

**Cooperating Objects  
Network of Excellence**  
FP7-ICT-2007-2-224053  
[www.cooperating-objects.eu](http://www.cooperating-objects.eu)



---

# Common Abstractions for Testbed Federation

EW Identifier: D33-20091201-CONET\_D33  
Date: 2009-12-01

Author(s) and company: Vlado Handziski, Claudio Donzelli, Irina Antonova (TUB)

Work package/task: WP3  
Document status: final  
Confidentiality: public  
Keywords: testbed, federation, API, REST  
Abstract: Requirements analysis and initial specification of common abstractions for federating Cooperating Object testbeds.

## DOCUMENT HISTORY

Version	Date	Reason of change
1	2009-08-20	document created
2	2009-10-27	draft version for SC meeting
3	2009-11-27	final draft
4	2009-12-01	final version

# Table of Contents

---

<b>1</b>	<b>Introduction</b> .....	<b>6</b>
<b>2</b>	<b>Testbeds for Cooperating Objects</b> .....	<b>8</b>
2.1	Design Challenges .....	8
2.2	Testbed Architectures .....	9
2.3	Testbed Directory .....	11
2.4	CONET Testbed Federation Survey .....	18
2.5	CONET Testbed Federation Prototyping Testbeds .....	39
<b>3</b>	<b>CONET Testbed Federation Platform</b> .....	<b>40</b>
3.1	Design Principles .....	40
3.2	System Architecture .....	41
3.3	Core Abstractions .....	48
<b>4</b>	<b>CONET Testbed Federation Platform APIs</b> .....	<b>57</b>
4.1	Common Features .....	57
4.2	Resource Model .....	61
4.3	Requests and Responses .....	73
4.4	Invocation Examples .....	78

# Acronyms

**AAA** Authentication, Authorization and Accounting

**ADC** Analog-to-Digital Converter

**AODV** Ad hoc On-Demand Distance Vector

**API** Application Programming Interface

**COTS** Commercial Off The Shelf

**CO** Cooperating Object

**CPU** Central Processing Unit

**CTFS** CONET Testbed Federation Server

**CTF** CONET Testbed Federation

**CTP** Collection Tree Protocol

**DC** Direct Current

**DSL** Digital Subscriber Line

**EMI** Electro-Magnetic Interference

**GENI** Global Environment for Network Innovations

**GPS** Global Positioning System

**GUI** Graphical User Interface

**HTTP** Hypertext Transport Protocol

**IBSS** Independent Basic Service Set

**IC** Identity Consumer

**IP** Identity Provider

**JSON** JavaScript Object Notation

**LED** Light Emitting Diode

**MAC** Medium Access Control

**MANET** Mobile Ad hoc Network

**MCU** Microcontroller Unit

**MIB** Management Information Base

**MIME** Multipurpose Internet Mail Extensions

**MIT** Massachusetts Institute of Technology

**NAS** Network Attached Storage

**OLSR** Optimized Link State Routing

**PCI** Peripheral Component Interconnect

**PC** Personal Computer

**PDA** Personal Data Assistant

**PHY** Physical Layer

**PoE** Power over Ethernet

**RAM** Random Access Memory

**REST** Representational State Transfer

**RF** Radio Frequency

**RPC** Remote Procedure Call

**SC** Service Consumer

**SNMP** Simple Network Management Protocol

**SOA** Service Oriented Architecture

**SP** Service Provider

**SSH** Secure Shell

**SSO** Single Sign-On

**SUT** System Under Test

**TA API** Testbed Adaptation API

**TCP** Transmission Control Protocol

**TF API** Testbed Federation API

**UI** User Interface

**URI** Universal Resource Identifier

**URN** Universal Resource Name

**URL** Universal Resource Locator

**USB** Universal Serial Bus

**UUID** Universally Unique Identifier

**VSG** Vector Signal Generator

**WSN** Wireless Sensor Network

**XML** Extensible Markup Language

# 1 Introduction

The design, implementation and evaluation of Cooperating Object (CO) systems is a challenging task that is further complicated by their distributed and heterogeneous nature and the tight coupling with the environment.

During the early design stages, valuable insight about the system can be gained using analytical modeling or simulation. In order to maintain tractability of the models, however, the designer is frequently forced to make simplifying assumptions about the behavior of the system components and their interaction with the environment. The associated loss of realism makes these approaches less suitable for the more advanced design stages, when the evaluation of the system performance, error resilience and other nonfunctional properties necessitate the use of real hardware, realistic environments and realistic experimental setups.

Testbeds provide convenient middle ground between simulation and full deployment on the realism axes. They allow for rigorous and controlled experimentation with the System Under Test (SUT). But similarly to full deployments, they lock the evaluation to one particular environment making it hard to differentiate between the intrinsic properties of the SUT and the effect of the specific features and external influences present at a given testbed site. One way of decoupling these influences is to cross-validate the functional and non-functional behavior of the SUT under various conditions as provided by different testbeds.

For example, Table 1.1, taken from [GFJ<sup>+</sup>09], shows a summary of an experimental evaluation of the functional properties of CTP on 12 different testbeds. Experimental studies like this one are becoming indispensable part of performing credible scientific research in the area of CO. Un-

Testbed	Platform	Nodes	Physical size [m <sup>2</sup> or m <sup>3</sup> ]	Degree		PL	Cost	PL/Cost	Churn
				Min	Max				
Tutornet (16)	Tmote	91	50 × 25 × 10	10	60	3.12	5.91	1.9	31.37
Wymanpark	Tmote	47	80 × 10	4	30	3.23	4.62	1.43	8.47
Motelab	Tmote	131	40 × 20 × 15	9	63	3.05	5.53	1.81	4.24
Kanseia	TelosB	310	40 × 20	214	305	1.45	-	-	4.34
Mirage	Mica2dot	35	50 × 20	9	32	2.92	3.83	1.31	2.05
NetEye	Tmote	125	6 × 4	114	120	1.34	1.4	1.04	1.94
Mirage	MicaZ	86	50 × 20	20	65	1.7	1.85	1.09	1.92
Quanto	Epic-Quanto	49	35 × 30	8	47	2.93	3.35	1.14	1.11
Twist	Tmote	100	30 × 13 × 17	38	81	1.69	2.01	1.19	1.01
Twist	eyesIFXv2	102	30 × 13 × 17	22	100	2.58	2.64	1.02	0.69
Vinelab	Tmote	48	60 × 30	6	23	2.79	3.49	1.25	0.63
Tutornet (26)	Tmote	91	50 × 25 × 10	14	72	2.02	2.07	1.02	0.04
Blazeb	Blaze	20	30 × 30	9	19	1.3	-	-	-

Table 1.1: CTP evaluation on multiple testbeds published in [GFJ<sup>+</sup>09]. The results illustrate how important protocol performance parameters like number of transmissions per successful delivery (Cost), average path length in hops (PL), or parent change rate (Churn) are influenced by the different testing environments.

fortunately, their realization is currently accompanied by significant overheads in configuring the experiments and collecting the results on the individual testbeds, since easy experiment migration is hindered by a lack of common management, experiment specification and control infrastructure.

The goal of the CONET Testbed Federation (CTF) Task is to address some of these roadblocks by developing a software platform that will enable convenient access to the experimental resources of multiple testbeds organized in a federation of autonomous entities.

Due to the specific nature of the wireless medium, we are primarily focused on facilitating experiment migration across the federation members, and not in combining the federation resources into a single “virtual” testbed. Our federation platform is built on top of a set of common abstractions for authentication and authorization; resource discovery and reservation; and experiment specification and control. The integration APIs are light-weight, scalable and extensible and aim to preserve high levels of autonomy for the member testbeds in the federation.

With the work on this task, CONET is contributing to the ongoing effort in the community towards federated testbeds solutions in this and related areas, as exemplified by the related activities under the FIRE (EU) and the GENI (USA) initiatives.

This document establishes the basic requirements for federating CO testbeds and presents the general design of the CTF platform. In the next chapter we first provide analysis of the main functional and architectural features of CO testbeds, an extensive directory of existing CO testbeds and their capabilities, as well as the results from the CONET Testbed Federation (CTF) Survey that collected information on the existing testbed capabilities of the CONET partners. This information then serves as basis for the requirements analysis driving the design of the federation platform.

In Chapter 3 we present the main features of the CTF platform. We start with the specification of the system architecture and the main characteristics of the Representational State Transfer (REST) architectural style that has been selected as guiding principle for designing the interactions between the major system components. In the second part of the chapter, following the REST approach, we document the CTF abstractions through the identification of the main resources, their representations and the possibilities for modification of the resource state through the application of the Hypertext Transport Protocol (HTTP) method set.

## 2 Testbeds for Cooperating Objects

In this chapter we provide an analysis of the architectural features of the individual CO testbeds as important input for the design of the CTF platform presented in the next chapter.

We start the discussion with a summary of the major challenges in designing CO testbeds and the main differentiating aspects among the existing CO testbed architectures. This analysis is complemented by a comprehensive directory of prominent CO testbeds which highlights their SUT resources and the architectural setup of their control and management planes.

In the second part, we turn our attention to the testbed resources of the partners in the CONET project. We present the results from the internal survey that we used to establish the feature set of the candidate base for the CONET Testbed Federation. Finally, we provide more detailed description of the selected CONET testbeds that will be used as testing ground for the CTF prototype.

### 2.1 Design Challenges

The primary goal of any CO testbed is to support the design, implementation, testing and evaluation of applications and protocols without forcing the investigator to make artificial assumptions about system components or the system environment (as often needed in analytical and simulation work). The design of testbeds for COs, is facing a number of challenges and constraints, some of them being:

- COs are large-scale distributed systems where each node can observe the network state only locally. For network debugging and testing purposes, however, often a global view on the network is required.
- The protocols used in these systems tend to be application specific, and there is (yet) no unified set of protocols or standardized equipment, nor are there any established mechanisms and tools for debugging and testing CO protocols and applications.
- COs are not only influenced by the built-in behavior of its software and hardware components, but also by external stimuli, since these systems are ultimately built to observe and to modify the state of the physical world.
- COs must handle ever-changing topologies because of node mobility or nodes dying on energy depletion or new nodes being added. They must also live with other limitations like scarce bandwidth, small amount of per-node memory, etc. The testbed infrastructure must be capable to track this dynamics.
- Many COs applications are not homogeneous and rely on nodes with different capabilities (in terms of processing power, energy supply and memory). This results in very specific distribution of functions like high-performance computation, data aggregation, data storage,

interfacing with external networks, etc., that also needs to be replicated in the testbed environment.

A successful testbed architecture needs to accommodate these specifics, in a scalable and cost-efficient way. To facilitate the design of the SUT, for example, it should be possible to implement different network architectures in the testbed, from which the designer can select the best one for his application. To support the implementation/test/debugging phase, functionality like node (re-)programming as well as collection, processing and displaying of debug data are needed. For evaluation purposes, users should have means to precisely and reproducibly control network topologies and to inject node faults. In addition to these development-oriented goals, the testbed should also be scalable enough to support large-scale SUT deployments with more than just a few dozen nodes while keeping the cost and management overheads at an acceptable level.

## 2.2 Testbed Architectures

In the following we provide an overview of the most important differentiating features among the current testbed frameworks and instances.

### 2.2.1 Main Features

- *Hardware architecture*: The testing functions can be organized in a single, flat architecture, or in several different, hierarchically arranged node classes. Since an application can also have one, two or more hardware tiers, the quality of the testing coverage crucially depends on the alignment between the testing system and SUT configurations.
- *Isolation of the communication channel*: Does the traffic generated by the testing functions (observed states, control traffic) has to share communication channels with the application traffic? If so, we call this *inband-signaling*, otherwise we speak about *out-of-band-signaling*. For example, a single-tier network where testbed infrastructure nodes are used entirely for testing functions, is regarded as using out-of-band signaling.
- *Isolation of the computational resources*: Another precious resource besides bandwidth are the computational and memory resources of SUT and testbed infrastructure nodes. When a node runs exclusively application functions or runs exclusively testing functions, we denote this as *processor/memory separation*. Passive snooping, for example, would be an example of processor/memory separation. If both application and testing functions run on the same processor/memory, we denote this as *processor/memory sharing*.
- *Software architecture*: The test-related software can function in a centrally organized manner (even in cases where the test hardware possesses multiple levels), or it can be distributed over different hierarchical levels. For example, leaf testbed nodes can generate a stream of observed information, which is filtered/preprocessed in higher levels of the testbed hierarchy.
- *Concurrency and vitalization*: In some testbeds it is possible to partition the network between different applications (and related observers). This partitioning can be truly orthogonal or partial where some resources have to be shared (e.g. radio resources like frequencies / codes).

