abstract entities) within the GENI suite.
- Each slice contains from zero to many resources (named "slivers"), obtained from zero to many aggregates. See Figure 3-1.

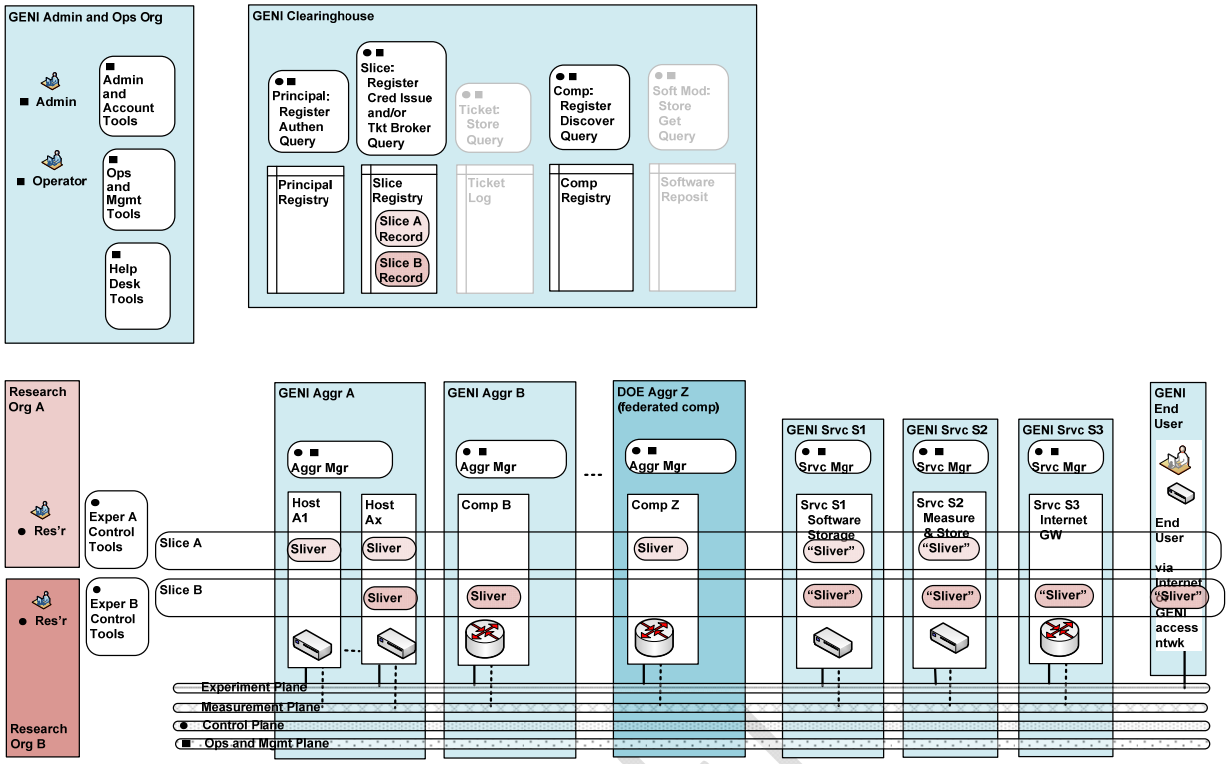

Figure 3-1. Two GENI Slices.

Associated with each slice are these actors:

- A slice authority that is responsible for the management of the slice, and the experiments conducted with the resources assigned to that slice, and which can delegate selected functions to other actors.
- A slice administrator is an actor who has been delegated the responsibility for getting resources assigned to a slice.
- A PI is an actor who has been delegated the responsibility for allowing certain researchers to request that resources be assigned to the slice, and then utilize the resources that have been assigned to the slice.
- A researcher is an actor who has been allowed to request that resources be assigned to the slice, and then utilize the resources that have been assigned to the slice to conduct an experiment.
- The slice authority can enter into agreements for the slice.
- A slice is almost always utilized by researchers at a research institution, e.g., a research group at a university; then that research group becomes the "slice authority" , perhaps in combination with other university officials and/or groups.

It is expected that most research institutions will both provide and utilize resources, and thus they will have both a management and a slice authority. This can be considered the "reference case", with exceptions for research institutions that only provide or utilize resources, but not both.

