

NEPI v3.0 User Manual

Contents

1 FAQ

1.1 What is NEPI?

NEPI is not a network simulator, nor an emulator or a testbed. NEPI is a Python library that provides classes to describe and run network experiments on different experimentation platforms (e.g. Planetlab, OMF wireless testbeds, network simulators, etc).

Imagine that you want to run an experiment to test a distributed application you just coded, on the Internet. You can use NEPI to deploy your application on PlanetLab nodes, run the experiment, and collect result files you might have generated during the experiment (e.g. pcap files from tcpdump).

Sure, you could do this by coding your own BASH script, but it will probably take more time and painful hours of debugging if you want to do it right. NEPI aims at providing a re-usable code base to run network experiments on target experimentation platforms, so to decrease the time you spend in developing platform specific scripts or programs, and debugging them.

In a nut-shell, NEPI is a network experiment management framework which provides a simple way of describing network experiments, and the logic to automatically deploy those experiments on the target experimentation environments. It also provides the means to control the resources used in the experiment (e.g. Nodes, applications, switches, virtual machines, routing table entries, etc) during experiment execution, and to collect results generated by the experiment to a local repository.

The experiment deployment and control is done by the Experiment Controller (EC) entity, which is responsible for the global orchestration of the experiment. The EC knows nothing about how to manage specific resources (e.g. how to configure a network interface in a PlanetLab node), instead it delegates those tasks to entities called Resource Manager (RM).

The RMs are responsible of controlling single resources (e.g. a Linux host, an Open vSwitched on PlanetLab nodes, etc). Different types of resources will be controlled by different RMs, specifically adapted to control them. All RMs implement a same external interface, that the EC uses to control them in a uniform way.

NEPI is not magical, it can not control all existing resources on all existing experimentation platforms by default. However, potentially any resource could be controlled by NEPI if the adequate Resource Manager is implemented for it. Fortunately, NEPI already provides several Resource Managers for different resources on a variety of testbeds, and new Resource Manager classes can be extended from existing ones, to control new types of resources.

The idea behind NEPI is to enable running network experiments on potentially any experimentation platform, using a single software tool, as opposite to using a dedicated software for each platform. An additional perk is that you don't have to deal with a lot of platform-specific gory details of setting up and configuring the resources (e.g. Creating a TAP device on Planetlab. If you ever had to do that, you know what I mean). Also, you could combine resources from different platforms in a same experiment, using just one script.

So, 'One ring to rule them all', sorry I meant, 'One tool to control them all'... or something like that. We thought it was a good idea to abstract platform details behind a common resource management interface, and let NEPI deal with the details and give you back the results.

1.2 What does a NEPI script look like ?

Here is a very simple experiment example, which runs a PING to "nepi.inria.fr" from a given host. Note that you will need to replace the hostname, username, and ssh_key variables with the actual values to run the example.

```
from nepi.execution.ec import ExperimentController

ec = ExperimentController(exp_id = "myexperiment")

hostname = # Host that can be accessed with an SSH account
username = # SSH user account on host
ssh_key = # Path to SSH public key file to access host

node = ec.register_resource("LinuxNode")
ec.set(node, "hostname", host)
ec.set(node, "username", user)
ec.set(node, "identity", ssh_key)

app = ec.register_resource("LinuxApplication")
ec.set(app, "command", "ping -c3 nepi.inria.fr")
ec.register_connection(app, node)

ec.deploy()
```

```
ec.wait_finished(app)

print ec.trace(app, "stdout")

ec.shutdown()
```

1.3 What does NEPI stands for?

It stands for: Network Experiment Programming Interface.