- *Naming and addressing*: For maximal flexibility, the testbed architecture should decouple between the addressing of the node in the SUT space and its naming in the testbed architecture. Testbeds also differentiate in the naming schemes that are exported to the testbed user.
- *Active control*: Virtually all testbeds will offer functions to passively observe the state of the SUT. However, they can differ in their ability to modify the state of the SUT during the test run. Examples of such control are explicit mobility commands, control of the power supply, injection of packets, emulation of sensor excitations, etc.
- *Command synchronization and time stamping*: Is the testbed capable of performing (possibly distributed) synchronized control actions and time stamping? Synchronization can for instance be useful to ensure that SUT nodes initialize at (almost) the same time, or to insert distributed failures in a tightly controlled fashion.
- *Self-configuration*: The testbeds differentiate in the levels of support that the management infrastructure provides to the testbed operators in the testbed provisioning phase, i.e. by automatic updating of the testbed configuration when adding new components to the testing system (detection of node attach/detach events, etc.)

### 2.2.2 Experiment Specification and Control

- *Execution model*: Testbeds differ in the models of test execution that is supported. In general, two levels of access are provided: an online (interactive) mode, that allows to observe and modify the SUT state during runtime; or offline (batch/scheduled mode) where the collected data (logs, trace dumps, etc.), are made available to the user only after the test run has been completed. Some testbeds also allow both online and offline modes.
- *Configuration and control API*: Testbeds also differentiate in the types of interfaces that are offered to users for configuration and control of the test runs. Possible choices can be, configuration languages, remote procedure calls, remote method invocation, web services, etc.
- *External configuration*: This feature deals with the possibility to configure the SUT nodes and to supply them with information which otherwise has to be obtained by running protocols within the application?. An example would be a means to configure an application node with its geographic position as derived from the testbed topology, instead of forcing the node to execute a localization protocol. Using testbed supported time synchronization of the SUT nodes is another typical example.
- *Data access API*: The data access APIs define how the test run data is made available to the end users. For example, they can be stored in a relation database, written into plain files, custom XML files, offered as an SNMP MIB tree, etc.
- *User customization*: This feature evaluates the degree of customization that the testbed provides to its users. Can users plug own filters, own control functions or other code into the testing functions to modify its behavior? Which components of the testing software can be customized and what is the associated programming model?

### 2.2.3 Access Policy, Economical and Ethical Issues

- *Access policy*: Testbeds differentiate in their access policies in terms of the user groups granularity, the priority between the different groups, and the levied access fees.
- *Level of openness*: The more information the user has about the software and the hardware features of a given testbed, the higher is his confidence in the obtained results. This testbed characteristic evaluates if the hardware and software architecture of the testbed made available to external parties. What is the licensing model, and can external parties easily replicate the testbed setup in order to validate the obtained results?
- *Privacy protection*: For management reasons or for additional deployment realism, testbeds are frequently located in environments where the SUT sensors are in position to extract privacy sensitive information. Does the testbed contain provisions to protect the private sphere of individuals who are monitored with or without consent?

## 2.3 Testbed Directory

In this section we provide a summary of the main features of several prominent CO testbeds. This testbed directory is continuously maintained and is included in the CONET Roadmap document.

### 2.3.1 APE

The *Ad hoc Protocol Evaluation (APE)* testbed [LLN<sup>+</sup>02] from *Uppsala University* is a full-scale testbed (no artificial attenuation of the RF signals) for comparative study of different MANET protocols. Node mobility is achieved through choreographed movement of human volunteers carrying the laptops containing the SUT wireless cards. The experimental work is supported by a set of logging and visualization tools. The trace collection is performed in a distributed fashion, relying on timestamps and an off-line aggregation step to recover the global ordering of the events. The APE software framework is publicly available under open-source license.

### 2.3.2 Casino Lab

The *Casino Lab* WSN testbed [cas] at the *Colorado School of Mines* consists of 52 Tmote Sky nodes, hung from the ceiling of a large open industrial space with concrete walls, pipes, ducts and fluorescent lighting. The dimensions of the room is 24.4m × 12.30m, and the nodes are deployed in a 4 × 13 irregular grid. The nodes are connected via USB to 26 Tmote Connect Ethernet bridges providing a wired out-of-band channel for control. The *TOSSIM Live* extension [Met07] to the TOSSIM simulator, allowing execution of simulations in real-time and their interaction with real testbed nodes was originally developed and validated on the Casino Lab. The testbed is not publicly available.

### 2.3.3 CitySense

*CitySense* [UT] is an urban-scale WSN testbed developed by *Harvard University* and *BBN Technologies*. It consists of a set of nodes fixed to rooftops and streetlights (currently 30 outdoor nodes,

with a final target of 100 nodes) deployed in Cambridge, MA. The computational unit on the nodes comes in two configurations, one with 533 MHz CPU and 258 MB RAM or with 233 MHz CPU and 128 MB RAM. The nodes run FreeBSD 7 as operating system. The CitySense nodes support an array of sensors relevant for the outdoor urban setting: weather, CO<sub>2</sub>, noise pollution, etc. Each node has two radios. The communication substrate is formed by mesh routing using OLSR on a Ubiquiti SR9 (900 MHz) radio interface cards. The Wistron CM9 (2.4/5 GHz 802.11a/b/g) radio, is not used by the testing infrastructure and is available for SUT applications. The management of the testbed is performed using a custom software framework called *CityMD*.

### 2.3.4 CMU-DSR

One of the original full-scale MANET testbeds, the *Carnegie Mellon University's* testbed for evaluation of the Dynamic Source Routing (DSR) protocol consisted of 5 mobile and 2 stationary nodes. The wireless nodes were comprised by an IBM Thinkpad laptop equipped with 900 MHz WaveLAN radio card. Mobility effects were evaluated by driving the mobile units in rented cars in an outdoor area with rough dimensions of 700m × 300m. The localization of the units was performed via GPS. One of the fixed nodes was used as gateway, connecting the test network back to the Field Office via a 2.4 GHz point-to-point link.

### 2.3.5 DES-MESH

The *DES-MESH* is a hybrid wireless mesh and sensor network testbed being developed at the *Freie Universität Berlin* geared towards long-term studies [GBS08]. It currently consists of 35 hybrid nodes installed in an office setting spanning three floors. The testbed architecture is organized in three tiers: backbone mesh routers, mesh clients and sensor nodes. The mesh routers are equipped with 500 MHz AMD Geode LX800 CPUs with 256 MB of RAM, and have three IEEE 802.11 cards attached via USB hubs. The sensor nodes have 60 MHz ARM 7 cores and Chipcon CC1100 transceivers in the 868 MHz band. The testbed management is realized by a combination of SSH-supported remote command execution and SNMP services. The experiment configuration and control is facilitated by a domain specific language called DES-CRIPT based on XML.

### 2.3.6 DOME

The *Diverse Outdoor Mobile Testbed (DOME)* by the *University of Massachusetts Amherst* is one of the longest-running urban-scale MANET testbeds. It is comprised of three categories of nodes: The DieselNet is formed by 40 transit buses with nodes containing a Hacom OpenBrick 1 GHz computing board, a GPS receiver and several different radio cards: a 802.11 a/b/g mini PCI card, 802.11g wireless access point, a 3G modem and 900 MHz RF modem; The "throwboxes" have similar setup but use batteries recharged by solar cells. They can be mounted on bicycles and can act as additional routers improving the connectivity of the DieselNet buses; The third group of nodes is comprised by 26 Cisco 1500-series WiFi access points, mounted on different buildings and light poles, that form a mesh network. The management software is modular and provides services like link management, remote updating, logging and maintenance monitoring.

### 2.3.7 DSN

The *Deployment Support Network (DSN)* is a testbed framework developed at *ETH Zürich* [DBK<sup>+</sup>07], that leverages a secondary multi-hop WSN optimized for connectivity and reliability as a testbed backbone. The DSN-nodes forming this backbone network are in turn connected to the SUT nodes via custom wired interfaces. Currently supported SUT node platforms include the BTnode, TinyNode, Tmote Sky and A80. The testbed backbone is used for SUT image file distribution, for transfer of logging and debug data, and for sending direct commands to the SUT nodes. The operation of the testbed is controlled by a DSN-server that exports the DSN-services via XML-RPC and web based interfaces towards the testbed user. The instance of the DSN framework at ETH Zürich uses the BTnode platform and its Bluetooth radio for the backbone network. The current configuration of the testbed consists of 30 backbone nodes and 30 Tmote Sky and TinyNode SUT nodes.

### 2.3.8 EWANT

The *Emulated Wireless Ad Hoc Network Testbed (EWANT)* [SBBD03], developed at the *University of Colorado at Boulder*, is a reduced-scale MANET testbed with emulated RF environment using in-line attenuation and RF multiplexing. Mobility is simulated by discrete switching between different antennas connected to the outputs of the 1:4 RF multiplexers attached to the wireless cards.

### 2.3.9 iWWT

The *Illinois Wireless Wind Tunnel (iWWT)* [VBV<sup>+</sup>05] is a reduced-scale testing environment for wireless networks implemented in an electromagnetic anechoic chamber at the *University of Illinois at Urbana-Champaign*. The main aim of the testbed is to create a realistic scaled version of the wireless environment maintaining full control over all relevant parameters that affect the performance of the wireless network like obstructions, interferers, etc. Mobility is supported by placing the wireless hosts (laptops, PDAs, sensor nodes) on top of remotely controlled cars. The scaled wireless environment is constructed by combining the effects of several building blocks: Power control module, Multipath module, Doppler module and Scattering Module. Despite these efforts for complete control of the RF environment, repeatability of small-scale experimental results remains elusive due to intrinsic randomness in the evaluated protocols and object positioning errors [ZV08].

### 2.3.10 Kansei

The *Kansei* testbed at *The Ohio State University* [AER<sup>+</sup>06] has been initially developed as a testing facility for the middleware services for the final demonstrator of the *Extreme Scale Wireless sensor Networking (ExScal)* project [ARE<sup>+</sup>05]. The ExScal demonstration is one of the largest hybrid deployments of sensor and wireless mesh networks ever attempted, with more than 1000 sensor nodes and more than 200 wireless mesh nodes distributed in a 1.3km × 300m area. The Kansei testbed is a testbed environment that replicates this heterogeneous architecture at a reduced scale. Its stationary array originally consisted of 210 dual nodes (a combination of one Extreme Scale Stargate (XSS) wireless mesh node and one Extreme Scale Mote (XSM) node) placed on a 15 × 14 rectangular grid with about 1 m spacing. The XSS nodes have IEEE 802.11b wireless cards, and the XSM nodes, derivatives of the UC Berkeley's Mica family, operate in the 916 MHz band. Recently, 150 nodes have been upgraded with Tmote Sky boards. In addition to the

stationary array, Kansei also has a portable array of 50 Trio sensor nodes and a mobile array of 5 robots from Acorname Inc.. The software architecture of the testbed is organized around the *Kansei Director* that provides interfaces towards the basic services of the testbed like experiment scheduling, deployment, platform monitoring and management as well as creation and management of testbed arrays and configurations. The Kansei testbed is designed for shared usage, and has open access policy for members of the research community.

### 2.3.11 MiNT

The *Miniaturized Wireless Network Testbed (MiNT)* [DRSC05] at the *Stony Brook University* consists of eight mobile nodes roaming in an  $3.66\text{m} \times 1.83\text{m}$  area. The mobile nodes are built from COTS hardware: Routerboard 230 mini PCs with 1-3 Atheros IEEE 802.11a/b/g cards, placed on top of iRobot's Roomba as the mobility platform. One of the wireless cards on each node, operated in RF monitoring mode, is dedicated to collecting traces that are transferred to a central node where they can be visualized in real time. The output from the remaining wireless cards is connected to low-gain external antennas via fixed signal attenuators providing about 60 dBm attenuation, limiting the communication range to about 0.6 m. The custom control GUI enables convenient node configuration, editing and execution of traffic generation scripts, mobility scripts and fault injection scripts. The GUI performs merging of the traces collected by the different nodes and extraction of different network statistics. MiNT also supports hybrid execution of unmodified ns-2 simulations over the MAC and PHY layers of the real testbed nodes.

### 2.3.12 Mirage

*Mirage* [CBA<sup>+</sup>05] is a testbed management system developed by the *Intel Research Berkeley (IRB)* that applies the concepts of microeconomic resource allocation to the problem of allocating nodes in a sensor network testbed. Users submit bids that are specifying their interest in terms of the nodes and the time they would like to be granted access to, combined with the price they would be willing to pay. The system periodically runs a sealed-bid auction to determine the successful bids based on the demand/availability, while aiming to maximize the aggregate utilization of the testbed. The Mirage framework is used to manage a 100 MicaZ and 50 Mica2Dot node testbed at IRB premises.