3.3 Opt-In and 3rd-Party Users

A GENI suite may optionally include these users:

- Opt-In users who choose to "opt-in" to a GENI experiment within the GENI suite.
- Opt-In users who are pulled into a GENI experiment within the GENI suite as a group, e.g., because a site decides to participate in an experiment or service.
- 3rd-party users, who are outside the GENI suite, but may be undesirably affected by a GENI experiment from within the GENI suite (this may be considered a case of "unintended" opt-in by such users)

4 Necessary Agreements

An underlying set of agreements is necessary between these principal actors in a GENI suite, if the GENI suite is to provide useful service to all of the associated research institutions:

- The management authorities of all aggregates, that are responsible for the management of all aggregates and for maintaining the operational readiness of those aggregates.
- A slice authorities of all slices, that are responsible for the management of all slices, and the experiments conducted with the resources assigned to those slices.

These agreements must include:

- A resource usage agreement, that would typically allow resources from any (or all) of the aggregates to be used in any (or all) of the slices, and the rules (or policies) for assigning resources, particularly any exceptions.
- Cooperative operations agreements.
- Acceptable use agreements, and other agreements on best practices for managing the suite and conducting experiments.

5 GENI Structure Options

There are many potential structures depending on the underlying set of agreements between the management and slice authorities in a GENI suite.

5.1 One-level Mesh

The simplest arrangement would be a mesh of agreements between the management and slice authorities; see Figure 5-1. However, this is not scalable, and it is not expected to be used in many cases.

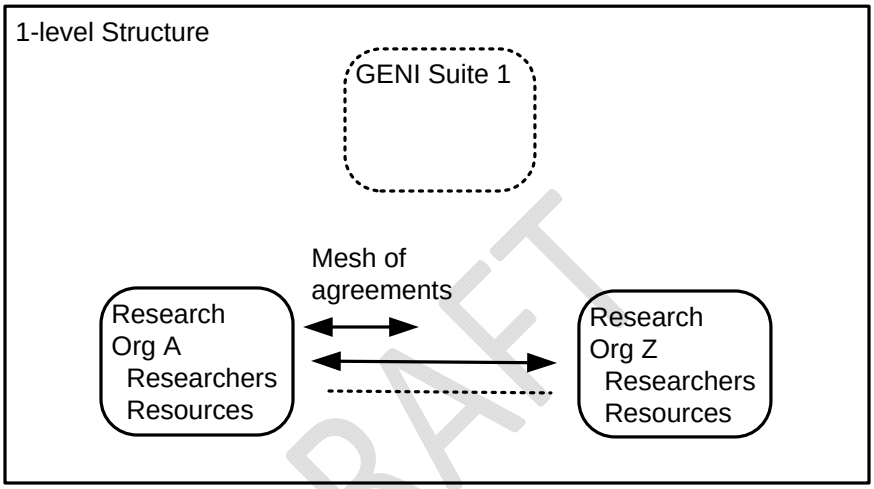

Figure 5-1. One-level Mesh.

5.2 Two-level Hierarchy

A straightforward alternative is to introduce a GENI suite authority, who in turn has agreements with all of the management and site authorities such that:

- All management authorities agree to follow a resource use policy that makes resources available to all slice authorities (all slices) in the GENI suite.
- All actors agree on cooperative operation policies.
- All actors agree to follow a common set of best practices.

See Figure 5-2.

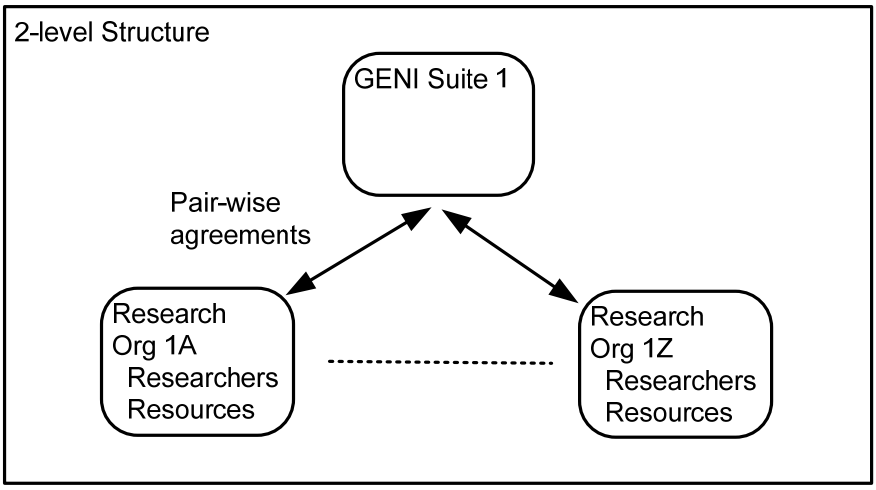

Figure 5-2. Two-level Hierarchy.