1.4 Who developed NEPI?

NEPI was developed at INRIA, Sophia Antipolis France. A first prototype was implemented in 2010. Versions 1.0 and 2.0 were released in 2011 and 2012, respectively. The current version is 3.0, and it was completely redesigned and rewritten to broaden the scope, and to include several cool new features, which will be described in detail in this document. The following people has contributed to the project:

- NEPI version 3.0: Alina Quereilhac, Julien Tribino, Lucia Guevgeozian Odizzio, Alexandros Kouvakas
- NEPI versions 1.0 and 2.0: Alina Quereilhac, Claudio Freire, Martin Ferrari, Mathieu Lacage
- NEPI prototype: Martin Ferrari, Mathieu Lacage
- Other contributors: Dirk Hasselbalch

1.5 Is it free?

Yes, NEPI is free software. It is free to use, free to to modify, free to share. NEPI v3.0 is licensed under GPL v3, so you can do whatever you want with it, as long as you keep the same license.

1.6 How can I contribute?

There are many ways you can contribute to the project. The first one is using it and reporting bugs. You can report bugs on the NEPI bugzilla page (**not yet set-up, so for now send the bugs by email to the users list**):

<https://nepi.inria.fr/bugzilla/>

You can also become a part of the NEPI community and join our mailing lists:

- To subscribe to the users mailing list at *nepi-users@inria.fr* you can send an email to *sympa@inria.fr* with subject *Subscribe nepi-users <put-your-user-name-here>*
- To subscribe to the developers mailing list at *nepi-developers@inria.fr* you can send an email to *sympa@inria.fr* with subject *Subscribe nepi-developers <put-your-user-name-here>*

To contribute with bug fixes and new features, please send your code patch to the *nepi-developers* list.

1.7 How can I report a bug ?

To report a bug take a look at the NEPI bugzilla page at:

<https://nepi.inria.fr/bugzilla/>

1.8 Where can I get more information ?

For more information visit NEPI web site at

<https://nepi.inria.fr>

Getting started 2

NEPI is written in Python, so you will need to install Python before before being able to run experiments with NEPI. NEPI is known to work on Linux (Fedore, Debian, Ubuntu) and Mac (OS X).

2.1 Dependencies

Dependencies for NEPI vary according to the features you want to enable. Make sure the following dependencies are correctly installed in your system before using NEPI.

Mandatory dependencies:

- Python 2.6+
- Mercurial

Optional dependencies:

- SleekXMPP - Required to run experiments on OMF testbeds

Install dependencies on Debian/Ubuntu

```
$ sudo aptitude install -y python mercurial
```

Install dependencies on Fedora

```
$ sudo yum install -y python mercurial
```

Install dependencies on Mac

First install homebrew (<http://mxcl.github.io/homebrew/>), then install Python.

```
$ brew install python
```

Install SleekXMPP

You will need *git* to get the SleekXMPP sources.

```
$ git clone -b develop git://github.com/fritzy/SleekXMPP.git
$ cd SleekXMPP
$ sudo python setup.py install
```

2.2 The source code

To get NEPI's source code you will need Mercurial version control system. The Mercurial NEPI repo can also be browsed online at:

```
http://nepsi.inria.fr/code/nepsi/
```

Clone the repo

```
$ hg clone http://nepsi.inria.fr/code/nepsi -r nepi-3.0-release
```

2.3 Install NEPI in your system

You don't need to install NEPI in your system to be able to run experiments. However this might be convenient if you don't plan to modify or extend the sources.

To install NEPI, just run *make install* in the NEPI source folder.

```
$ cd nepi
$ make install
```

If you are developing your own NEPI extensions, the installed NEPI version might interfere with your work. In this case it is probably more convenient to tell Python where to find the NEPI sources, using the `PYTHONPATH` environmental variable, when you run a NEPI script.

```
$ PYTHONATH=$PYTHONPATH:<path-to-nepsi>/src python experiment.py
```

2.4 Run experiments

There are two ways you can use NEPI to run your experiments. The first one is writing a Python script, which will import NEPI libraries, and run it. The second one is in interactive mode by using Python console.