### 2.3.13 Mobile Emulab

The *Mobile Emulab* [JSF<sup>+</sup>06] is a robotic wireless and sensor network testbed developed by the *University of Utah* leveraging their widely used *Emulab* platform for running network testbeds. The testbed is comprised by four Garcia robots from Acorname, each carrying a Stargate node that is interfaced with a Mica2 node. The mobile nodes are roaming in an area of about  $8\text{m} \times 3.5\text{m}$ . The tracking and the identification of the nodes is handled by a vision-based tracking system, using six ceiling-mounted video cameras aimed down towards the floor. A central control daemon is responsible for plotting the movement paths of the robots, so that they can safely reach the user-specified end positions, while maneuvering around any static and dynamic obstacles encountered during the motion. In addition, the testbed is equipped with 25 static Mica2 nodes arranged on the ceiling in a rough 2 m grid and on the walls near the floor. The testbed management software

uses the standard Emulab services to provide a batch queued or interactive first-come, first-served usage through a web-based, GUI-driven or programmable XML-RPC user interface.

### 2.3.14 MoteLab

*MoteLab* [WASW05] is a very popular testbed solution from *Harvard University*. In its original design, the testbed was comprised from Mica2 nodes, each connected to Ethernet backbone via dedicated Crossbow interface boards, providing TCP forwarding for the serial ports. The current deployment is one of the largest publicly accessible testbeds with 190 Tmote Sky nodes deployed over 3 floors of Harvard's Engineering building. The SUT nodes are connected to the testbed backbone via Tmote Connect units. The testbed server provides a web interface that lets users monitor the status of the testbed and register jobs. The system uses a quota system to limit the time that each user has at its disposal for the outstanding test jobs. At one given time, only a single user has control over the complete resources of the testbed. The jobs are started by a job scheduler that configures the SUT according to the job description (installing and configuring the SUT images, etc.) and starts logging of the SUT output to a local database. The users can also get raw data access by connecting to the TCP forwarded serial ports of the SUT nodes during the assigned job slot.

### 2.3.15 Motescope

The *Motescope* testbed [mot] at the *University of California, Berkeley* is an update of the *sMote* testbed installed in the Soda Hall. The original 78 Mica2Dot nodes in *sMote* have been replaced with MicaZ nodes in *Motescope*. The testbed provides convenient web interface for configuration and control of the experiments. The testbed has open access policy for the members of the academic research community.

### 2.3.16 ORBIT

The *Open Access Research Testbed from Next-Generation Wireless Networks (ORBIT)* [OSSS05] at *Rutgers University* is massive indoor grid of 400 wireless nodes each equipped with two IEEE 802.11a/b/g mini-PCI cards, deployed in a 20 by 20 grid with 1 m spacing. The testbed provides full remote control over the nodes during the assigned time slots, including installation of custom system images. The topology control is performed by a topology generator that leverages a wired remote control infrastructure to switch some of the nodes on or off. Similarly, mobility is simulated by transferring the state of a "mobile" node from one grid node to another. ORBIT provides extensive libraries for collecting, analyzing and accessing measurement data.

### 2.3.17 Omega

The *Omega* testbed [ome] is another testbed at the *University of California, Berkeley*. It consists of 28 TelosB nodes, connected via daisy-chained USB hubs to the central control server. This wired back-channel is used for powering, programming and debugging of the SUT nodes. The testbed has open access policy for the members of the academic research community.

### 2.3.18 Re-Mote

The *Re-Mote* framework [Dat] developed by the *Datalogisk Institut (DIKU)* at the *University of Copenhagen* is comprehensive testbed management suite with four main software components: mote control infrastructure, database scripts, web services and GUI client. The database component uses the MySQL relation database and tracks the static and dynamic state of the testbed through the mote, mote host site and user session models. The core of the architecture is formed by the mote control infrastructure that is implemented in C++, separated in a low-level Mote Control Host (MCH) and Mote Control Server (MCS) parts. The web services component provides a loose interface between the testbed services and the clients. The Java-based client interacts with the Java web services component and the MCS allowing users to authenticate to the testbed and to program and monitor the SUT nodes. The Re-Mote framework has been tested in a deployment at DIKU with 36 Freescale DIG528-2 development boards as SUT nodes connected to 5 host PCs.

### 2.3.19 RoofNet

The *RoofNet* testbed [BABM05] at the *Massachusetts Institute of Technology* is an experimental multihop IEEE 802.11b Internet access network comprised of nodes positioned on the roofs of various buildings in Cambridge, Massachusetts, covering an area of about 4 square kilometers. In the original design, each node was built out of a mini PC and a 802.11b card connected to a roof-mounted omnidirectional antenna. The cards are operated in simplified IBSS ad-hoc mode which does not use beacons. The RoofNet routing protocol *Srcr*, implemented in Click [MKJK99], uses a combination of link-state and on-demand mechanisms and selects routes that minimize the estimated transmission time. The RoofNet deployment demonstrated that urban mesh-based Internet access networks can provide sustained DSL-level performance. The original software framework for RoofNet was open sourced and serves as basis for many RoofNet replicas like the Berlin Roof Net [Hum] and the NetEquality Portland deployment [Net]. The RoofNet technology was commercialized and is being further developed by Meraki [mer].

### 2.3.20 Sensei

*Sensei* [RHF<sup>+</sup>08] is a nomadic testbed that can be used in a lab, for in-situ experiments and as prototype deployments. Physical sensors are normally attached to a Linux host consisting of a stationary computer, a laptop, a broadband router or a Personal Data Assistant (PDA) but there is also a cellphone implementation that enables using a cellphone as sensor proxy or as a sensor itself. The testbed supports usage mobile sensor nodes, carried around on random or predefined paths by either humans or robots. Control and management of the testbed can be handled remotely via a GUI that is highly customizable for different purposes. The nomadism supports testing the same experiment in different locations or at different testbed installations. The support for repeatable mobility enables repeating similar tests including mobile nodes multiple times, and the highly flexible GUI makes it possible to create experiment-specific visualizations of testbed activities.

### 2.3.21 Tutornet

*Tutornet* [tut] in the Ronald Tutor Hall at the *University of Southern California* is a tiered testbed with 13 node clusters, formed by a Crossbow stargate cluster-head, connected to several mote-class

nodes via USB cables. The stargate tier uses the EmStar development platform [GSR<sup>+</sup>04] and operates as a multihop wireless network, using AODV-based routing over IEEE 802.11b links. The mote tier is currently populated with 81 Tmote Sky and 13 MicaZ nodes. The testbed has open access policy for the members of the academic research community.

### 2.3.22 WASAL

The WASAL testbed [dL] at the *École Polytechnique Fédérale de Lausanne* consists of 25 TinyNode nodes connected to a wired testbed back-channel via custom serial-to-ethernet devices that act as passive communication bridges and range extenders. The WASAdmin management tool uses an XML-based configuration language and concurrently executes a separate shell instance and script parser for each target node in the SUT.

### 2.3.23 WINTeR

The *Wireless Industrial Sensor Network Testbed for Radio-Harsh Environments (WINTeR)* [SMP<sup>+</sup>08] is a testbed facility for the Canadian *Petroleum Applications of Wireless Systems (PAWS)* project. It aims to replicate the harsh RF conditions of an offshore oil platform, where the dense piping and other large metallic structures create complex multipath effects, combined with strong noise and interference and challenging environmental conditions. To this end, a platform mockup can be crated out of six modules each with dimensions of 2.44m × 1.83m × 2.74m. The mockup includes typical industrial metal structures like beams, pipes and tanks. The testbed uses a customized node platform consisting of a mote, DC power source, embedded CPU, micro power meter and programmable attenuator placed in an industrial enclosure. The emulation of the RF environment is further supported by an VSG-based EMI generator. The software architecture of the testbed is based on a modified version of Harvard's MoteLab suite.

## 2.4 CONET Testbed Federation Survey

In this section we present the results from the survey that we used to establish the feature set among the existing CONET testbeds as potential candidates for the future CONET Testbed Federation (CTF). The selected questions cover a range of hardware, software and policy issues that are of interest for the CONET Testbed Federation (CTF) Task 3.2.

The answers collected by the survey form the basis for the CTF platform requirements analysis. They were also used to compile a single point of reference for the existing testbed resources at disposal of CONET partners, which is published on the CONET web site.

Table 2.1 lists the testbeds covered in the survey, their hosting institution and the public web site (where available).

Testbed Name	Institution	Public Website
DSN	ETHZ	<a href="http://tec-pc-btnode.ethz.ch/dsngui/index.php?act=GetDSN">http://tec-pc-btnode.ethz.ch/dsngui/index.php?act=GetDSN</a>
SOCRADES	SCHNEIDER	...
Piaggio	UNIPI	<a href="http://www.piaggio.ing.unipi.it/">http://www.piaggio.ing.unipi.it/</a>
Powerbench	TUD	<a href="http://www.tudelft.nl">http://www.tudelft.nl</a>
NUIG-DERI	NUIG	<a href="http://www.deri.ie/">http://www.deri.ie/</a>
TWIST	TUB	<a href="http://www.twist.tu-berlin.de">http://www.twist.tu-berlin.de</a>
UBonnTestbed	UBONN	<a href="http://iv.cs.uni-bonn.de/">http://iv.cs.uni-bonn.de/</a>
IoT@SAP	SAP	<a href="http://www.sap.com">http://www.sap.com</a>
SICS-TB	SICS	<a href="http://www.sics.se">http://www.sics.se</a>
ART-WiSe	ISEP	<a href="http://artwise.cister-isep.info/searchrescue.php">http://artwise.cister-isep.info/searchrescue.php</a>
VSTL	BRTE	<a href="http://www.boeing.com/companyoffices/aboutus/brief/pw.html">http://www.boeing.com/companyoffices/aboutus/brief/pw.html</a>

Table 2.1: CONET Testbeds participating in the survey


In the rest of the section we provide a summary of the responses to the individual questions, grouped in eight categories. Each item consists of the original survey question, graphical and tabular summary of the responses <sup>1</sup>, followed by a short commentary on the impact that the information has on the desing of the CTF platform.

### 2.4.1 General Information

#### Operational Status

*Please specify the operational status of the testbed...*

#### Responses



<sup>1</sup>Some questions allow for multiple answers, so the percentage points can add up to more than 100 %.

Option	Number	Percentage
Planning phase	3	27 %
Set-up phase	2	18 %
Stable and running	6	55 %


### Comment

The responses confirm that testing facilities play an important role in the activities of all CONET partners. It shows that a sufficient testing technology base already exists that can both contribute towards and benefit from the work in the CTF Task.

### Technology Focus

*Please specify the technological focus of the testbed...*

### Responses



Option	Number	Percentage
Wireless Sensor Networks	10	91 %
Autonomous Robots	2	18 %
Other	4	36 %

### Comment

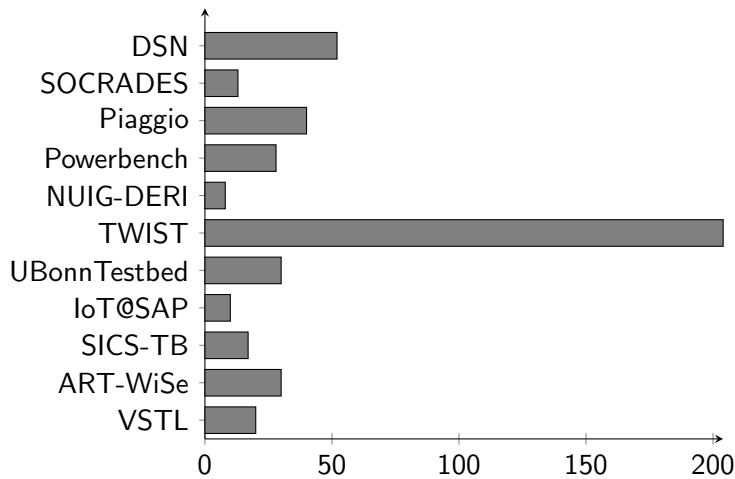
The responses show that the majority of the existing testbed resources in CONET target the area of Wireless Sensor Networks (WSNs). This confirms the central role that WSN testbed features play in the design and implementation of the CTF platform, but also reaffirms the importance of the Integrated Testbed, developed in Task 3.1, in filling a real gap in the testing capability of the partners.

