

1 Appendix 1: Experiment Lifecycle Support Tools and Services in GENI Spiral 1

This appendix describes experimenter tools and services expected to be available from each of the five Spiral 1 clusters at the end of Spiral 1. These tools and services are classified according to the taxonomy described in the body of this document.

The objective of this appendix is to understand the different approaches to experimenter tools and services, inform the Spiral 1 clusters about tools they may be able to leverage from other clusters, and identify areas that require additional tools to be developed.

The remainder of this appendix describes tools for experimenter registration, experiment planning, experiment deployment, experiment execution and experiment sunsetting. The appendix includes a summary table at the end.

1.1 ORBIT

ORBIT is a testbed for experimentation with wireless networks. It consists of multiple testbeds (aggregates) of which the indoor ‘grid’ testbed at WINLAB, Rutgers University is the largest. It contains 400 nodes, each equipped with at least two 802.11 radios while some also have additional radios, such as bluetooth, Gnu Radios and other experimental radios attached. The current configuration can be obtained through an “inventory” service of the Aggregate Manager (AM). Nodes also have two wired Ethernet interfaces, one for experiment data (and under full control by the experimenter), and another one exclusively reserved for control and measurement traffic. The ORBIT facility at Rutgers University also hosts a mobile testbed, including mobile wireless nodes with 802.11 radios. For operational reasons, the mobile testbed requires special arrangements with the testbed operator.

In addition to the nodes, the experimenter on the ORBIT “grid” has access to other devices, such as RF signal generators and spectrum analyzers. For instance, the signal generator can be used to raise the noise floor and therefore adjust the RF space and with it the topology of the testbed.

ORBIT provides experimenters with a set of command line tools called OMF (Orbit Control, Management and Measurement Framework) to describe and execute experiments and to collect results (Figure 1). OMF also includes tools for testbed operators to manage and operate the testbed.

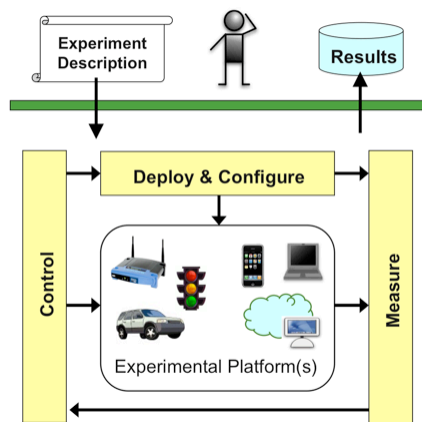


Figure 1. OMF from the experimenter's point of view.

ORBIT users are not required to use the entire OMF toolkit though it is highly recommended they do so. The following sections assume the use of the OMF toolkit.

The material for this section comes from the Rutgers ORBIT and OMF web pages (<http://www.orbit-lab.org> - ORBIT @ Rutgers specific and <http://omf.mytestbed.net> - OMF generic).

1.1.1 Experimenter Registration

ORBIT is open to experimenters from research organizations that belong to the ORBIT Consortium. This consortium is open to any research organization or lab in the world.

To use the ORBIT testbed, the researcher must get an account on the testbed. Accounts are requested using a web form. The form includes a drop-down list of ORBIT consortium members; account requestors must select their organization from this list. The account requested is sent to a designated person at the selected organization who authenticates the requestor and approves the request.

1.1.2 Experiment Planning

Currently, the ORBIT facility at Rutgers University gives experimenters exclusive access to any of the available testbeds at any given time (time sharing). This avoids RF interference between experiments from different experimenters. While the nodes in the testbeds are fairly homogenous, the attached radios change. The “inventory” AM service (currently a REST service) allows for basic discovery of available configurations.

ORBIT provides multiple “sandboxes” that experimenters can use to develop and test their experiments before deploying them on the ORBIT testbed. These sandboxes consist of two nodes and a console and it is strongly recommended that experimenters test their software on a sandbox before reserving time on the larger ORBIT testbed. There are a limited number of sandboxes; ORBIT supports a calendar for scheduling time on the sandboxes and the larger testbed. In addition, various research groups have built their own testbed and using OMF to manage it, reserving the use of the ORBIT “grid” for experiments requiring large scale.