This is the choice of several similar projects, including PlanetLab; see Figure 5-3 and Section 6-1.

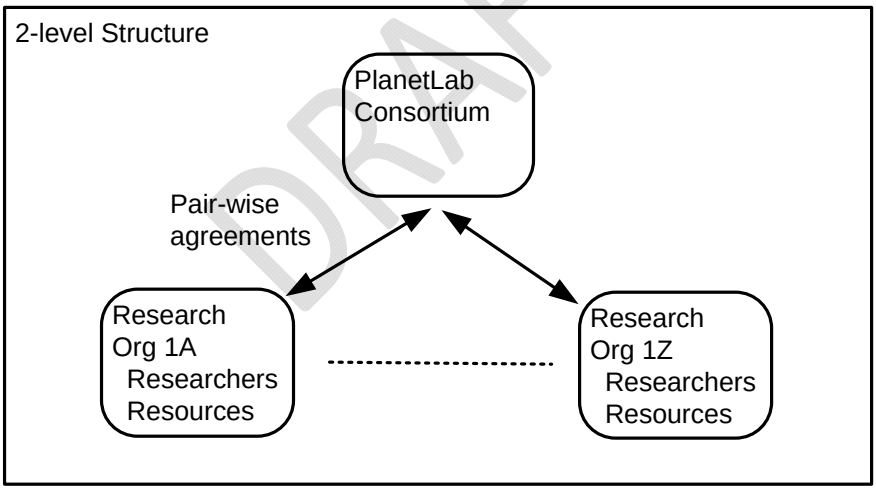

Figure 5-3. Basic PlanetLab Structure.

This is the choice that has often been assumed in defining GENI. See the GENI System Overview document at http://www.geni.net/docs/GENISysOvrvw092908.pdf for more details, which includes a system diagram, repeated here as Figure 5-4.

In this case, a "GENI clearinghouse" is the entity associated with the GENI suite authority.