## 2.4.2 Deployment Configuration

### Number of SUT Nodes

Please enter the total number of nodes in the SUT that are supported by the testbed...

#### Responses



Testbed	Nodes
DSN	52
SOCRADES	13
Piaggio	40
Powerbench	28
NUIG-DERI	8
TWIST	204
UBonnTestbed	30
IoT@SAP	10
SICS-TB	17
ART-WiSe	30
VSTL	20


#### Comment

With a median size of 28 nodes, the responses show that the majority of CONET testbeds support relatively small SUT configurations. This reflects the significant financial and operative difficulties associated with building and maintaining testbed resources for large-scale experimentation. TWIST, one of the largest academic testbeds in Europe, stands out with its 204 node capability. Despite this large discrepancy in size, the scalability of the developed CTF solutions in terms of the number of SUT nodes, remains a primary design requirement.

**Deployment Scenario**

*Please specify the deployment scenario for the SUT supported by the testbed...*

**Responses**



Option	Number	Percentage
Indoor partitioned space	9	82 %
Indoor continuous space	2	18 %


**Comment**

All existing testbeds in CONET are oriented towards indoor SUT deployments, due to the challenging management and operational conditions for permanent outdoor testing installations. The design of the CTF platform, however, will also take into consideration some of the requirements of urban sensor testbeds like CitySense (Section 2.3).

**Deployment Pattern**

*Please specify the nature of the node deployment pattern...*

**Responses**



Option	Number	Percentage
Uniform 2D	5	45 %
Non-uniform 2D	3	27 %
Uniform 3D	1	9 %
Other	2	18 %


**Comment**

The responses show that the majority of testbeds support uniform 2D deployments, only TWIST has a uniform 3D deployment pattern. The regular configurations allow for rule-based addressing of the SUT nodes by the testbed infrastructure, and simplify the assessment of the node density impact on the observed SUT metrics.

## Mobility Support

*Please specify the supported mobility for the SUT nodes...*

### Responses



Option	Number	Percentage
Fixed nodes	11	100 %
Mobile nodes	3	27 %

### Comment


Since mobility is an important system factor in CO systems, the CTF platform will contain support for mobile nodes, although currently the majority of testbeds support only fixed SUT nodes. Several CONET testbeds are currently evaluating the possibility of extending the current installations with mobile units.

### 2.4.3 SUT Properties

#### Variability of Supported SUT Platforms

Please specify the number of SUT platforms that the testbed supports...

#### Responses



Option	Number	Percentage
Single - homogeneous SUT	4	36 %
Multiple - heterogeneous SUT	7	64 %


#### Comment

The heterogeneity of the CO systems is an important characteristic that has to be reflected in the capabilities of the testing infrastructure. To this end, it is important that the CTF platform offers generic services for uniform management of the different SUT platforms.

#### Primary SUT Platform

Please specify the name of the primary/largest supported SUT platform...

#### Responses



Option	Number	Percentage
Tmote Sky	6	55 %
MicaZ	1	9 %
TNOde	1	9 %
Other	3	27 %


#### Comment

In line with the wider research community, the primary SUT platform on the majority of testbeds uses MCUs from the TI MSP430 or Atmel ATmega128 MCU families, and radio transceivers with IEEE 802.15.4 PHY.

**Additional SUT Platform**

*Please specify the name of the second largest supported SUT platform (if applicable)...*

**Responses**

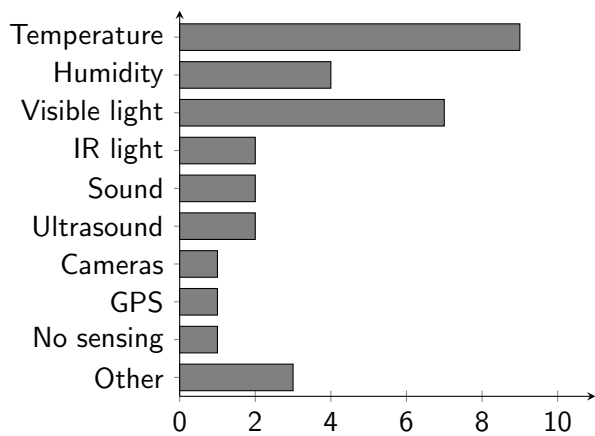


Option	Number	Percentage
Tmote Sky	1	14 %
MicaZ	1	14 %
BTnode	1	14 %
eyesIFX	1	14 %
TNOde	1	14 %
Other	2	29 %

**SUT Sensors**

*Please specify the types of sensor inputs available on the SUT...*

**Responses**



Option	Number	Percentage
Temperature	9	82 %
Humidity	4	36 %
Visible light	7	64 %
IR light	2	18 %
Sound	2	18 %
Ultrasound	2	18 %
Cameras	1	9 %
GPS	1	9 %
No sensing	1	9 %
Other	3	27 %


### Comment

The responses show that almost all testbeds have SUT platforms with basic sensing capabilities like temperature, IR/visible light, humidity, etc. Some testbeds also have more capable sensors like cameras and GPS units, enabling rich application level testing.

### SUT Actuators

*Please specify the types of actuation available on the SUT...*

### Responses



Option	Number	Percentage
LED	7	64 %
IR-LED	1	9 %
Ultrasound	2	18 %
Position	3	27 %
No actuation	1	9 %
Other	2	18 %

### Comment


In contrast to the sensing, the actuation capabilities of the current testbeds are more limited. Most of the SUT nodes have indication LEDs, and some also have sound outputs. The actuation of the

position in the deployment scenarios with mobility is important capability that should be considered in the design of the CTF experiment control abstractions.

### Power Supply of SUT Nodes

*Please specify the supported modes of supplying power to the SUT nodes...*

#### Responses



Option	Number	Percentage
Battery	6	55 %
Power over USB	8	73 %
Mains powered	4	36 %
Other	1	9 %

#### Comment


The control of the SUT power supply and the ability to switch between battery and externally powered setups is very useful capability for a CO testbed, enabling much richer experimental scenarios. The CTF experiment specification and control abstractions have to provide a unified way for describing and executing such actions.

### 2.4.4 Management Architecture

#### Control/Monitoring Plane Separation

*Please specify how the control/monitoring plane is implemented with respect to the SUT...*

##### Responses



Option	Number	Percentage
In-band	5	45 %
Out-of-band	6	55 %


##### Comment

Clean separation between the SUT and the testing infrastructure is important for minimizing the contamination of the experimental results by the testing activity. Unfortunately, many of the existing CONET testbeds don't have sufficient level of resource separation between the SUT and the testbed infrastructure. This makes the cross-validation of the experiments on the different testbeds of the future CTF even more valuable.

#### Control/Monitoring Plane Communication Technologies

*Please specify the communication technologies/standards used on the control/monitoring plane...*

##### Responses



Option	Number	Percentage
802.15.4	4	36 %
Bluetooth	2	18 %
WLAN	3	27 %
USB	6	55 %
Ethernet	5	45 %
Serial	2	18 %
Other	2	18 %


**Comment**

The majority of testbeds with separate SUT and testbed channels use USB or Ethernet as communication technology for the control/monitoring plane. Both of these standards enable combined data access and power supply control (via the USB 2.0 standard and PoE). This simplifies the hardware architecture of the testbed and reduces the costs.

**Control/Monitoring Plane Architecture**

*Please specify the type of backbone architecture...*

**Responses**




Option	Number	Percentage
Flat and centralized	4	36 %
Hierarchical and distributed	5	45 %
Other	2	18 %

**Comment**

Maintaining adequate performance in SUT deployments with several hundred nodes requires decentralization of the testing infrastructure. This, however, transforms the testbed in a large-scale distributed system with the associated challenges in concurrency management.

**Control/Monitoring Plane Gateway Platforms**

*Please specify the hardware platforms used on the control/monitoring plane in the case of a hierarchical architecture...*

**Responses**

Option	Number	Percentage
BTnode	2	22 %
Motiv Tmote Connect	1	11 %
Custom NSLU2	1	11 %
Laptop PC	3	33 %
Desktop PC	3	33 %
Other	3	33 %


**Comment**

Although they provide ample computational resources, the usage of PC-class devices as control/monitoring gateway platforms can limit the scalability of the configuration due to their relative high costs. In many cases the required functionality of this tier can be achieved with more constrained but energy-efficient embedded platforms like Linksys's NSLU2. This also opens the possibility of using some of them as part of the SUT since their capabilities are very similar to a "high-performance" node class in a typical WSN scenario.

**Job Execution Model**

*Please specify the supported types of test job execution...*

**Responses**



Option	Number	Percentage
Batch	3	27 %
Interactive	7	64 %
Other	3	27 %


**Comment**

The surveyed testbeds support a mixture of batch and interactive job execution models. Since both models have their own value, the CTF platform will provide support enabling the user to specify the wanted execution model for each experiment.

**Job Concurrency**

*Please specify the supported types of testbed resources sharing between multiple concurrent test jobs...*

**Responses**



Option	Number	Percentage
No concurrency	5	45 %
Per platform type	3	27 %
Per node ID	4	36 %
Per comm channel	1	9 %
Other	2	18 %


**Comment**

Due to the specifics of the wireless channel, concurrent access to the SUT resources based on node ID can be dangerous. In transparent testbeds, allowing full control over the SUT nodes and their radio interfaces, there is no reliable way to prevent intentional or unintentional interference between the concurrent jobs. This is why, in the first instance, the CTF platform will allow concurrent access only in the cases when the resources allocated to the separate experiments have disjoint spectral footprints.

**Job Registration and Scheduling Interface**

*Please specify the supported interfaces for test job registration and scheduling...*

**Responses**



Option	Number	Percentage
Web UI	2	18 %
GUI client	4	36 %
Programmable Client APIs	2	18 %
None	2	18 %
Other	3	27 %


**Comment**

Most of the testbeds interact with the users via some sort of GUI, either in the form of a custom desktop client or a web UI. Only a small number of testbeds offer client-side libraries for experiment specification and orchestration. The CTF Task will also concentrate first on the design and the implementation of server-side APIs, that will simplify the building of custom-side tools in later stages of the work.

**Authentication and Authorization Architecture**

*Please specify the supported forms of authentication and authorization...*

**Responses**



Option	Number	Percentage
Basic HTTP authorization	1	9 %
HTTPS with custom database	1	9 %
SSH public-key	1	9 %
None	6	55 %
Other	3	27 %

**Comment**

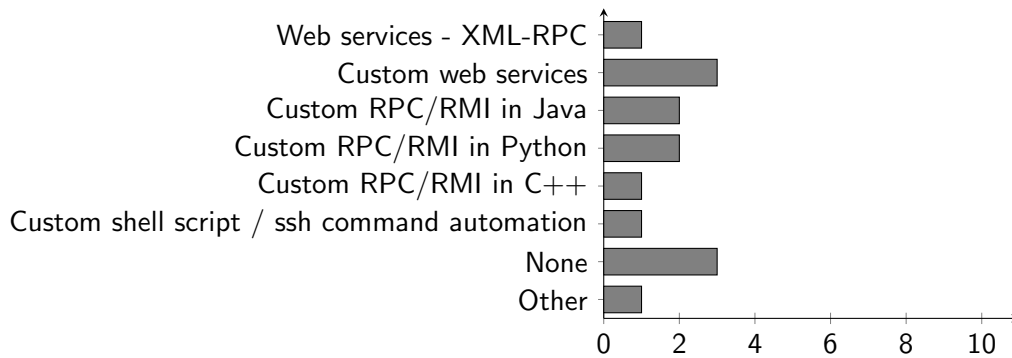
The fact that only small number of the existing CONET testbeds have mature authorization and authentication services, can be impediment for their integration in the future CTF. To address this capability gap, during the CTF implementation activities, we will evaluate the possibility of providing a reusable AAA subsystem that can be integrated with the existing management systems of the testbeds lacking AAA support.

### 2.4.5 Job Registration and Control

#### Job Control Interface

*Please specify the supported interfaces for interactive test job control...*

#### Responses



Option	Number	Percentage
Web services - XML-RPC	1	11 %
Custom web services	3	33 %
Custom RPC/RMI in Java	2	22 %
Custom RPC/RMI in Python	2	22 %
Custom RPC/RMI in C++	1	11 %
Custom shell script / ssh command automation	1	11 %
None	3	33 %
Other	1	11 %


#### Comment

The responses show large heterogeneity in the implementation technologies for the existing job control APIs. This strongly motivates a move towards programming language agnostic abstractions in the design of the CTF platform that will decouple the choice of the implementation language from the offered services.