Experimentation on ORBIT begins with the development of a script that describes the experiment. The script describes the resources required for the experiment, a description of how they are initialized, connections between resources, and actions to be performed on or with the resource during the experiment. The script is written in a language based on the Ruby scripting language. The OMF experiment scripting language provides common abstractions for configuring network resources and wrappers for various commonly used tools and applications to generate, receive, and measure network traffic flows. Experimenters can deploy and integrate their own applications, or extend the OMF resource controllers to add new functionality.

1.1.3 Experiment Deployment

As mentioned earlier, experimenters reserve (lease) the entire ORBIT testbed or an ORBIT sandbox to deploy and run experiments. ORBIT provides an on-line calendar. Experimenters with access to ORBIT can log into this calendar, find a time when the testbed is available, and reserve the testbed for up to a maximum of two hours. Reservations are considered pending until they are approved by a scheduling service. It is possible for reservation requests to conflict. ORBIT employs a conflict resolution algorithm based on a number of factors including how much time the experimenter got on the testbed over the previous two weeks. An email notification is sent to the experimenter when the requested timeslot is approved.

During an experimenter’s approved time slot, he can use ssh to log into the console of the testbed or sandbox for which he has a reservation. The experimenter then uses the OMF load command to install

a disk image on the nodes in the testbed. This can be a standard baseline disk image provided by ORBIT or an experimenter provided disk image.

1.1.4 Experiment Execution and Instrumentation

After a disk image has been loaded into the testbed node, the experimenter copies the experiment script into the console machine and runs this script using the OMF “exec” command. This command starts an Experiment Controller that runs the experiment and controls its execution. This is illustrated in Figure 3.

The Experiment Controller (EC) communicates with the relevant Aggregate Managers and Resource Controllers to configure resources, install applications and services, and ultimately orchestrate the experiment itself. It also receives status and error messages, which are reflected in a dynamic XML document maintained by the EC. An experiment description can define triggers through XPath expressions and associated action handlers. A configurable logging mechanism will print relevant messages to the console, file, or remote logging facility, depending on experimenter’s preferences. The EC also includes a web server which allows for a web browser based interaction with the experiment.

Most resource configuration options are dynamic, meaning they can be changed during an experiment. This ranges from changing the sending rate of a traffic generator, to switching the frequency of a wireless card, or even turning a node on or off.

The EC also allows for temporarily disconnected experiments, where some of the controlled resources are temporarily not connected to the control network. This is quite typical for many mobile experiments, but also of great interest for testbeds without a dedicated control network where control messages are temporarily stopped as not to interfere with the experiment itself.

OMF includes a measurement collection framework called OML (ORBIT Measurement Framework). For every experiment OML creates a database with tables for each type of measurement specified for the experiment.

OML consists of a central collection server that collects and stores measurements taken at the nodes and other resources, such as spectrum analyzers, in the experiment. Various applications, services, and daemons have been instrumented with OML “measurement points”. A measurement point is essentially a source for a stream of tuples. The tuple’s size and item types are defined for each measurement point. Various language bindings exist to register and later use measurement points from inside a program. While the measurement points are defined by the developer of an application or service, the experimenter defines at run time which of the measurements available are of interest and what additional local pre-processing should be performed before the resulting “measurement stream” is sent towards the OML Collection service. These configurations, described in an XML file, and the resulting processing steps take place inside the OML Client Library, which is normally linked to the respective application or service.

A resource resident OML Proxy Server can be used to intercept and temporarily store measurements to support the above described temporarily disconnected experiments.

The OML collection service stores all the measurement streams it receives in a database. The streams include meta data describing their respective schema from which the collection service will automatically create database tables in which to store the stream elements. Collected measurements can be accessed and analyzed during and at the end of the experiment through SQL queries on the experiment’s database. The AM provides services to execute the query on behalf of an authorized user and returns the result in various formats for further processing. The EC contains experimental support to define measurement queries and visualize the results through the embedded web server.

The AM also provides a service that allows the experimenter to download the entire database for local processing with tools, such as Matlab and Excel or by user developed scripts.

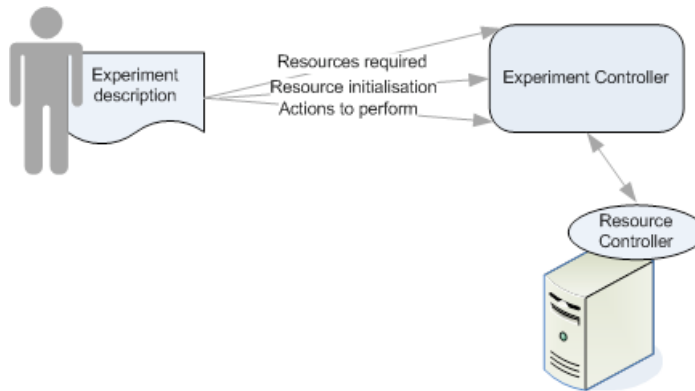


Figure 3. From experiment description to experiment execution in ORBIT.