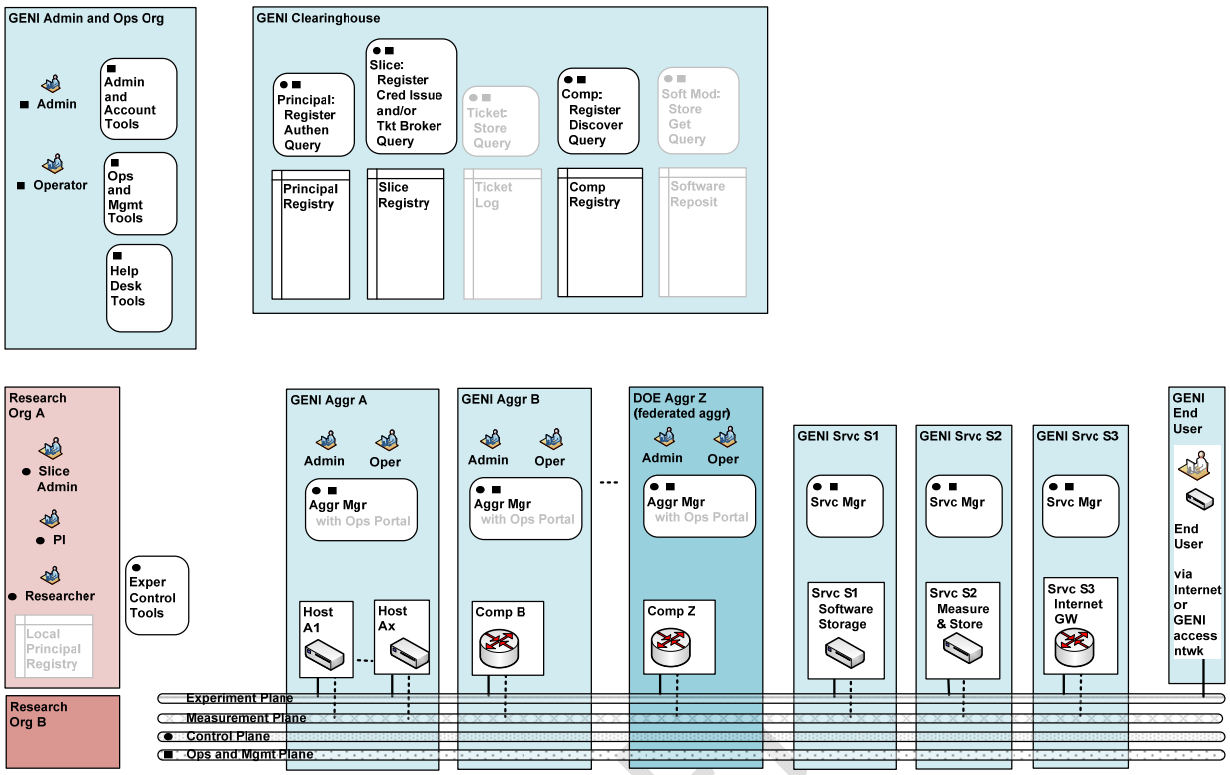

Figure 5-4. GENI System Diagram.

5.3 Two-level Hierarchy with Federation

There will also be situations where two (or more) GENI suites are federated as peers; see for example Figure 5-5.

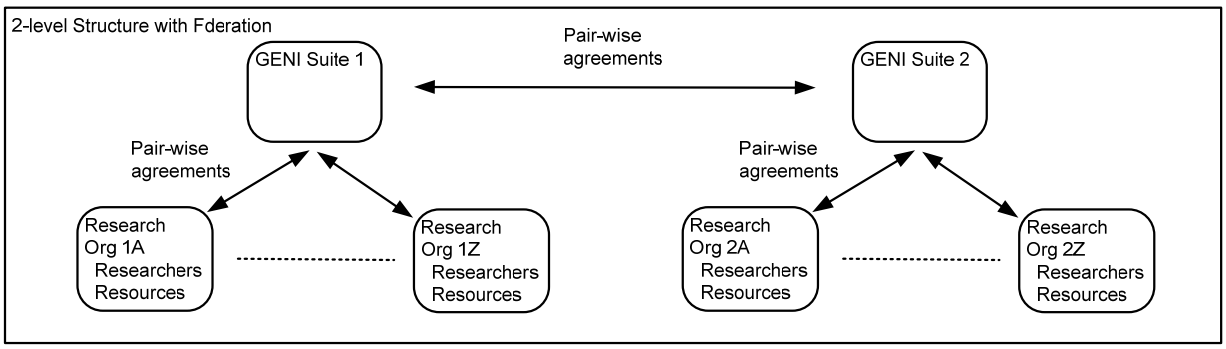

Figure 5-5. Two Federated GENI Suites.

This structure must be supported by the GENI architecture, and particularly by the control framework.

This is also included in the current GENI system overview, which shows how two GENI suites can be federated.

PlanetLab is just beginning to explore this situation; see Figure 3.6 in http://www.cs.princeton.edu/~llp/geniwrapper.pdf .

5.4 Three-level Hierarchy with Local Suites

There will also be situations where one (or more) "local" GENI suites are affiliated into a larger GENI suite as a peer to that suite's other management and slice authorities; see for example Figure 5-6.