#### SUT Node Control Services

*Please specify the basic SUT node control services provided by the testbed...*

**Responses**



Option	Number	Percentage
Power cycling	4	36 %
Resetting	6	55 %
Image programming	5	45 %
Image erasing	4	36 %
Data collection	7	64 %
Data injection	3	27 %
Power consumption monitoring	2	18 %
Other	7	64 %

**Comment**


The responses provide a nice summary of the essential SUT control services that the CTF platform should provide. They motivate a more generic approach for the specification and management of the control tasks, that can be easily extended with new services without effecting the existing experiment specifications.

**2.4.6 Management of Experimental Data**

**SUT Data Access**

*Please specify the type of access to data from the SUT supported by the testbed...*

**Responses**



Option	Number	Percentage
Direct raw access	5	63 %
Indirect access via client library APIs	4	50 %
Indirect access via "shell" abstraction	1	13 %
Indirect access via data persistence	3	38 %


**Comment**

The responses confirm the importance of having capability for both immediate, "raw" access to the experimental data, as well as for its long-term storage on the testbed servers.

**SUT Data Persistence**

*Please enter the supported modes of persisting the data generated during the job execution...*

**Responses**




Option	Number	Percentage
File-system storage in raw text format	3	33 %
Database storage in raw text format	3	33 %
None	2	22 %
Other	1	11 %

**Comment**

The majority of the testbeds persist the output from the experiments either in plain text format or in relational database tables. The CTF platform will initially concentrate on providing persistence in structured text formats, eliminating the need for user-provided dissectors for transforming SUT node output into database records.

**Time Stamping Services**

*Please specify the provided time stamping services...*

**Responses**

Option	Number	Percentage
Source-based time-stamping	1	11 %
Distributed time-stamping	3	33 %
Centralized time-stamping	1	11 %
None	3	33 %
Other	1	11 %

**Comment**


Time synchronization and time stamping are important services for the logging and tracing components. With sufficient fidelity they can enable causality analysis even in the case when the collection of the SUT output is performed on a distributed testbed infrastructure. They also form the basis for more advanced testbed services like snapshotting and recovering of SUT state.

### 2.4.7 Emulation and Simulation Support

#### Integration with Simulators and Emulators

*Please specify the types of integration with simulation/emulation frameworks that the testbed supports...*

#### Responses



Option	Number	Percentage
Integration of simulated nodes in experiments	3	30 %
Integration of testbed nodes in simulations	2	20 %
None	7	70 %
Other	2	20 %

#### Comment


The responses show that the level of direct integration between the existing CONET testbeds and simulation and emulation frameworks is currently limited, demonstrating the importance of the integration activities between Task 3.3 and Task 3.2 in the later phases of the project.

### 2.4.8 Usage Policy

#### Remote Use Policy

*Please specify the current remote use policy for the testbed...*

#### Responses



Option	Number	Percentage
Free access for project partners	2	20 %
Free access for academia	1	10 %
No remote access	8	80 %

#### Comment

The fact that the majority of CONET testbeds don't provide remote access to their resources is rather concerning from the point of view of the CTF implementation activities. The current situation can be traced back to two major causes: lack of technical capability and undefined remote usage policy. To mitigate the first problem, we are encouraging the partners to migrate to one of the open-source testbed management frameworks (like TWIST or MoteLab) that provide support for remote access. We are addressing the second problem by establishing a tentative unified CTF remote usage policy until we are in position to benefit from the legal infrastructure developed in the FIRE projects Panlab II and OneLab II.

## 2.5 CONET Testbed Federation Prototyping Testbeds

Based on the results from the CTF Survey, two mature CONET testbeds have been selected as primary testing ground for the prototype that will be implemented in the next phase of the CTF Task. In the following we provide a more detailed description of their current capabilities and architectural characteristics.

### 2.5.1 PowerBench

*PowerBench* [HHP<sup>+</sup>08] deployed at the *Delft University of Technology* is capable of parallel recording of the power consumption of all SUT nodes via low-cost custom interface boards with shunt resistors and an Analog-to-Digital Converter (ADC) that is sampled by a modified Linksys Network Attached Storage (NAS) device. The setup achieves sampling rates of about 5 kHz and a resolution of about 30  $\mu\text{A}$  (after calibration). The current deployment consists of 24 TNode nodes installed on the ceilings of 4 rooms and 8 Linksys NSLU2 devices. The software architecture includes tools for controlling the allocation of nodes to test runs, sampling and converting the power consumption output from the ADC as well as a TinyOS library for printf-like debugging over the serial port.

### 2.5.2 TWIST

The *TKN Wireless Indoor Sensor Network Testbed (TWIST)* [HKWW06] is a multi-platform, hierarchical testbed architecture developed at the *Technische Universität Berlin*. The TWIST instance at the TKN office building is one of the largest remotely accessible testbeds with 204 SUT sockets, populated with 102 eyesIFX and 102 Tmote Sky nodes. The SUT nodes are deployed in a 3D grid spanning 3 floors of the office building. The TWIST architecture introduces a layer of “super-nodes” between the SUT and the testbed server that play role in modeling hierarchical and hybrid SUT setups and in decentralizing the testbed management functions. TWIST relies on COTS hardware and fully leverages the features of the USB 2.0 standard. The SUT nodes are connected to the super-nodes via USB hubs which act as concentrators and provide a power supply management capability. This enables active SUT topology control and node fault injection modeling through selective powering on and off of SUT nodes. The software architecture is designed for easy remote access. The primary user interface is web-based and provides support for registration, configuration and monitoring of test jobs. The testbed resources can be shared among multiple users as long as each accesses different SUT platform. The web-interface is coupled to a background testbed server that uses a RPC architecture to distribute the testbed management tasks to the super-nodes, where they are executed in parallel on the attached SUT nodes. TWIST provides automatic trace collection and centralized time stamping service, as well as raw access to the serial I/O of the SUT nodes.

## 3 CONET Testbed Federation Platform

The integration of CO testbeds, in themselves large heterogeneous distributed systems, is an intrinsically complex task. The CTF platform presented in this chapter should be thus considered as a first step in a cyclic design refinement process that will continue for the remaining duration of the CTF Task. In particular, we expect significant additions and corrections to be included after the completion of the first prototyping cycle that will be documented in the upcoming Deliverable 3.6 (Testbed Federation Tool Suite).

The content of the chapter is divided in two main parts. In the first, we summarize the general objectives for the CTF platform, and we provide a description of the main architectural features by focusing on the relationship between the main components and the design principles behind the core abstractions.

In accordance with the selected architectural style, in the second part of the chapter, we specify the CTF platform through a detailed resource<sup>1</sup> model comprised of their representations and the possibilities for accessing and modifying the resource state using the HTTP method set.

### 3.1 Design Principles

The primary goal of the CTF is to enable convenient access to the resources of multiple CO testbeds, when organized in a loose federation of testbeds. The CTF platform provides an abstraction over the federation members and offers additional APIs that operate in the context of the federation aggregate in order to support the experimenters during the full experiment life-cycle.

In the design of the platform we have followed a set of guiding principles:

**Specialization** Although it shares some common characteristic with other testbed federation frameworks, the proposed CTF platform is carefully tuned to the specific needs of the CO testbeds. This specialization has enabled us to leave out some complex features that are not applicable or not important for the target domain, at the same time giving us opportunity to focus on the more important aspects like the impact of the wireless communication on the resource allocation problem, or the inclusion of tester controlled SUT mobility.

**User-centricity** In contrast to many other testbed federation frameworks that take an institution centric approach, our platform puts the individual user in the center of the design. It decouples service levels and policies from the question whether a particular user belongs to a federation member institution or not. All aspects of the service should be configurable and controllable at the granularity of a single user. This design principle, for example, has direct implications

---

<sup>1</sup>Following the established practice of using the word *resource* to refer both to the real physical assets of the testbeds, as well as to the main abstract concept in REST, we apply the word in both contexts in this chapter, clarifying its meaning only when it can not be readily deduced from the surrounding text.

on the architecture of the AAA abstractions and has significant impact on the openness of the platform.

**Scalability** A typical CO testbed is a large heterogeneous distributed system which makes the task of integrating several such systems behind a common federation platform particularly challenging. Addressing these challenges requires careful engineering that leverages the best-practices learned from other successful large distributed systems, like the Web. High scalability of the solution has to be maintained by promoting stateless interactions and using caching whenever possible.

**Extensibility** To be successful, the CTF platform has to be kept as simple as possible, but also modular and extensible, so new features and capabilities can be organically added when they are needed. In contrast to many other testbed federation frameworks, an explicit goal of the CTF platform is the openness towards external service providers. The same APIs that are used internally to build the higher-level federation services will be also made available to external entities to promote integration with their services.

**Coexistence** By focusing on a set of common, testbed-independent APIs, the CTF platform will necessarily lack some specific services of the individual testbeds that are valuable to the end-users. Thus, it is crucial to enable parallel use of the native interfaces of the member testbeds and the federation APIs. The CTF platform should not limit the autonomy of the members of the federation by imposing only a single access-path to the resources of the individual testbeds.

**Flexibility** The usage of the CTF platform should not impose unnecessary constraints on the development process on the user side. In particular, the platform should be accessible using various programming languages and the users should have freedom in selecting the level of abstraction overhead. This should be achieved by offering a basic, language-agnostic set of API, that enables building client-side solutions for raising the level of abstraction.

In the rest of the chapter we discuss how these high level principles have been converted into concrete requirements for the overall architecture of the CTF platform and the features of the core abstractions.

## 3.2 System Architecture

### 3.2.1 Overview

To illustrate the main architectural features of the CTF platform and to establish some common vocabulary, we start by revisiting our motivating use case of performing a cross-validation experimental study on a large number of CO testbeds, introduced in Chapter 1.

Figure 3.1 depicts our baseline scenario  $S_0$ , reflecting the existing service level. The user  $U$  wants to perform a study comprised of a set of experiments  $E_1, E_2, \dots, E_N$  that need to be executed on a set of CO testbeds:  $T_1, T_2, \dots, T_N$ . Each testbed provides a set of SUT resources  $R_1, R_2, \dots, R_N$  necessary for the corresponding experiment.

As it can be seen, without a federation substrate, the user can access the resources of the individual testbeds only over their native APIs:  $T_i$  API. This means that for each experiment, and

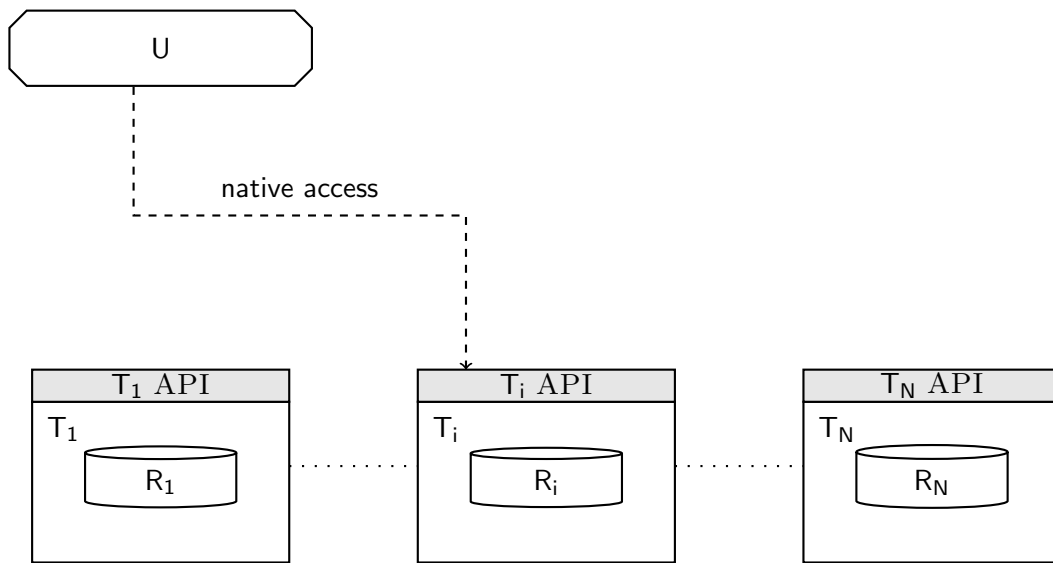


Figure 3.1: Baseline scenario  $S_0$ . The user can access the resources of the individual testbeds only over their native interfaces

after completing the native authenticating and authorization process, she needs to use a proprietary way to discover and reserve the required resources, to define and control the experiment, and finally, to collect the results.