1.1.5 Experiment Sunsetting

ORBIT does not have any formal archive for experiments where others can find them to replicate or extend.

1.1.6 Other Tools

Both ORBIT and OMF have wiki pages with tutorials for using the testbed and the toolkit. The OMF wiki includes forums for news and discussions about developing and using OMF.

1.2 ORCA/BEN

ORCA is an extensible architecture for on-demand networked computing infrastructure. Its purpose is to manage the hosting of diverse computing environments (guests) on a common pool of networked hardware resources such as virtualized clusters, storage, and network elements. Examples of guests include network application services, virtual machines for personal computing, managed virtual clusters for large-scale computation, or other networked systems. The hardware resources may be owned and administered by multiple infrastructure providers, who control the terms of their use.

ORCA is based on a foundational abstraction of resource leasing. A resource lease is a contract between a resource provider and a resource consumer (guest). The contract grants the consumer access to some resource for a specified period of time, with additional contract terms defining the nature of the resource and its configuration. Orca resource consumers use open leasing protocols and programmatic interfaces to negotiate contracts and acquire and coordinate the underlying resources, optionally assisted by brokering intermediaries.

Figure 4 shows the major components of ORCA. The *Site Authorities* are akin to GENI aggregate

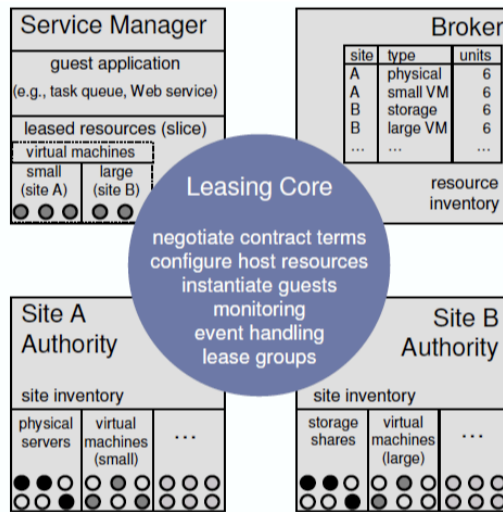


Figure 4. Major Components of ORCA.

managers: they control infrastructure resources at a particular site, administrative domain, or component aggregate comprising a set of servers, storage units, network elements, or other components under common ownership and control. *Brokers* are analogous to GENI clearinghouses: they mediate resource discovery and arbitration by controlling the scheduling of resources at one or more sites over time. Finally, *Service Managers* provide a runtime environment for *guest applications*. Guest applications are distributed applications running on ORCA. They use Service Manager services to procure lease contracts to use resources from one or more sites, learn about changes in resource status and adapt to these changes.

The ORCA/BEN project is extending ORCA to be the control framework based on BEN, a metro-scale experimental optical network in Research Triangle Park, North Carolina.

The above description of ORCA and the material in the remainder of this section is drawn from the ORCA web site (<http://nicl.cod.cs.duke.edu/orca/about.html>), from the paper “Sharing Networked Resources with Brokered Leases” (<http://www.cs.duke.edu/nicl/pub/papers/shirako06.pdf>) and from the technical note “ORCA Technical Note: Guests and Guest Controllers” (<http://www.cs.duke.edu/nicl/pub/papers/control.pdf>).

1.2.1 Experimenter Tools and Services

Experimenter tools and services in ORCA are envisioned to be guest applications similar to any other distributed applications that run on ORCA. These applications run in the context of the Service Manager shown in Figure 4.

At this time the ORCA project does not have any officially supported experimenter tools and services other than the services provided by the Service Managers. The PlanetLab Plush tool has however been ported to be an ORCA guest application. It is expected by the end of Serial 1 the Gush tool (Section 1.4.4) will be ported to be an ORCA guest application.

1.3 ProtoGENI

ProtoGENI is based on the University of Utah Emulab testbed. In Spiral 1 the ProtoGENI project is primarily focused on control framework issues. The experimenter workflow tools described in this section are therefore primarily Emulab tools i.e. they do not take advantage of the added ProtoGENI functionality. Much of the material in this section has been gathered from documentation at <https://users.emulab.net/trac/emulab/wiki>.