Run from script

Writing a simple NEPI experiment script is easy. Take a look at the example in the FAQ section ???. Once you have written down the script, you can run it using Python. Note that since NEPI is not yet installed in your system, you will need to export the path to NEPI's source code to the PYTHONPATH environment variable, so that Python can find NEPI's libraries.

```
$ export PYTHONPATH=<path-to-nepi>/src:$PYTHONPATH
$ python first-experiment.py
```

Run interactively

The Python interpreter can be used as an interactive console to execute Python instructions. We can take advantage of this feature, to interactively run experiments with NEPI. We use the *ipython* console for the example below, you can install it with *sudo yum/aptitude install ipython* on Fedora/Debian.

```
$ export PYTHONPATH=<path-to-nepi>/src:$PYTHONPATH
$ ipython
Python 2.7.3 (default, Jan 2 2013, 13:56:14)
Type "copyright", "credits" or "license" for more information.

IPython 0.13.1 -- An enhanced Interactive Python.
?          -> Introduction and overview of IPython's features.
%quickref  -> Quick reference.
help       -> Python's own help system.
object?    -> Details about 'object', use 'object??' for extra details.
```

With *ipython*, if you want to paste many lines at once you will need to type *%cpaste* and finish the paste block with *.*

The first thing we have to do is to import the NEPI classes we will use. In particular we need to import the *ExperimentController* class.

```
from nepi.execution.ec import ExperimentController
```

Then we need to create an *ExperimentController* instance. The *exp-id* argument serves as a human readable identifier for the experiment to be ran.

```
ec = ExperimentController(exp_id = "<your-exp-id>")
```

Next we will define two Python functions, one to register *LinuxNode* resources and the other to register *LinuxApplication* resources. A *LinuxNode* resource (or ResourceManager) will serve as an abstraction to a Linux host resource, that can be accessed using SSH key authentication. A *LinuxApplication* resource represents anything that can be executed on a Linux host as a BASH command.

```
%cpaste
def add_node(ec, host, user, ssh_key):
    node = ec.register_resource("LinuxNode")
    ec.set(node, "hostname", host)
    ec.set(node, "username", user)
    ec.set(node, "identity", ssh_key)
    ec.set(node, "cleanHome", True)
    ec.set(node, "cleanProcesses", True)
    return node

def add_app(ec, command, node):
    app = ec.register_resource("LinuxApplication")
    ec.set(app, "command", command)
    ec.register_connection(app, node)
    return app
--
```

The method *register_resource* declares a resource instance to the Experiment Controller. The method *register_connection* indicates that two resource will interact during the experiment. Note that invoking *add_node* or *add_app* has no effect other than informing the EC about the resources that will be used during the experiment. The actual deployment of the experiment requires the method *deploy* to be invoked.

The Resource Managers (RM) that abstract the concrete resources expose configuration attributes. In the *LinuxNode* RM we set the *hostname* and *username* as attributes, while in the *LinuxApplication* RM we set the *command* attribute.

Apart from the *command* attribute, the *LinuxApplication* ResourceManager exposed several other attributes that permit to upload, compile and install arbitrary sources, to run any application might be needed to run an experiment. More details will be given in the following sections of this document.

Lets now use these functions to describe the experiment we will run. Choose a host where you have an account, and can access using SSH key authentication. Define string variables with the right values for the *hostname*, *username* and path to the SSH public key in as *ssh_key*, and then type the following lines.

```
node = add_node(ec, hostname, username, ssh_key)
app = add_app(ec, "ping -c3 nepi.inria.fr", node)
```

Now that we have described our simple PING experiment, we can deploy it. Invoking the *deploy* command will not only configure the resource but also automatically launch the applications.

```
ec.deploy()
```

After some seconds, we should see some log messages informing about the progress of the deployment. If you now open another terminal and connect to the host through SSH, you should find some directories created by NEPI. You should see a directory named *nepi-exp*, and under that directory you will find another with the identifier you gave when you created the experiment controller (the *<exp-id>*). Under the experiment directory, you will find directories for each of the resources deployed (e.g. *node-1*, *app-2*). The resource directories are named with a short string that identifies the type of resource (e.g. 'node', 'app', etc), followed by a unique number that uniquely identifies a given resource within the experiment, the global unique identifier (guid).