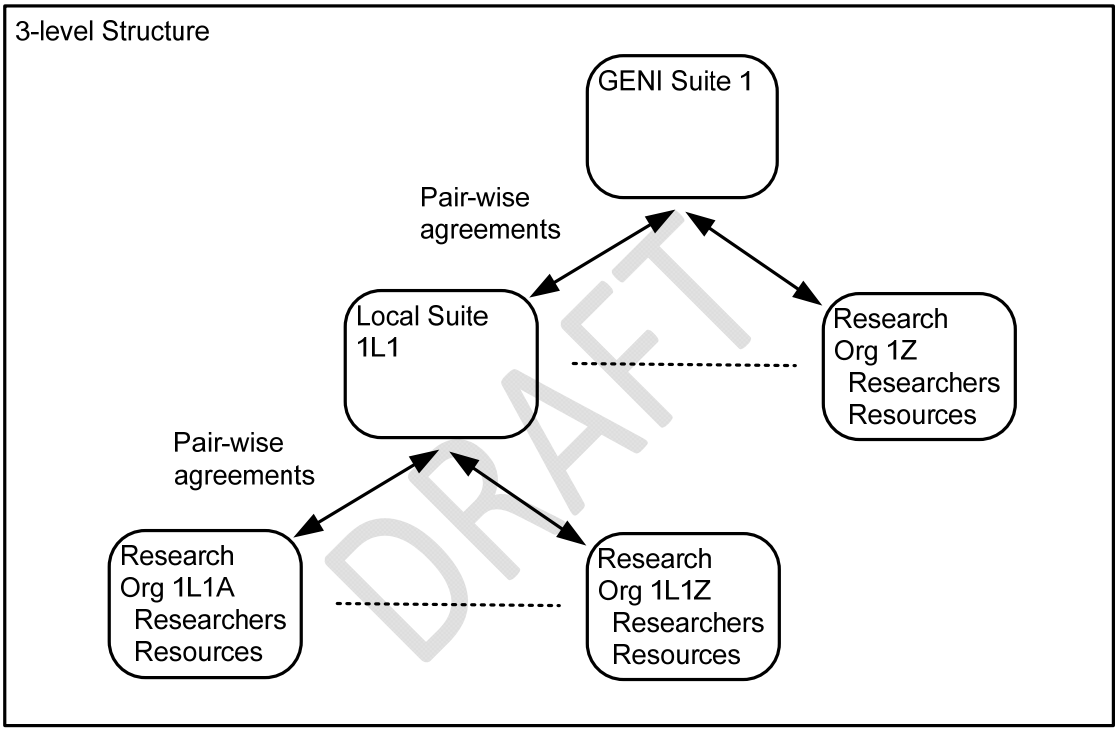

Figure 5-6. Two GENI Slices.

Local suites may form to serve particular groups of research institutions, defined by geography and/or function. However, they will want to affiliate with a larger GENI suite to gain access to resources not in their local suite.

This structure must be supported by the GENI architecture, and particularly by the control framework. In general, a wide variety of structures must be supported by the GENI architecture, and particularly by the control framework.

6 Structures of Similar Projects

There are several existing projects with goals similar or related to GENI. Their structures are instructive to GENI.

6.1 PlanetLab

PlanetLab provides compute resources to many research organizations and has been used to inspire the GENI design. See http://www.planet-lab.org .

6.1.1 Structure

In PlanetLab, there is a central PlanetLab Consortium that has agreements with many separate research organizations (or sites), each of with both provide and utilize resources.

A description of the PlanetLab Consortium is at http://www.planet-lab.org/consortium .

A more detailed view of this two-level structure with hierarchy is shown in Figure 6-1.

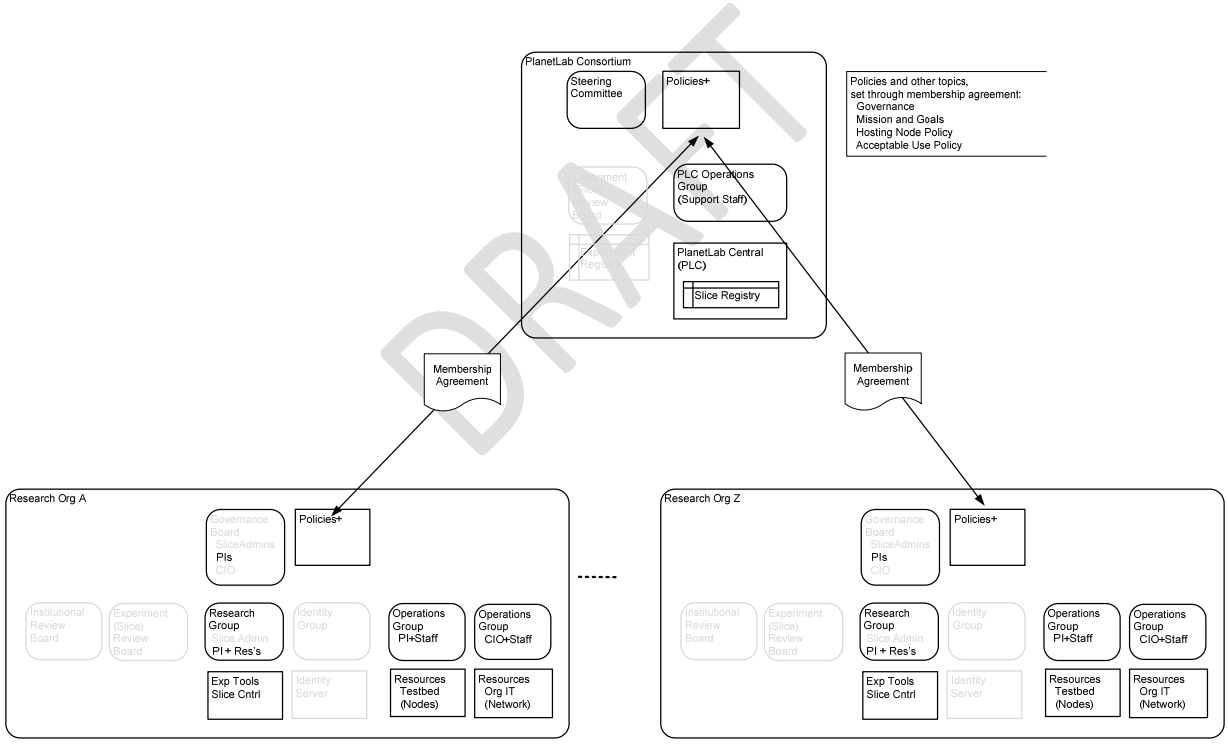

Figure 6-1. PlanetLab Structure.