1.3.1 Experimenter Registration

The ProtoGENI testbed is accessed by a web interface (<http://www.emulab.net/>). Experimenters wishing to use the testbed must first register to create an account. Typically, someone who will be the project leader requests permission from Testbed Ops/Admin, via the web interface, to create a project. In academic parlance, a project leader is a "principal investigator." That person is expected to be someone who is responsible, whose position is more or less verifiable, and is therefore accountable. The project leader is held responsible for the actions of members of the project. People working on the project (students, staff, etc.) will request permission from the project leader, also via the web interface, to join the project. Once project members have been authorized by the leader, they can use the Web interface to login, start experiments, reserve and configure nodes, etc..

1.3.2 Experiment Planning

An experimenter logged into Emulab can use a browsing tool, implemented as a Java applet, to view available resources. The tool displays a resource name and a description. In the case of compute resources, the default OS on that node is also listed.

Network topologies for experiments on Emulab are described using a scripting language very similar to that used by the ns-2 network simulator. The experiment description includes defining nodes and links in the network. The definition of a node includes giving it a name that is meaningful to the experiment and the operating system to be loaded on the node. The definition of a link includes the nodes it connects and its bandwidth, latency, queue type and packet loss rate. The script may also used to specify how packets are to be routed in the network. Currently supported routing protocols are Session (dynamic routing using the OSPF protocol running on all nodes in the experiment), Static (routes are pre-computed by a distributed route computation algorithm running on the experiment nodes) and Manual (experimenter specifies per node routing information).

Emulab provides a graphical tool that can be used to describe the network for an experiment; the tool generates a corresponding ns file. The graphical tool is limited in its expressiveness and can be tedious to use to specify large networks.

1.3.3 Experiment Deployment

Experimenters use the Emulab web interface to run experiments. They log on to Emulab and create an experiment in the context of a project. They provide an experiment name, a short text description and an ns file that describes the network topology for the experiment (Section 1.3.2). The testbed processes the ns file to map the network described in the script to physical nodes and links. It also assigns IP addresses to nodes.

When Emulab is done processing the ns script it sends the experimenter an email along with a listing of the nodes in the network, physical machines to which these nodes have been mapped, IP addresses and link characteristics. The Emulab configuration system will ensure these nodes are configured with the specified operating systems and are ready to use. The network interfaces on the

nodes are configured appropriately; the project's NFS directory is mounted on every node so files can be easily shared amongst the nodes; user accounts are created for project members; and the `/etc/hosts` file is set up so other nodes can be referenced by name rather than IP address.

On receipt of this email the experimenter can use `ssh` to log into any of the nodes in the experiment. The experimenter can then install any software needed for the experiment. Experimenters have root access to their nodes and can change system configurations if needed. Experimenters can also have the system automatically start up a program on each node from their project's NFS mounted directory when the node boots up; this is specified in the `ns` description of the experiment.

1.3.4 Experiment Execution and Instrumentation

NEED INFORMATION on Experiment control, instrumentation and measurement.

1.3.5 Experiment Sunsetting

Emulab does not have any formal archive for experiments where others can find them to replicate or extend.

1.3.6 Other Tools

The Emulab project maintains a wiki page with tutorials and other documentation to help experimenters use the testbed.

1.4 PlanetLab

PlanetLab consists of a collection of machines hosted by members of the PlanetLab consortium. All PlanetLab machines run a common software package that includes a Linux-based operating system; mechanisms for bootstrapping nodes and distributing software updates; a collection of management tools that monitor node health, audit system activity, and control system parameters; and a facility for managing user accounts and distributing keys. All PlanetLab related information in this section comes from the PlanetLab web pages (<http://www.planet-lab.org/>).

PlanetLab provides a base set of tools such as PlanetLab shell (`plcsh`), a command-line tool, to set up and deploy experiments. However, PlanetLab supports an `xml-rpc` interface that allows other experimenter tools and services developed.

Gush and Raven are examples of two tools being developed in GENI Spiral 1 to support experimentation with PlanetLab. This section describes experimentation with PlanetLab using these tools. Information on Gush and Raven comes from <http://gush.cs.williams.edu/trac/gush> and <http://www.cs.arizona.edu/stork/> respectively.