The user needs to do this on each testbed separately. For  $N$  testbeds, she has to potentially use  $N$  different interfaces and processes. Without a federation substrate, there is no way to reuse the authentication and authorization credentials, no way to perform resource discovery across the multiple testbeds, no way to reuse the experiment specification, no way to reuse the client-side code for the on-line control of the experiment or for storing and post processing the results. In addition, there is no way to share this content with other users, so they can repeat the experiment and validate the results. There is also no common way to provide access to the results to external service providers that can provide storage or post processing (statistical analysis, plotting, etc.)

The CTF platform tries to address these limitations. Due to the large variability in the services offered by the individual testbeds this requires a two-step approach as described in the following.

As a first step, a common abstraction over the existing capabilities of the testbeds needs to be defined. This abstraction exports an interface, called *Testbed Adaptation API* (TA API), that can be used by the users to access the resources on the individual testbeds via a standardized interface. In this process of interface unification, some specific features of the individual testbeds will be undoubtedly remain unrepresented. Due to this, as well as to our coexistence and autonomy principles, we expect that this common interface will be used in parallel, and not instead of the native interfaces  $I_{T_i}$ .

The TA API offers a new service level to the users,  $S_1$ , leading to the scenario is depicted on Figure 3.2. Instead of using the various native testbed interfaces  $I_{T_i}$ , as in the baseline scenario, users can now access the resources over a standardized Application Programming Interface (API). When they need some testbed-specific capabilities they can always leverage the native interfaces.

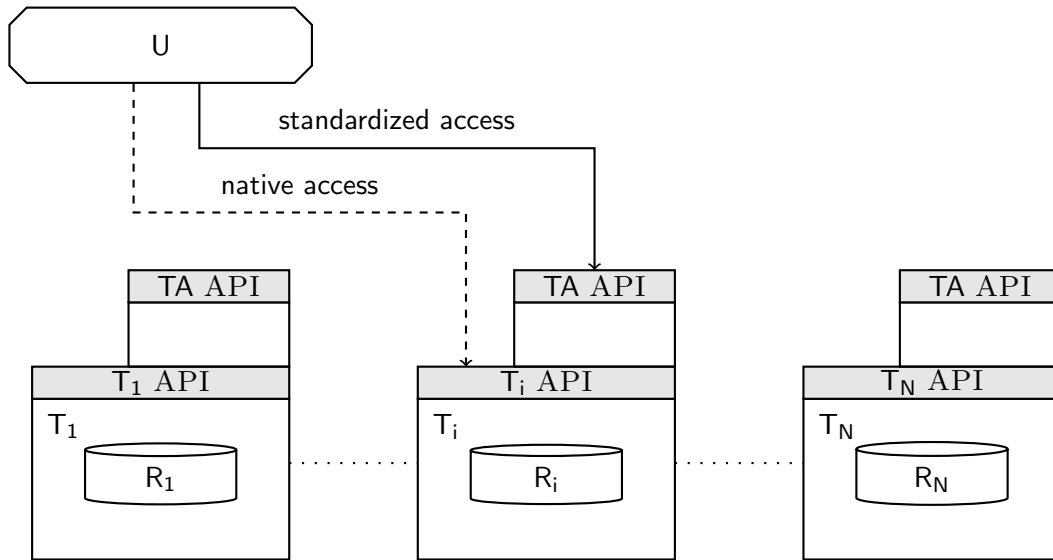


Figure 3.2: Scenario S<sub>1</sub>. The Testbed Adaptation API enables standardized access to the resources on the individual testbeds in addition to the native interfaces.

By incorporating adequate arbitration mechanisms, the implementation of the TA API will make sure that no conflicts in the access of the resources between the native and the federation users can occur.

Although this service level provides a significant convenience gain with respect to the S<sub>0</sub> scenario, by using only the Testbed Adaptation API (TA API) we are still short of providing some very useful services that abstract over the individual testbeds and operate in the context of the global federation aggregate. For example, we can provide a central place to store the definitions of the experiments, introduce a centralized discovery and reservation system, a central repository of traces and results, etc. For this we need an additional entity, a central *CONET Testbed Federation Server* (CTFS) that will offer a second interface class, the *Testbed Federation API* (TF API).

The CONET Testbed Federation Server (CTFS) leverages the services of the individual testbeds, exported through the standardized TA API, to build a higher level abstraction over the federation aggregate, thus offering new service level to the user, S<sub>2</sub>.

Figure 3.3 illustrates the relation between the individual testbed servers exporting the native and adaptation interfaces, the federation server and the users (both native and federation) plus external services in the S<sub>2</sub> scenario. The presented architecture has the flexibility to satisfy the major requirements outlined in our motivating use-case, and represents the basis for the CTF platform solution.

### 3.2.2 REST Architectural Style

The overall architecture of the CTF platform can be realized using many different architectural patterns that support rich interaction between a set of distributed components. Taking into consideration the main design principles presented in Section 3.1, we have decided to base the CTF

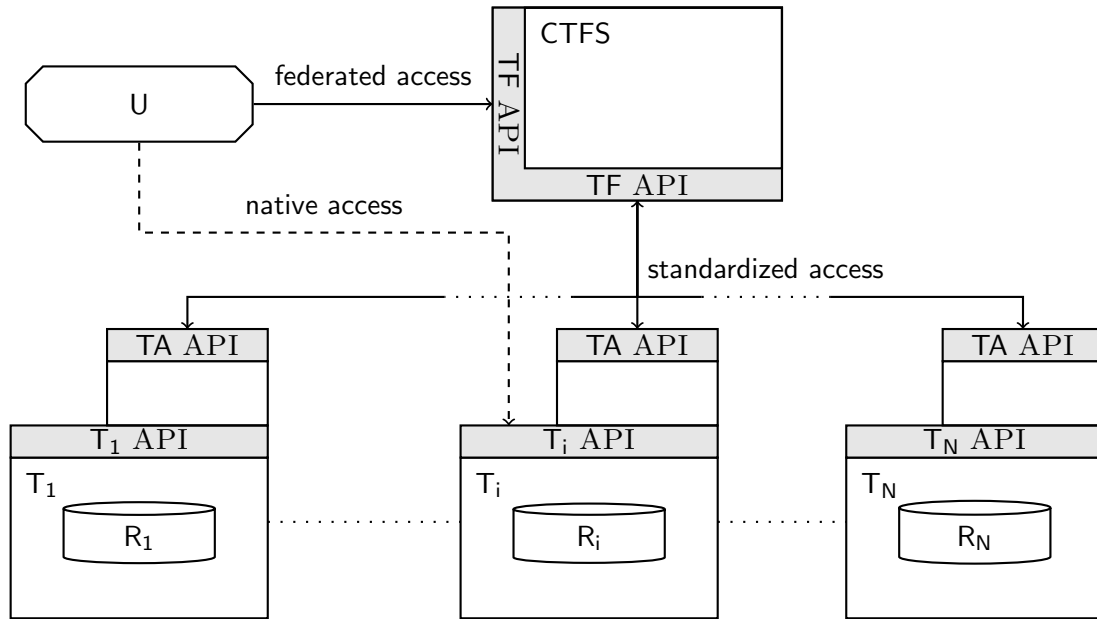


Figure 3.3: Scenario S<sub>2</sub>. The CONET Testbed Federation Server (CTFS) leverages the services of the individual testbeds, exported through the standardized Testbed Adaptation API, to build a higher level abstraction over the federation aggregate. The new services are made available to the user over the Testbed Federation API (TF API).

design on the *Representational State Transfer* (REST) architectural style [Fie00].

REST is a set of design constraints for developing rich resource-oriented systems that mirror the scalability and the flexibility of the Web. In the following we provide a short overview of the core properties of systems following the REST architectural style (so called RESTful systems) and their benefits in the context of the CTF platform.

### Resources

In contrast to Service Oriented Architecture (SOA) and other Remote Procedure Call (RPC)-based architectural styles where the data is kept private, encapsulated and hidden behind the processing components, in REST, the state and the nature of the data elements play a central role.

The *resource* is the main abstraction of information in a RESTful system. The *resource representation* captures the current or intended state of the resource. The components (clients, servers, etc.) act on the resources by transferring and modifying their representations. A *resource identifier* is used to uniquely identify the resource involved in the interaction.

### Resource Representations

Resources are abstract entities that can only be manipulated through their representations. A resource representation is a sequence of bytes, accompanied by representation metadata that describes those bytes. The components in a RESTful system use *media types* to differentiate between

the different possible representations of a resource. One scheme for specifying the media types in the system are Multipurpose Internet Mail Extensions (MIME) types [FN96].

The agreement on a set of common media types, together with the remaining architectural constraints, give messages in a RESTful system a “self-describing” property. Using only the media type as indication, the components in the system can safely perform many useful operations without having to look into the body of the message.

The resource model and the definition of the media type(s), used for representing the resources and driving resource and application state, forms the main cognitive effort in designing a RESTful system. Since all components in the system have to agree on the media types used for representing the resources, there is clear benefit in reusing a well defined media type from the MIME register whenever possible. Unfortunately, due to the specifics of the problem that the CTF platform is addressing, we were not able to follow this approach. Instead, we have decided to use a custom media types encoded with in a standard data interchange format.

In contrast to the existing testbed federation frameworks that use documents based on Extensible Markup Language (XML) (WISEBED uses WiseML, ProtoGENI uses Rspec, etc.), we have opted for JavaScript Object Notation (JSON) [Cro06] as the serialization method for our resources.

JSON is more light-weight and readable than the equivalent XML serialization and is especially suitable for exchanging general data structures. Code for parsing JSON-encoded data exists for large number of programming languages. The use of JSON is also very convenient when developing rich in-browser client side applications for interacting with the CTF platform. Because all JSON text is legal JavaScript code, it is very easy for a JavaScript program to convert the serialized data into an active object. Instead of using a heavy-weight parser, like in the case of XML, one can just use the built-in `eval()` function, passing the JSON encoded representation (after a security check) as parameter.

The default media type for JSON-encoded representations is `application/json`. Following the approach used in the *Sun Cloud API* [Ken], for our custom JSON-based representations we use the media type designation `application/ctf.{resource}+json`, where `{resource}` stands for the name of the particular resource that is being represented.

## Resource Identifiers

This ability to uniquely identify any resource in the system is a crucial characteristic of RESTful systems and represents the basis for their openness and composability. Although REST does not impose a specific naming scheme for the resource identifiers, in practice they typically take the form of Universal Resource Identifiers (URIs), as defined in the RFC 3986 [BLFM05].

There are two major subtypes of URIs: those that also include the “location” of the resource, so that the URI can be immediately dereferenced to get the representation of the resource (i.e. Universal Resource Locators (URLs)) and those that only provide a unique name (Universal Resource Names (URNs) or other Universally Unique Identifier (UUID)-based solutions) without the location part.

The URN identifiers are more stable (for example, a URL can be rendered invalid if the domain name of the server exporting the resources is changed). Many existing testbed federation frameworks like ProtoGENI [Pro] and WISEBED [Wis], use URNs as their default identification mechanism. The URNs stability, however, comes at a cost of increased administration overheads and loss of flexibility. The URN namespaces typically need to be registered with an external register [DVGIF02]

like IANA [IAN] and one loses the standardized dereferencing capability of URL.

We believe that the benefits that URLs carry as common addressing and naming scheme outweigh their shortcomings. The CTF platform uses URLs as resource identifiers exclusively. Any problems associated with non-persistent URLs on the individual testbeds will be handled using simple URL rewriting and redirecting rules in the federation components.

### Uniform interface

REST mandates a *uniform interface* between the components in the system. This is the most fundamental differentiation from the other networking styles. In contrast to the rich “verb” space in SOA and other RPC architectural styles, where each object can export different set of methods that operate on its state, in RESTful systems all resources are manipulated with exactly the same method set. In a similar way like the use of standard media types, the use of a generic interface constraints the design freedom but brings significant benefits in return.

Since all resources in the system can be manipulated with the same method set, the components don’t have to implement specialized code for accessing different resources in the system. Also, the semantics of each operation is well defined and uniform across all resources and components. This leads to an interface that is easy to understand and simplifies the interoperability between large number of uncoordinated actors. The uniform method set also opens the possibility of using a standardized set of return values to inform the caller about the success of the method invocation.

In RESTful systems implemented using Web technologies, the standard HTTP method set serves the role of a uniform interface. The HTTP 1.1 standard [FGM<sup>+</sup>99] defines eight methods: OPTIONS, GET, HEAD, POST, PUT, DELETE, TRACE and CONNECT, out of which the subset: GET, PUT, DELETE and POST are most commonly used in practice. Correspondingly, the HTTP status codes serve the role of return codes.

Table 3.1 provides a succinct overview of the semantics of the most common HTTP methods when applied in a RESTful context and Section 4.1.3 illustrates the use of HTTP status codes as return values.