The figure shows a detailed view of the entities and actors that have evolved to allow PlanetLab to operate efficiently.

6.1.2 Entities and Actors

The PlanetLab Consortium runs an important physical entity, PlanetLab Central (PLC), which in turn contains a Slice Registry (and several other important functions).

In each of the research organizations, agreements have typically been made with PIs, who are the authorities responsible for slices and some of the resources (testbed nodes).

It has been recognized that some of the resources are the responsibility of the campus CIO+staff, and existing agreements should be extended to include these authorities and actors.

6.1.3 Agreements

Policies and other topics are set when a research organization joins the PlanetLab Consortium.

First and foremost, an organization must contribute resources (at least two nodes) before it can utilize resources.

And it must abide by an acceptable use policy.

See Understanding and Resolving Conflicts in PlanetLab, by Larry Peterson - Princeton Univ, at

http://www.cs.princeton.edu/~llp/policy.pdf and slides from GEC3 at http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%202a%20%20llp_policy.ppt

Furthermore, it must agree to common best practices for management. Here is a summary of what is expected:

Appendix: Framework for an Agreement

The following summarizes the expectations placed on researchers, hosting sites, and network testbed operators. We propose it as a framework for a community-wide agreement about managing wide-area network testbeds.

Researchers

• Follow established best practices. This includes first running experiments on a researcher's home site, and then ramping up slowly on the greater testbed.

• Be responsive. This includes monitoring experiments (and support mailing list), and not leaving experiments running on auto-pilot.

• Honor opt-out. This means maintaining blacklist addresses for third-party sites that object to experimental traffic.

• Keep any logs containing user-identifying information private. This also involves being responsive to legal requests to access these logs.

Hosting Sites

• Place nodes in a DMZ. Do not filter ports or ICMP ping packets.

• Forward complaints and concerns to the testbed operators.

• Do not disconnect nodes as this affects all researchers.

• If a compromise of the substrate is suspected, bring nodes into *safe* mode pending resolution.

Network Testbed Operators

• Actively manage the testbed. This includes keeping nodes up-to-date with the latest security patches and monitoring the tesbed for suspicious behavior.

• If a compromise of the substrate is supected, bring suspicious nodes into *safe* mode for further investigation.

• Be responsive to hosting site and third-party complaints.

• Suspend slices that do not adhere to best practices. Suspend users that consistently ignore best practices.

• Engage the research community in defining best practices for each type of experiment. Educate researchers as to these practices.

• Honor opt-out. This means maintaining blacklist addresses for third-party sites that object to experimental traffic.

• Allow hosting sites to opt-out of unwanted slices.

6.2 TeraGrid

TeraGrid provides compute resources for science experiments; see http://www.teragrid.org/ and http://www.teragrid.org/about/ :

TeraGrid is an open scientific discovery infrastructure combining leadership class resources at eleven partner sites to create an integrated, persistent computational resource.

Using high-performance network connections, the TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities around the country. Currently, TeraGrid resources include more than a petaflop of computing capability and more than 30 petabytes of online and archival data storage, with rapid access and retrieval over high-performance networks. Researchers can also access more than 100 discipline-specific databases. With this combination of resources, the TeraGrid is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research.

TeraGrid is coordinated through the Grid Infrastructure Group (GIG) at the University of Chicago, working in partnership with the Resource Provider sites: Indiana University, the Louisiana Optical Network Initiative, National Center for Supercomputing Applications, the National Institute for Computational Sciences, Oak Ridge National Laboratory, Pittsburgh Supercomputing Center, Purdue University, San Diego Supercomputer Center, Texas Advanced Computing Center, and University of Chicago/Argonne National Laboratory, and the National Center for Atmospheric Research.

TeraGrid follows a two-level structure with hierarchy.

TeraGrid has a centralized operations organization.