1.4.1 Experimenter Registration

Researchers wishing to use PlanetLab get accounts on the system through an institution (typically their home institution) that is a member of the PlanetLab Consortium. Institutions join the PlanetLab Consortium by signing a membership agreement and connecting two or more nodes to the PlanetLab infrastructure.

Researchers get an account by reading the PlanetLab Acceptable User Policy and completing a form on the PlanetLab website. The form requires the user to select a PlanetLab consortium member institution (typically the researcher's home institution); the account request is sent to the PI of the

institution for approval. When the PI approves the request the user gets an email confirming the creation of an account.

PlanetLab uses an ssh public-key based authentication mechanism to allow researchers to connect to PlanetLab resources (nodes). Experimenters create their own key pair and upload their public key to the PlanetLab database.

Experiments in PlanetLab run within the context of a slice. A slice is a collection of resources distributed across multiple PlanetLab nodes. Experimenters are given a new slice or added to an existing slice by the PlanetLab PI at their institution.

1.4.2 Experiment Planning

There are two different sets of tools available to PlanetLab experimenters for discovering resources for experiments: CoMON and SWORD. CoMON collects statistics on PlanetLab nodes so experimenters can query for lightly loaded nodes, most reliable nodes, popular nodes or nodes at a site. The SWORD tool is based on CoMon in that it uses data collected by CoMON. The query language of SWORD is more flexible and expressive than the CoMon query language.

1.4.3 Experiment Deployment

As mentioned earlier, PlanetLab experimenters must first obtain a slice to which they add resources for experimentation. Experimenters are given a new slice or added to an existing slice by the PlanetLab PI at their institution.

Gush users describe their experiment using the Gush experiment description language, based on xml. The experiment description includes a description of the PlanetLab resources needed by the experiment and the software packages to be installed in these resources.

After experimenters obtain PlanetLab slices they start up the Gush command line tool and use the “load” command to load a file with their experiment specifications. They then use the Gush “run” command to start the experiment. Gush parses the experiment specification, connects to the PlanetLab central server to obtain information about available slices, finds resources that match the requirements in the experiment specification, adds these resources to the slice if needed, and connects to these resources. Gush optionally uses SWORD to find the appropriate PlanetLab resources. Finally, Gush installs the specified software packages on these resources and starts the execution.

PlanetLab users may also use Raven/Stork, a more full-featured software deployment tool. Raven/Stork is especially useful for long-running experiments where new versions of software may have to be deployed during the course of the experiment or software re-deployed on nodes that might have re-booted.

1.4.4 Experiment Execution and Instrumentation

Gush can also be used to monitor and control a running experiment. For example, the Gush “info” command can be used to get status of the slice, nodes in the slice or activity such as file transfers.

Gush can also be used to specify actions to be taken if parts of the experiment fail. By default, Gush reacts to process failures and resource failures. The standard recovery options include restarting failed processes, reconnecting to failed resources, finding and configuring replacement resources, and aborting the entire experiment. Custom application-specific failure recovery behaviors can also be described within the experiment specification.

INSTRUMENTATION TOOLS/SERVICES ON PLANETLAB?

1.4.5 Experiment Sunsetting

PlanetLab does not have any formal archive for experiments where others can find them to replicate or extend.

1.4.6 Other Tools

The PlanetLab project is also developing a command-line tool called SFI for use by experimenters.

1.5 TIED/DETER

The TIED/DETER federation system allows a researcher to construct experiments that span testbeds by dynamically acquiring resources from other testbeds and configuring them into a single experiment. As closely as possible that experiment will mimic a single DETER experiment. The DETER testbed is a public facility for medium-scale repeatable experiments in computer security. Built using Utah's Emulab software, the DETER testbed has been configured and extended to provide effective containment of a variety of computer security experiments, including defense against attacks such as DDoS, worms, viruses, and other malware, as well as attacks on the routing infrastructure.

TIED/DETER supports creation of cohesive experiments (slices) from independently administered resources (components/aggregates). Because resources are independently administered and serve different communities, the authorization system needs to support a rich delegation structure, formal semantics, efficient negotiation, and clear auditing. The Attribute Based Access Control (ABAC) [http://www.isso.sparta.com/research_projects/security_infrastructure/abac_overview.html] system meets those requirements; TIED is integrating this into the federation system.

To make use of widely distributed components it is helpful to establish guaranteed network connections between them. TIED is addressing this by federating with testbeds that represent dynamically allocatable wide-area network resources. The prototyping plan is to use DRAGON interfaces to configure these resources.