From the ipython console, we can check the deployment status of each resource by querying the EC with the method *state*. The argument *hr* stand for 'human readable', and will return a string state instead of a state number.

```
ec.state(app, hr=True)
```

Once the LinuxApplication is STARTED, we can retrieve the PING output using stored as a *trace* file on the host. For this we use use *trace* method, and specify the resource and the trace name (i.e. stdout).

```
ec.trace(app, "stdout")
```

That is it. We can terminate the experiment by invoking the method *shutdown*.

```
ec.shutdown()
```

Introduction to NEPI

3

During the past decades, a wide variety of platforms to conduct network experiments, including simulators, emulators and live testbeds, have been made available to the research community. Some of these platforms are tailored for very specific use cases (e.g. PlanetLab for very realistic Internet application level scenarios), while others support more generic ones (e.g. ns-3 for controllable and repeatable experimentation). Nevertheless, no single platform is able to satisfy all possible scenarios, and so researchers often rely on different platforms to evaluate their ideas.

Given the huge diversity of available platforms, it is to be expected a big disparity in the way to carry out an experiment between one platform and another. Indeed, different platforms provide their own mechanisms to access resources and different tools to conduct experiments. These tools vary widely, for instance, to run a ns-3 simulation it is necessary to write a C++ program, while to conduct an experiment using PlanetLab nodes, one must first provision resources through a special web service, and then connect to the nodes using SSH to launch any applications involved in the experiment.

Mastering such diversity of tools can be a daunting task, but the complexity of conducting network experiments is not only limited to having to master different tools and services. Designing and implementing the programs and scripts to run an experiment can be a time consuming and difficult task, specially if distributed resources need to be synchronised to perform the right action at the right time. Detecting and handling possible errors during experiment execution also poses a challenge, even more so when dealing with large size experiments. Additionally, difficulties related to instrumenting the experiment and gathering the results must also be considered.

In this context, the challenges that NEPI addresses are manifold. Firstly, simplifying the complexity of running network experiments. Secondly, simplifying the use of different experimentation platforms, allowing to easily switch from one to another. Thirdly, simplifying the use of resources from different platforms at the same time in a single experiment.

The approach proposed by NEPI consists on exposing a generic API that re-

searchers can use to *program* experiments, and providing the libraries that can execute those experiments on target network experimentation platforms. The API abstracts the researchers from the details required to actually run an experiment on a given platform, while the libraries provide the code to automatically perform the steps necessary to deploy the experiment and manage resources.

The API is generic enough to allow describing potentially any type of experiment, while the architecture of the libraries was designed to be extensible to support arbitrary platforms. A consequence of this is that any new platform can be supported in NEPI without changing the API, in a way that is transparent to the users.

3.1 Experiment Description

NEPI represents experiments as graphs of interconnected resources. A resource is an abstraction of any component that takes part of an experiment and that can be controlled by NEPI. It can be a software or hardware component, it could be a virtual machine, a switch, a remote application process, a sensor node, etc.

Resources in NEPI are described by a set of attributes, traces and connections. The attributes define the configuration of the resource, the traces represent the results that can be collected for that resource during the experiment and the connections represent how a resource relates to other resources in the experiment.