See http://www.teragrid.org/userinfo/ for user information.

6.3 Open Science Grid

The Open Science Grid provides compute resources for science experiments; see http://www.opensciencegrid.org/ :

The project goal is to stimulate new discoveries by providing scientists with effective and dependable access to an unprecedented national distributed computational facility: the Open Science Grid (OSG). The project will achieve this through the work of the Open Science Grid Consortium: a collaboration of computing

facilities in DOE and NSF, computer scientists, information technology engineers, physicists, biologists, astrophysicists and researchers from other domains to maintain and operate a premier distributed facility, to provide education and training opportunities in its use, and to expand its reach and capacity to meet the needs of the stakeholder organizations. Together the stakeholders in this consortium will sustain and use a shared distributed computing environment that transforms simulation and experimental science in the US.

The Open Science Grid follows a two-level structure with hierarchy; see http://www.opensciencegrid.org/About/Learn_About_Us/OSG_Organization :

The OSG Consortium builds and operates the OSG project. Consortium members contribute effort and resources to the common infrastructure, with the goal of giving scientists from many fields access to shared resources worldwide. See org chart for OSG Consortium (being revised) and the project .

The OSG model of operation is that of a distributed facility which provides access to computing and storage resources at various sites in the US and abroad.

- Resource owners register their resource with the OSG.
- Scientific researchers gain access to these resources by registering with one or more Virtual Organizations (VOs).
- The VO administrators Register their VOs with the OSG. All members of the VO who have signed the acceptable use policy (AUP) are allowed to access OSG resources, subject to the policies of the resource owners.
- Each resource and each VO is supported by a designated, and in some cases shared, "Support Center (SC)," determined at registration time.

There is a collaborative wiki for OSG management activities. The Consortium Council governs the consortium. The OSG Consortium Governance Procedures and By-laws explain how the OSG Consortium works. The Executive Team manages the project. Within the OSG, work is organized into Technical Activities, often joint projects between the OSG project and members of the consortium. Access to more detailed information is available here.

The Open Science Grid has a centralized operations organization: http://www.grid.iu.edu/ .

6.4 Open Cirrus Cloud Computing

The Open Cirrus Cloud Computing Testbed provides federated data centers for open source system and server research; see https://opencirrus.org/ and http://www.usenix.org/events/hotcloud09/tech/full_papers/campbell.pdf .

It offers both physical and virtual machines, and it offers services for: single sign-on (using a global db of usernames and access keys); the ability for each site to admit (or not) each researcher with a sign-on account; monitoring services; storage services; and job submission services.

The Open Cirrus Cloud Computing Testbed follows a two-level structure with hierarchy; see for the governance model:

- Different kinds of participants, which include:
  - *sponsors* (HP, Intel, Yahoo)

- *partners* (i.e. IDA, UIUC, KIT, MIMOS, ETRI, RAS)
- *research members* (who will use the testbed)
- *NSF*

- Committees and sub-teams, including:
  - a Steering Committee that is in overall charge of the Open Cirrus program
  - an operational management organization, responsible for the smooth running of the collaboration and providing support services to all the participants. Initially this will be staffed by Hewlett Packard.
  - a technical team, responsible for getting Open Cirrus up and running
  - a research community, comprising all the Open Cirrus users and developers
- Agreed-upon policies and processes, including:
  - Requesting information on how to become a partner or researcher
  - Applying to conduct research on the testbed
  - Gaining access to the Open Cirrus Portal
  - Gaining access to individual Testbed Sites

The Open Cirrus Cloud Computing Testbed has ?? a centralized operations organization.

6.5 InCommon Federation

The InCommon federation provides identity services; see http://www.incommonfederation.org/ :

InCommon eliminates the need for researchers, students, and educators to maintain multiple passwords and usernames. Online service providers no longer need to maintain user accounts. Identity providers manage the levels of their users' privacy and information exchange. InCommon uses SAML-based authentication and authorization systems (such as Shibboleth®) to enable scalable, trusted collaborations among its community of participants.

The InCommon federation follows a two-level structure with hierarchy, detailed at: http://www.incommonfederation.org/policies.cfm .

The InCommon federation has ?? a centralized operations organization.

7 Reference Single-Suite GENI Structure

7.1 Structure

A straightforward single-suite structure for GENI is based on a two-level hierarchy (see Section 5.2), and includes a GENI suite authority. The GENI suite authority has agreements with all of the management and site authorities such that:

- All management authorities agree to follow a resource use policy that makes resources available to all slice authorities (all slices) in the GENI suite (perhaps with some exceptions).
- All actors agree on cooperative operation policies.
- All actors agree to follow a common set of best practices.