Much of the information in this section comes from <http://www.isi.deterlab.net>, <http://groups.geni.net/geni/wiki/TIEDClearinghouse>, and <http://seer.isi.deterlab.net/trac>.

1.5.1 Experimenter Registration

Creation of a TIED/DETER user account is very similar to that with ProtoGENI/Emulab: The user goes to a web page (<http://www.isi.deterlab.net/>) and requests an account. As with Emulab, the user can request to join an existing project or create a new project.

A registered TIED user has access to the various aggregates TIED works with and can create slices across them.

1.5.2 Experiment Planning

Experiment planning is similar to ProtoGENI/Emulab: Experimenters can log into their TIED/DETER accounts and browse available resources.

1.5.3 Experiment Deployment

For TIED, slice management is the creation and manipulation of experiments that span multiple facilities. The facilities are aggregates, each of which controls access to a set of resources. Currently TIED supports creation of slices across testbeds with an Emulab interface and that run instances of a daemon that is the TIED/DETER aggregate manager.

TIED's slice management features can be accessed through a simple web interface that makes common actions straightforward, and a command line interface that is more expressive for users with sophisticated needs.

Experimenters can create a slice either by activating one of the existing slice descriptions, which will create the slice based on its last configuration, or by going directly to the slice creation page and providing a slice description and a name. Slice descriptions are ns files similar to the ProtoGENI/Emulab network description, extended to describe the embedding of nodes in the various aggregates.

TIED keeps a record of slices created and allows users to recreate them or terminate running ones through a web interface. This web interface is shown in **Error! Reference source not found.** The image associated with each slice listed in the figure shows the logical topology of the slice. The master testbed column indicates which testbed is exporting its user environment into the experiment; for TIED slices, this will usually be the DETER testbed. The project column will be TIED for TIED

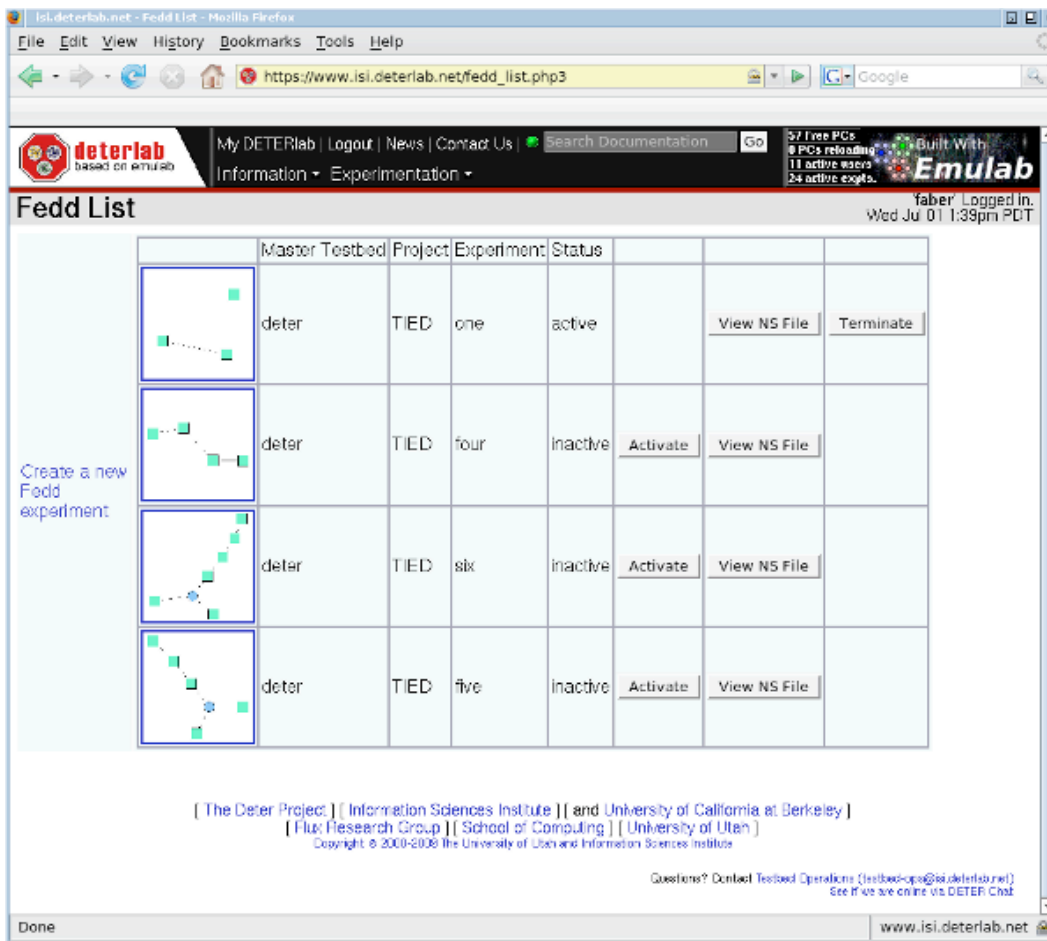


Figure 5. TIED Slice management interface.