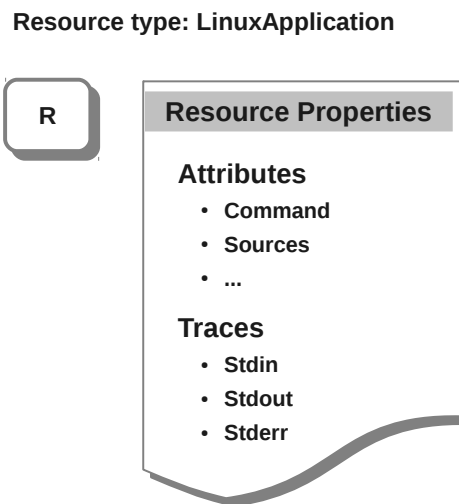


Figure 3.1: Properties of a resource of type LinuxApplication

Examples of attributes are a linux host host name, an IP address to be assigned to a network interface, a command to run as a remote application. Examples of traces

are the standard output or standard error of a running application, a tcpdump on a network interface, etc.

Resources are also associated to a type (e.g. a Linux host, a Tap device on PlanetLab, an application running on a Linux host, etc). Different types of resources expose different attributes and traces and can be connected to other specific types (e.g. A resource representing a wireless channel can have an attribute SSID and be connected to a Linux interface but not directly to a Linux host resource) Figure ?? exemplifies this concept.

There are two different types of connections between resources, the first one is used to define the *topology graph* of the experiment. This graph provides information about which resources will interact with which other resources during the experiment (e.g. application A should run in host B, and host B will be connected to wireless channel D through a network interface C). Figure ?? shows a representation of the concept of topology graph to describe the an experiment.




Figure 3.2: A topology graph representation of an abstract experiment

The second type of connections (called conditions to differentiate them from the first type) specifies the *dependencies graph*. This graph is optional and imposes constraints on the experiment workflow, that is the order in which different events occur during the experiment. For instance, as depicted in Figure ?? a condition on the experiment could specify that a server application has to start before a client application does, or that an network interface needs to be stopped (go down) at a certain time after the beginning of the experiment.

It is important to note, that the *topology graph* also defines implicit and compulsory workflow constraints (e.g. if an application is *topologically* connected to a host, the host will always need to be up and running before an application can run




Figure 3.3: A dependencies graph representation involving two applications resources in an experiment

on it). The difference is that the *dependency graph* adds complementary constraints specified by the user, related to the behavior of the experiment.

This technique for modeling experiments is generic enough that can be used to describe experiments involving resources from any experimentation environment (i.e. testbed, simulator, emulator, etc). However, it does not provide by itself any information about how to actually deploy and run an experiment using concrete resources.

3.2 Experiment Life Cycle

The Experiment Description by itself is not enough to conduct an experiment. In order to run an experiment it is necessary to translate the description into concrete actions and to perform these actions on the specific resources taking part of the experiment. NEPI does this for the user in an automated manner.

Given that different resources will require performing actions in different ways (e.g. deploying an application on a Linux machine is different than deploying a mobile wireless robot), NEPI abstracts the life cycle of resources into common stages associated to generic actions, and allows to plug-in different implementation of these actions for different types of resources. Figure ?? shows the three main stages of the network experiment life cycle, *Deployment*, *Control* and *Result (collection)*, and the actions that are involved in each of them.

In order to be able to control different types of resources in a uniform way, NEPI assigns a generic state to each of these actions and expects all resources to follow