HTTP Method	Properties	Semantics
GET	Safe, Idempotent	Retrieve a representation of the resource identified by the Request-URI without client-relevant side effects
HEAD	Safe, Idempotent	Same as GET, but only retrieve the header information without the message-body
PUT	Idempotent	Update the resource identified by the Request-URI with the representation contained in the message-body
DELETE	Idempotent	Delete the resource identified by the Request-URI
POST	Not safe, Not idempotent	Accept the representation in the message-body as a new subordinate of the resource identified by the Request-URI; Create a new resource without knowing the final URI; Append to the state of the resource identified by the Request-URI;

Table 3.1: Uniform interface formed by the HTTP method set

The “Property” column in the table illustrates how having a common set of methods with fixed semantics that is applicable to all resources in the system also leads to a more scalable and robust system.

Issuing a HTTP GET fetches the representation of the resource identified by the URI without any “strings attached”. Just like HEAD, this method is *safe*. It does not result in any change of client-relevant state at the server. Due to this properties, GET can support very efficient and sophisticated caching schemes. At the same time, the GET method is *idempotent*. The client is allowed to issue the same call one or ten times, if he wishes to do so, and the effects to the state of the identified resource at the server is guaranteed to be the same as if he made the call only once. The same idempotent property is shared by the PUT and the DELETE methods. If the client attempts to create or delete a resource and does not receive a positive status code, she knows that she can simply reissue the same request without relevant side-effects.

The HTTP POST method is neither safe nor idempotent. Making two POST requests to a collection resource will likely result in creation of two separate subordinate resources in the collection. The same applies when POST is used in a non-RESTful manner as a way to tunnel data to an arbitrary data-handling process. Because of this, the use of the POST requests has to be limited only to those scenarios where the safe or idempotent methods are not sufficient. In many cases a POST request can be avoided by reorganization of the resource model or by introducing a new resource representing the “result” of the indented activity of the original POST request.

The CTF platform uses the reduced HTTP method set described in Table 3.1 as its uniform interface. Apart from the scalability and robustness benefits explained previously, this also allows clients to cleanly separate the code that deals in generic way with the issuing of the request and the interpretation of the status codes with the service specific aspects that are part of the handling of the resource representations.

For example, fetching information about a particular node in a testbed would be implemented by a GET request on a node resource URI. The same code that handles the GET request can also be used for fetching information on a given platform, only the target will now be a platform resource URI. Similarly, the code that handles the PUT request for creating a new node resource at a given node URI can be shared with the code for creating a new binary image at a given image resource URI.

This should be compared with classical SOA and RPC styles where there are multiplicity of methods taking the role of the uniform method set. For example, the WISEBED API has methods like `getNodeList`, `getCapabilities`, `flashImagesToSensorNodes`, etc. Every service effectively speaks a different language and both the client and the intermediary components (caches, proxies, etc.) can not make any generic assumptions about their semantic properties.

### Statelessness

All REST interactions are *stateless*. In Web-based RESTful systems this means that each HTTP request should happen in full isolation and the client has to provide all of the information necessary for the server to understand the request, independent of any requests that may have preceded it.

This constraint, effectively moves the burden for maintaining the *application state* from the server to the client. The *resource state* is maintained on the server and it is the same for every client. If one client changes a particular resource state, this change is made visible to all other clients. It is the responsibility of the client to maintain any state that is specific only to that client.

The statelessness constraint frees the server from the need to retain application state between the individual requests, improving the scalability of the system. It enables parallel processing of the requests without further coordination apart from the resource state. It also allows intermediaries to

view each request in isolation. For example, a cache server can make a decision whether to cache or not a result from a request without fearing that state from some previous requests might affect its validity.

### Connectedness

Resource representations in RESTful systems contain links to other related resources allowing the client to *navigate* to them instead of using out-of-band information about the right URI where a given resource representation can be accessed.

A client should be able to effectively use a RESTful APIs without any prior knowledge beyond the initial URI and understanding of the standard media types appropriate for the specific application domain. From the initial URI, all application state transitions must be initiated by the client by selecting among a set of valid next states which are provided by the server as part of the representations of the manipulated resources.

Following this REST constraint, all CTF resources are interlinked and the servers guide the clients through the application state. The testbed servers remain in full control of their URI namespaces. This allows unconstrained evolution of the server-side services by limiting the effects that changes have on client-side code.

The connectedness principle is the main reason why the CTF platform is defined through its resource model and not through the URI space as usual in traditional non-RESTful Web services. As long as the client understands the media types used for representing the CTF resources, they can effectively access all services without prior knowledge of the URI hierarchy.

## 3.3 Core Abstractions

### 3.3.1 Authentication and Authorization

In order to securely identify the user, to grant the user access to all testbeds parts of the federation and to manage levels of access users should have to secured resources, the Federation of CO testbeds requires authentication and authorization facilities.

Single Sign-On (SSO) allows the user to authenticate on several servers by means of a single credential, so that the user doesn't need to use different passwords, often difficult on a large and complex network. On the other hand, we need a support for granting access to private resources according to customized policies and without exposing those resources to vulnerability.

Several protocols address the problem of authentication and authorization in computer networks. The Kerberos [Ker] protocol published by Massachusetts Institute of Technology (MIT), offers a solid and reliable support to authentication and authorization. Kerberos was released under open-source license and soon become a de facto standard and was adopted as a security protocol by the main operating systems.

At the time of writing, most of the applications run in the Internet and are based on the web which proved to be the answer to the increasing need of scalability. Currently, a large and growing number of web-oriented protocols address user identity trust and authentication, while much of the academic community seemed to have chosen Shibboleth [Shi].

Shibboleth is an institution-centric approach to identity management in which an organization is elected as a trusted verifier of the authentication data. When a user attempts to authenticate

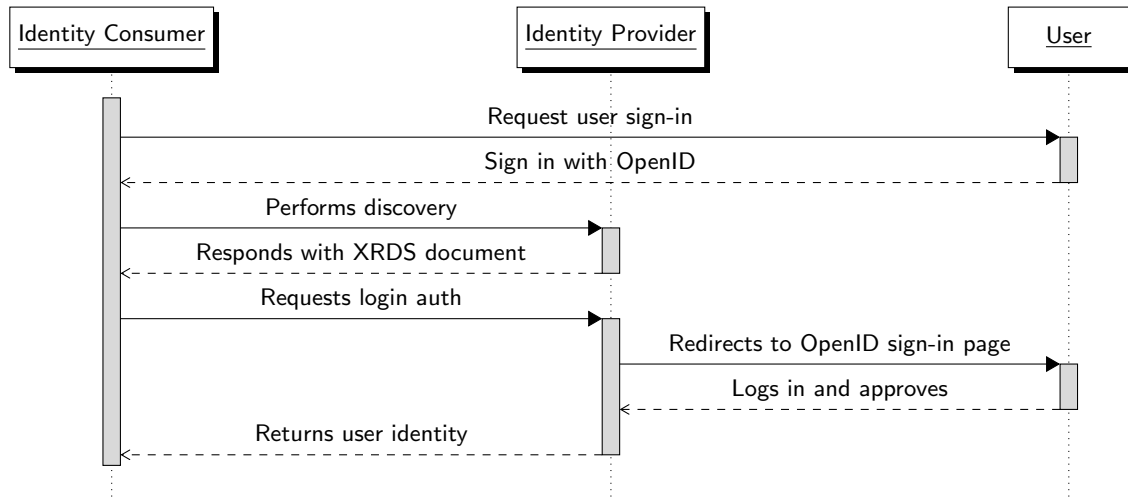


Figure 3.4: OpenID sequence diagram

against a server supporting Shibboleth, he basically asks his organization to certify his identity. This adheres to the paradigm of delegated authentication in which both the user and the web service tightly bind to the existence of a particular institution.

This is not acceptable in our case: users should be able to control their identity over the Internet and every testbed should be given autonomy and freedom to decide how to administer access policies. For this reason we need a user-centric identity solution.

OpenID [Ope] is arguably the leading protocol supporting a user-centric approach and a life-long and sustainable solution to identity management. There are currently over 120 million OpenID accounts since OpenID is gaining adoption by large organizations like Google, Yahoo!, Facebook, AOL, Symantec, VeriSign, Mozilla, and Novell.

The OpenID authentication process involves three different entities: the user, the Identity Provider (IP) and the Identity Consumer (IC). When a user attempts to gain access to an IC, he is redirected to his IP. OpenID is a URL, which means that the IC can easily use it to determine the location of the IP without recourse to some kind of directory service. The IP must authenticate the user credentials and redirects the user back to the IC to allow the user access.

There are several benefits deriving from adopting OpenID. From the user’s perspective, OpenID is a portable identity all over the federation and there is no need to open a new account on each of the testbeds of the federation. This also relieves each single testbed from the responsibility of storing and managing user identities while accelerating and simplifying the registration process.

While OpenID offers a light-weight HTTP-based solution to federated authentication, designed from the same creators of OpenID, OAuth [OAU] addresses the problem of allowing a user to grant a first server called Service Consumer (SC) access to a specific resource stored on a second server called Service Provider (SP), without sharing any credentials with the SC. To achieve this, OAuth introduces tokens in place of user credentials. While user credentials grant unlimited access to user’s private resources, tokens can restrict access to a limited set of resources or time.

The SC contacts the SP, asking for a request token. The SP verifies that the SC is registered and responds with an unauthorized request token. The SC directs the end user to an authorization page located on the SP, referencing the request token. On the authorization page, the user is prompted

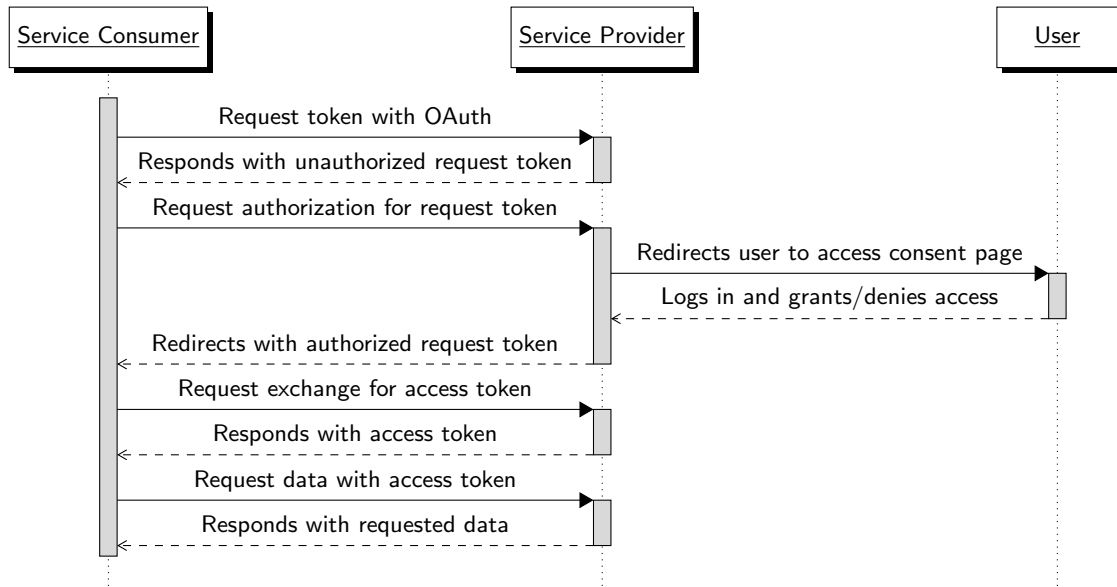


Figure 3.5: OAuth sequence diagram

to log into their account and then either grant or deny limited access to their data by the SC. If the user grants access, the SC sends a request to the SP to exchange the authorized request token for an access token. The SP verifies the request and returns a valid access token. The SC sends a request to the SP. The request is signed and includes the access token. If the SP recognizes the token, it supplies the requested data.

To better understand how OpenID and OAuth meet our needs enabling authentication and authorization in the CTF platform, we describe here two use cases.

In the first use case, involving OpenID, we can imagine to provide the Federation with an OpenID server acting as an IP, implementing the OpenID protocol for third-party service authentication. The Identity Provider may be local to the Federation Server or the Federation Server itself may delegate authentication to an external existing IP, like Google [Goo] or AOL for example, by means of the OpenID discovery capabilities. When the user needs to access one the Testbed Servers, he needs to authenticate against them, so that they can keep a local representation of every user of the federation. To obtain this, instead of forcing the user to open a new account on every Testbed Server, we enable them to be OpenID IC.