experimenters; it indicates that the DETER TIED project is being exported into the slice. The name of the slice is a mnemonic name attached to the slice. There is also a button to retrieve the slice description used to create this slice, either to inspect it or to copy it for extension. Finally there is an

indication of the slice's status (active, inactive, swapping, or unknown). Based on the current status there are buttons to create or terminate it. Termination removes the nodes from active use, but the description remains in this database, so a user may recreate it at a later time.

As mentioned earlier, resources are added to a slice by uploading a slice description.

1.5.4 Experiment Execution and Instrumentation

The SEER experiment control tool is a useful way to manipulate a TIED experiment to both carry out the experiment and to examine real time data from the experiment. Other Emulab-style measurement gathering and event systems may also be used with TIED.

To run and manage an experiment using SEER, experimenters must install the Java-based SEER GUI on their workstations. Once the experiment is swapped in (similar to experiment swap-in in Emulab), the experimenter starts up the SEER GUI and attaches to the experiment. Users must supply their TIED username, password and ssh passphrase to connect to their experiments.

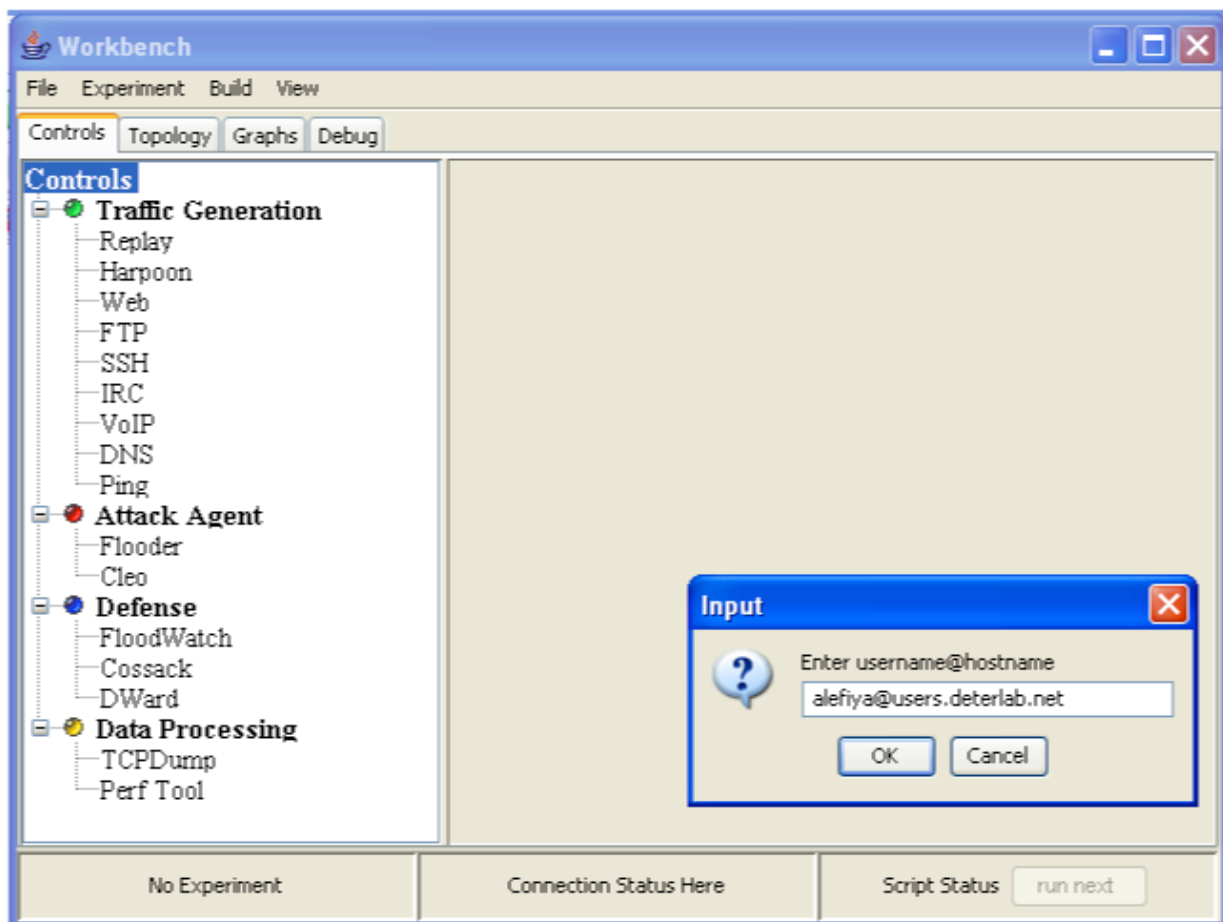


Figure 6. SEER Graphical User Interface.

Additionally, experimenters must modify the ns file describing their experiment to: (1) Create an additional unconnected node in the experiment called 'control' to gather and display statistics and act as the remote GUI client, (2) Specify the SEER agents to run on each node in the topology, and (3) Have the script run a SEER provided python script on each node when it starts up.

The SEER GUI includes a number of different views (implemented as tabs in the interface) to control and visualize the experiment (Figure 7). The “controls” tab on the GUI is used to associate different types of traffic generators to the nodes in the network. Examples of traffic generators include Web, VoIP and FTP traffic. Different types of attack traffic and network defenses can also be associated with these nodes. Finally, the “controls” tab allows different packet analysis tools such as TCPDump to be started up on nodes in the network.

Another tab, labeled “Topology” on the SEER graphical user interface can be used to visualize the network and traffic on the network links. Clicking on the “Topology” tab causes the SEER interface to obtain logical information about the slice and populate the topology window. To get a real-time summary of traffic at a node, the experimenter places the mouse pointer on the node in the topology window, right clicks to access the node menu, and requests a graph. The graph appears under the "graphs" tab.