See Figure 7-1.

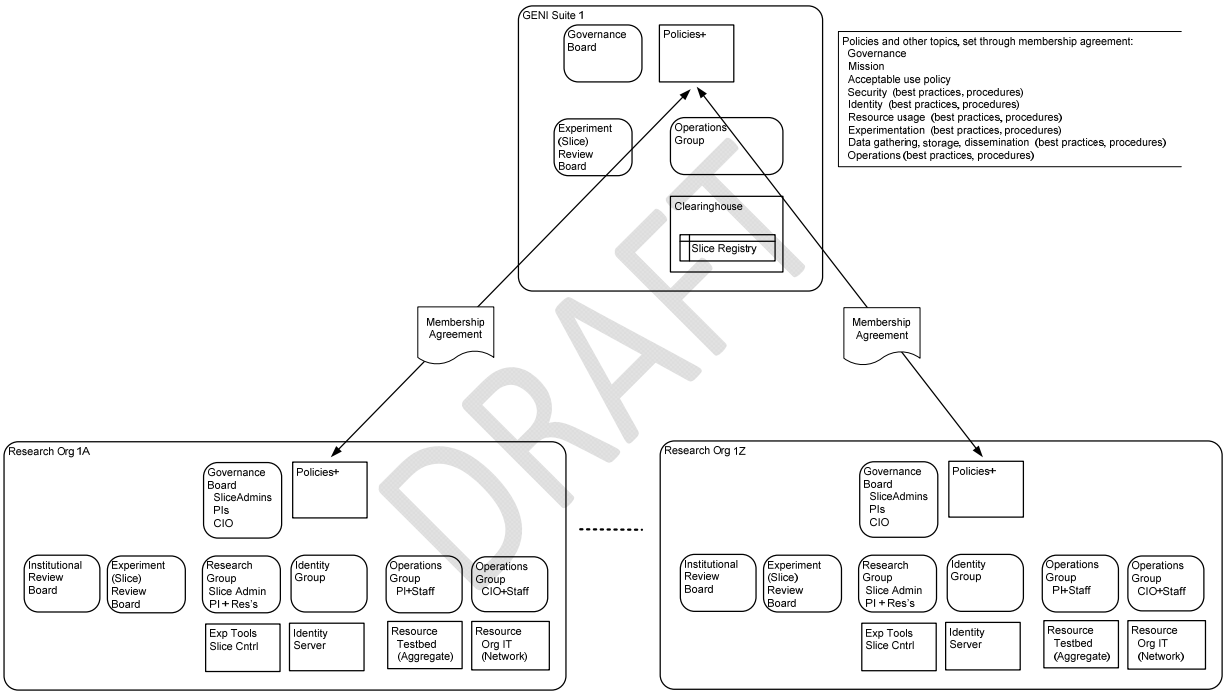

Figure 7-1. Reference Single-Suite GENI Strucutre.

This is the choice that has often been assumed in defining GENI. See the GENI System Overview document at http://www.geni.net/docs/GENISysOvrvw092908.pdf for more details.

When multiple suites are federated and/or affiliated, a more complicated set of agreements is required.

7.2 Entities and Actors

The GENI Suite Authority runs an important physical entity, the GENI Clearinghouse, which in turn contains a Slice Registry (and several other important functions).

In each of the research organizations, agreements should be made with a governance board that includes Slice Administrators, PIs and the CIO.

Several actors have been identified in each research organization, representing functions that will be required for efficient operation, including extensions for end-user opt-in.

These actors include:

- Operations group of CIO and staff.
- Operations group of PI and staff.
- Identity provider group. (new)
- Research group, including Slice Administrators, PIs and Researchers.
- Experiment review board. (new, suggested by PlanetLab experience, and important to screen opt-in experiments)
- Institutional review board. (new, driven by certain opt-in experiments)

7.3 Agreements

Agreements between the GENI Suite Authority governance board and the research organizations governance boards are expected to include:

- Governance
- Mission
- Acceptable use policy
- Security (best practices, procedures)
- Identity (best practices, procedures)
- Resource usage (best practices, procedures)
- Experimentation (best practices, procedures)
- Data gathering, storage, dissemination (best practices, procedures)
- Operations (best practices, procedures)

8 Glossary

The current GENI glossary can be found at: **http://groups.geni.net/geni/wiki/GeniGlossary**