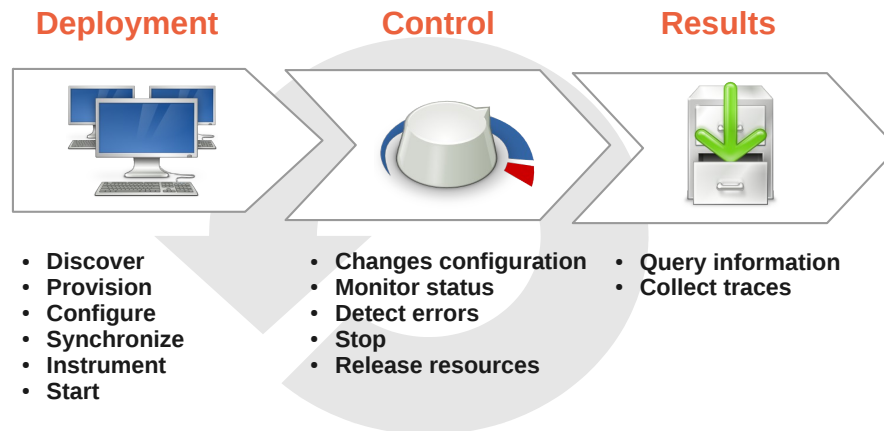


Figure 3.4: Common stages of a network experiment life cycle

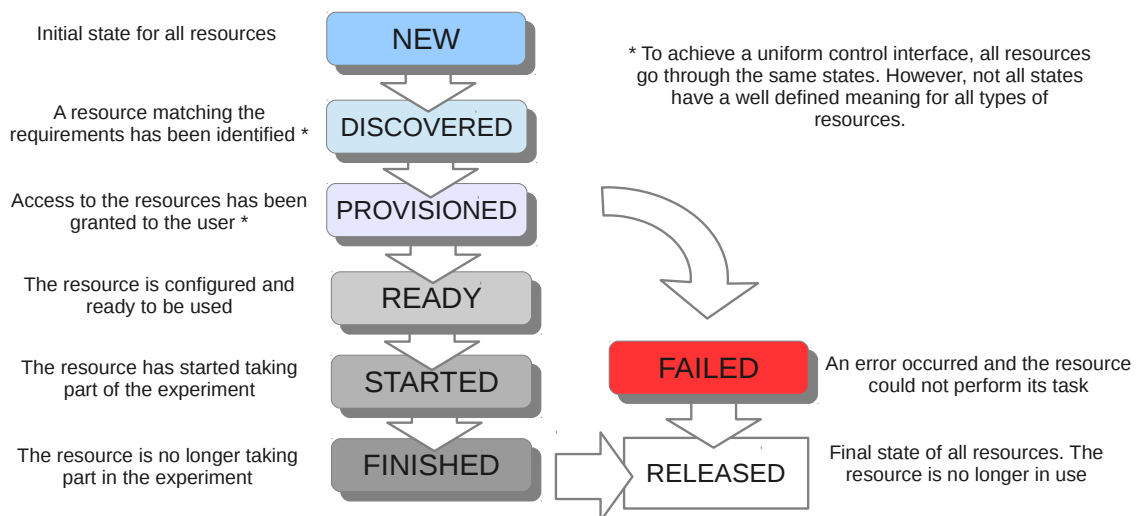


Figure 3.5: Resources state transitions

the same set of state transitions during the experiment life. The states and state transitions are depicted in Figure ??.

It is important to note that NEPI does not require these states to be globally synchronized for all resources (e.g. resources are not required to be all ready or started at the same time). NEPI does not even require all resources to be declared and known at the beginning of the experiment, making it possible to use an *interactive deployment* mode, where new resources can be declared and deployed on the fly, according to the experiment needs. This interactive mode can be useful to

run experiments with the purpose of exploring a new technology, or to use NEPI as a adaptive experimentation tool, that could change an experiment according to changing external conditions or measurements.

3.3 Resource Management: The EC & The RMS

The Experiment Controller (EC) is the entity that is responsible for translating the Experiment Description into a running experiment. It holds the *topology* and *dependencies* graphs, and it exposes a generic experiment control API that the user can invoke to deploy experiments, control resources and collect results.