1.5.5 Experiment Sunsetting

TIED does not have any formal archive for experiments where others can find them to replicate or extend.

1.5.6 Other Tools

The TIED project maintains a web page with documents and videos to help experimenters get started with TIED and SEER.

2 Summary of Experimenter Tools and Services Available

	ORBIT	ORCA	ProtoGENI/Emulab	PlanetLab	DETER
Experimenter Registration	ORBIT consortium members register on ORBIT web site. Approved by designated person at member institution.		Project leader applies for account on web site. Account approved by Emulab admin. Project team member accounts approved by project leader.	PlanetLab consortium members can apply for accounts. Approval by PlanetLab PI at member institution.	Similar to Emulab.
Experiment Planning and Specification	Test and debug small scale version of testbed. Experiment described using script written in Ruby-based language.		NS-2 like scripting language to describe experimental network including OS images to be loaded on network nodes.	CoMON and Sword for resource discovery. GUSH experiment specification language.	Similar to Emulab. Web tool to create slices from scratch or one like a previously created slice

<p>Experiment Deployment</p>	<p>On-line calendar to request and get access to testbed. Log into testbed console to upload software. Experiment controller interprets experiment description, loads software into nodes and loads software into nodes.</p>		<p>Experimenters have ssh access to nodes: can upload any software onto these nodes.</p>	<p>GUSH interprets experiment specification to obtain resources, and install software on nodes. Raven/Stork may also be used for deploying software on nodes---it is especially useful for long-running experiments.</p>	<p>Same as Emulab.</p>
<p>Experiment Control</p>	<p>Commands to turn nodes on and off.</p>	<p>GUSH?</p>		<p>GUSH: To start, stop, pause experiments. GUSH also supports monitoring of experiment and specification of handlers that are invoked if nodes in the experiment crash.</p>	<p>SEER</p>
<p>Instrumentation and Measurement</p>	<p>Measurement system (OML) has clients on each node that collect measurements and send to a central collection server. Measurements include those</p>				<p>SEER can be used to visualize network traffic to/from a node in the experiment.</p>

	generated by the experiment software.				
Sunsetting Tools	None		None	None.	None
Other Tools	Wiki pages with tutorials. News and discussion pages on OMF wiki.		Wiki pages with tutorials and other documentation.	Wiki pages with tutorials and other documentation.	Wiki pages with tutorials and other documentation