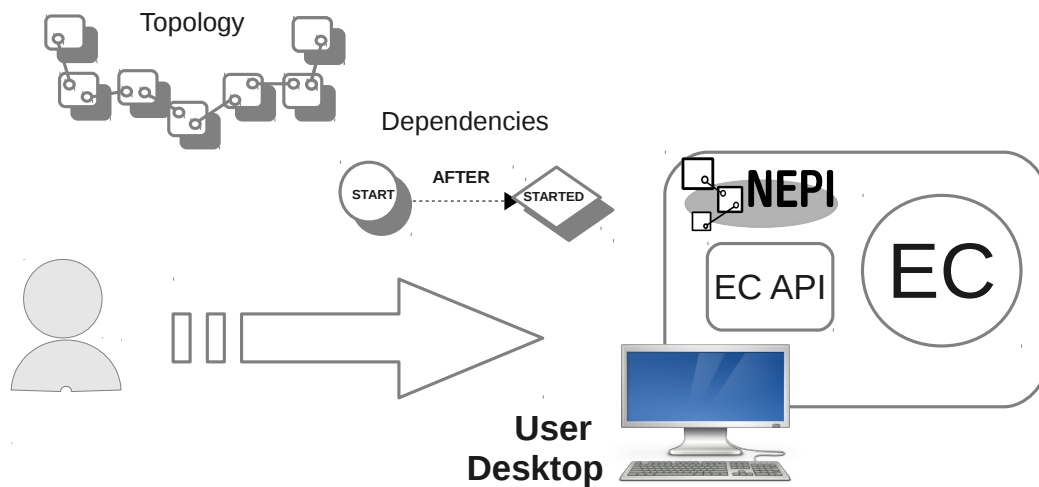


Figure 3.6: User interacting with the Experiment Controller

As shown in Figure ??, the user declares the resources and their dependencies directly with the EC. When the user requests the EC to deploy a certain resource or a group of resources, the EC will take care of performing all the necessary actions without further user intervention, including the sequencing of actions to respect user defined and topology specific dependencies, through internal scheduling mechanisms.

The EC is a generic entity responsible for the global orchestration of the experiment. As such, it abstracts itself from the details of how to control concrete resources and relies on other entities called Resource Managers (RM)s to perform resource specific actions.

For each resource that the user registers in the *topology graph*, the EC will instantiate a RM of a corresponding type. A RM is a resource specific controller and

different types of resources require a different type of RM, specifically adapted to manage them.

The EC communicates with the RMs through a well defined API that exposes the necessary methods (actions) to achieve all the state transitions defined by the common resource life-cycle. Each type of RM must provide a specific implementation for each action and ensure that the correct state transition has been achieved for the resource (e.g. upon invocation of the START action, the RM must take the necessary steps to start the resource and set itself to state STARTED). This decoupling between the EC and the RMs makes it possible to extend the control capabilities of NEPI to arbitrary resources, as long as a RM can be implemented to support it.

As an example, a testbed X could allow to control host resources using a certain API X, which could be accessed via HTTP, XMLRPC, or via any other protocol. In order to allow NEPI to run experiments using this type of resource, it would suffice to create a new RM of type host X, which extends the common RM API, and implements the API X to manage the resources.

Figure ?? illustrates how the user, the EC, the RMs and the resources collaborate together to run an experiment.




Figure 3.7: Resource management in NEPI

The ExperimentController API

4

- EC API
 - register resource
 - RM States (State transition DIAGRAM)
 - RM Actions
 - Attributes
 - The critical attribute
 - Traces (how to collect traces to the local repo, talk about the Collector RM)
 - deploy (interactive deployment)
 - workflow - register condition
- Resource Factory
 - Populate factory (happens automatically)
 - How to discover available rm types and their resources

Supported resources

5

5.1 Linux resources

- SSH
- The directory structure
- Linux Node (Clean home, etc)
- The api (run, execute, x11, etc)
- Traces and collection
- Linux Application
- LinuxPing, LinuxTraceroute, etc
- CCNx

5.2 Planetlab resources

- how to get an account
- The vsys system
- python-vsys
- TAP/TUN/TUNNEL
- Note on PL instability
- differences between PLE and PLC

5.3 OMF resources

- available OMF testbeds
- how to get an account
- the concept of resource reservation

6 Debugging

TODO