In the second case we can imagine a user launching a job involving several resources on a specific testbed. The job physically runs on the Testbed Server but the information about the experiment specification and the document containing the experiment results should be stored in the Federation Server or in any other external storage service. To achieve this, the user must be enabled to allow the Testbed Server to read and write his private information on his behalf. This could be obtained with an implementation of OAuth, i.e. enabling the Testbed Server to be a SC and the Federation Server to be a SP.

### 3.3.2 Resource Discovery, Reservation and Scheduling

As every system allowing authenticated users access to limited and distributed resources, the CTF requires discovery, reservation and scheduling capabilities, to provide the users with a view of the federation, to avoid conflicts, optimize usage and improve the users experience.

Discovery, reservation and scheduling must be *highly available* as if they becomes unavailable the entire access to resources is compromised.

*Scalability* is also a strong requirement: the federation will include several testbeds, the number of users and nodes may increase considerably to potentially hundreds users and thousands of nodes and a continuous monitoring of the target infrastructure may lead to a significant load.

As we show in the next sections, Resource Discovery, Reservation and Scheduling take and inherit full advantage of the benefits of a HTTP RESTful architecture, such as identifying resources with URIs and caching. In the last section we present two notable cases of related work, emphasizing the main differences with our approach.

#### Resource Discovery

In a federation of testbeds, where both physical and virtual resources are distributed across a worldwide heterogeneous network, an integrated mechanism for resource discovery<sup>2</sup> is necessary to provide users with an overview of the state of the underlying resources.

With the general term of *resources* we here indicate every single entity in the federation that can be addressed and accessed by applications. This does not include only nodes, which are the final objective of the whole CTF, but also any other entity that may take part to the experiment for any purpose, such as measuring, monitoring or simply forwarding information (i.e. channel scanners, spectrum analyzers).

While every experiment requires a specific subset of resources, we don't want to force the user to browse for them on every testbed. On the contrary we need to provide the user with an abstract aggregated view of the current state of the federation through a minimum number of effort and a priori assumptions.

Moreover, resource discovery must allow tracking of both relatively static and frequently changing resource characteristics and loosening and tightening of the granularity of the searching filter, according to the user's needs, by means of an expressive and adaptive query language.

Discovery happens at two different levels: *federation level* and *testbed level*. At *federation level* discovery is in charge to gather information coming from the different testbeds, while at *testbed level* the information is retrieved from the testbed resources. Such a hierarchical structure allows the discovery query processing to be distributed among all testbeds avoiding significant load at the federation level.

A discovery query issued to the CTFS may be split in *resource discovery* and *availability discovery*. The *resource discovery* answers to the question *What is present?* while *availability discovery* answers to *What is available?*. Since the result of *resource discovery* is likely slowly changing over time, it can be cached in the CTFS, behaving here like a shared cache (proxy). During *availability*

---

<sup>2</sup>With the term *discovery* in broad-sense we indicate here both *resource discovery* (the process of exploring/browsing for testbed resources) and the *availability discovery* (the process of retrieving information about resources *free* from other users reservations). Later in this section we will show how this decoupling can improve scalability.

*discovery* the system retrieves a collection of which resources are currently available in a specified time slot for a specified user. No detail about resources is returned in this phase, but only whether each resource is available for use or not. This is context-sensitive as it might be subject to access and priority policies. Decoupling information about resources and their availability allows discovery to inherit full advantage of REST, allowing optimization of caching.

This discovery request contains information about the searching parameters that could appear as HTTP GET query parameters or be uploaded (HTTP POST) as a document describing the experiment to be performed including the needed resources.

An example list of parameters follows here:

- *platform type*
- *radio technology*
- *amount of nodes per node type*
- *duration of experiment*
- *mobile node required*
- *sensors*
- *actuators*

### Resource Reservation

In order to avoid collision and to allow concurrent access at same time, the CTF includes a global Resource Reservation system.

With collision, we mean the conflict generated by concurrent access to the a shared portion of space, frequency and time. In our case, if we assume that the different testbeds are far apart enough to grant no electromagnetic interference during concurrent experiments, we can conclude that collision is avoided as long as experiments run on different time slots and on different portion of the spectrum.

A first simple approach would consist in reserving an entire testbed for a given slot of time. This would be enough to avoid any kind of conflicts during concurrent access to the federation by different users. Most of reservation systems for testbeds (and federation of testbeds) seems to have chosen this approach. In CONET we want to move further increasing the granularity of resource reservation allowing a bigger number of users to simultaneously access the CTF and optimizing the resource usage rate on a single testbed.

Accordingly with the federation philosophy, none of the federation features must prevent each testbed to reduce its autonomy and functionality. We assume that every testbed in the federation is already capable of managing reservations through the TA API. Where the native API does not allow concurrent access, we introduce a locking system at the *federation level* which temporarily isolates the testbed from the federation during access by native users.

We envision a web interface to let users reserve subsets of resources.

## Resource Scheduling

The CTFS introduces another layer of optimization which strongly impacts the user experience. A complete knowledge of the status of the underlying resources allows the Federation to support additional functions like solving dependent reservation requests featuring inter-testbed parameters and optimizing the usage of resources by providing the user with a series of alternatives (both in resources and in time) whenever the request for reservation could not succeed.

In the ProtoGENI implementation and deployment of Global Environment for Network Innovations (GENI), a virtual laboratory for at-scale networking experimentation, RSpec defines concepts like *advertisements*, *requests* and *tickets* which constitute an ad-hoc solution to the problem of discovery and reservation. While *Advertisements* describe resources provided by component managers, *requests* describe the resources that a client wants and *tickets* are resource reservations which are promised by a component manager to a client or broker. In our case *advertisements* are simply responses to aggregated requests about the state of the federation (mostly related with slowly varying information), *requests* are simply discovery queries and *tickets* are simply URLs pointing at the resource created out of a successful request for a new reservation.

### 3.3.3 Experiment Specification and Control

The main goals of the CTF is to facilitate running a given experiment on different testbeds. To do this conveniently the CTF must provide facilities for specifying the experiment setup in a testbed independent manner, as well as a facility for issuing the control commands necessary for preparing the experiments environment, executing the experiment and collecting results.

#### Experiment Specification

An experiment that is to be run on CTF is defined in a generic way so that it can be reused on all testbeds participating in the federation. It must define the types of hardware and software resources, as well as other properties of the infrastructure of the testbed in which the experiment should be run. The static types of hardware and software resources are described independently from the participating testbed resources in form of HTTP REST resources. Thus new types of device resources can easily be added to the federation by creating new resources without interfering with resources already in use. Each experiment contains one or more *property sets* that define what kind of devices are needed and how many of them. Each property set defines a unique combination of properties of a node type needed in the experiment. Nodes must have at least these properties but may offer other properties that are not defined in the property set. The node count in each property set defines how many of those nodes are needed, leading to disjoint property sets in the sense that when creating a property set with node count  $n$ , it is imperative to create exactly  $n$  virtual nodes that reference this property set. This part of the experiment specification is matched with the available resources from resource discovery.

The second part that needs to be specified are the tasks of the experiment, their order of execution and on which nodes they have to be executed. When executing an experiment there are two options:

In *batch mode* the user specifies all necessary tasks and setup for experiment execution in the experiment definition. In this case no further user interaction is necessary. The experiment is executed and the results are stored.

In *interactive mode* a subset of the experiment tasks are specified in the experiment definition, but the user may interact with experiment execution during runtime by adding further tasks on the fly. This allows the user to dynamically modify the course of the experiment.

This approach allows reusing experiments as templates. Thus by modifying part of an existing experiment (adding, removing, modifying or rescheduling tasks or virtual node groups) the user can easily create new experiments.

There are 2 kinds of task and node resources: Virtual tasks, virtual nodes (forming virtual node groups) and tasks and nodes (forming node groups). The former are needed to define the experiment independently of the executing testbed.

*Virtual nodes* are needed for naming each node participating in the experiment. Each virtual node must correspond to exactly one real node found by resource discovery. So when a task is defined a reference on which nodes it needs to be done can be added. Some tasks need to be done at the same time on several different virtual nodes, so these virtual nodes are grouped forming *virtual node groups*. A virtual node holds a name and is part of a property set. A virtual node group is formed for grouping all nodes that need to be addressed for executing one virtual task at a certain point in time. That is why a virtual node group may contain virtual nodes of different property sets. The characteristic grouping the nodes is the task and not the properties of the nodes. A consequence of this is that virtual nodes in a virtual node group are neither homogeneous nor exclusively found in one group. Consequently virtual node groups are not necessarily disjoint.

*Virtual tasks* describe which steps exactly have to be executed for the experiment. Each experiment step must define time of execution (relative to start of the experiment), which operation has to be executed, which nodes are involved and which input is needed and where it is located. As we use the REST architectural style, all information needed is stored within the federation in form of referenced resources or directly within the task resource. Besides some administration fields the virtual tasks resource contains the following fields:

- *eta* - time of execution of task, relative to start time
- *method* - Verb defining type of HTTP request
- *headers* - Headers defining media type/representation of reply and request, cache options, authorization options, definition of host
- *payload* - JSON representation of the needed input information for this operation, may contain references to other resources of fixed name-value pairs involved in the operation
- *target* - The reference to the virtual node group resource that references all virtual nodes that participate in this task.

The actual task is then constructed by assembling those fields into an HTTP request that is issued to the real node group. As the virtual task is constructed for the purpose of defining an experiment step independent from any testbed, the virtual task will not be translated into an HTTP request but is used as a generic variant describing what operation needs to be executed on what kind of nodes. This approach allows us to be independent of a fixed set of operations/experiment tasks or node types. Thus the experiment specification can be extended at any point without having to change existing federation resource types and can be used to create a testbed dependent experiment description by mapping the virtual node groups to testbed dependent node groups.

An example set of experiment tasks:

- *load image on node*
- *load image on node\_group*
- *reboot node*
- *reboot node\_group*
- *delete image from node*
- *delete image from node\_group*
- *move a node*
- *move a node\_group*
- *inject message on node*
- *inject message on node\_group*
- *extract message on node*
- *extract message on node\_group*
- *start battery-power node*
- *stop battery-power node*
- *start usb-power node*
- *stop usb-power node*
- *start monitoring node*
- *stop monitoring node*
- *start tracing*
- *stop tracing*
- *pick trace*

## Experiment Control

After having specified the experiment setup, we need to define how the experiment specification can be translated into a testbed dependent experiment version. And what is needed on the testbed-side for being able to translate the experiment tasks coming from the CTF into native testbed tasks that are executed at the time defined in the experiment.

Virtual tasks cannot be executed. When a matching testbed is found and a job has been scheduled the user is sent references(URI) to the reserved nodes. Those nodes must have a one to one mapping to the virtual nodes. Node groups corresponding to the virtual node groups must be created as well. The testbed-dependent task is then created by getting the virtual task representation and replacing the virtual node group references with the real node groups references. Those testbed dependent tasks are then processed by the CTFS. The CTFS translates each of the tasks into HTTP requests that are sent to the testbed server. Each testbed server must implement the TA API. Thus the tasks coming from the CTFS are understood by the testbed and queued in a tasks execution engine. According to the scheduling of the tasks received by the CTFS will the task execution engine trigger executing the testbed native tasks. The initial request of the CTFS receives a response with HTTP status code 202(see Section 4.1.3) and a reference to a status resource. Completion of the request is then announced to the status resource that can be polled for further information.

### 3.3.4 Logging and Tracing

The CTF must provide a facility for logging and storing information about events that occur while running an experiment. There are two kinds of events. Events that reflect changes in the testbed setup during an experiment and events that are directly triggered by the SUT. The former are referred to as logs, the latter as traces.

Experiment control on the CTFS is done by issuing HTTP requests to dedicated target resources, e.g. node groups. The change in the state of a resource of the CTF reflects changes in the testbed setup. Those changes need not only be reported to the tasks resource that caused them, but are also logged by creating a representation of a log resource. Similar to the concept of the OASIS Base event, where a predefined state change of a source component is reported by another component the CTF introduces the Log resource. The Log resource binds the finishing of a task on a dedicated resource to its status resource and adds a timestamp.

The Trace resource is triggered by the SUT, by the image that is burned on the nodes under test. As the CTF federates testbeds of different infrastructure it must be able to deal with nodes providing their traces via heterogeneous conduits and protocols. In order to be able to create a homogeneous trace files the nodes of each testbed provide information via what type of interface the output is provided. Besides the trace content the trace resource contains reference to the job, testbed, project and user.

## 4 CONET Testbed Federation Platform APIs

This section specifies the RESTful API of the CTF platform. The structure of the specification is inspired by the *Sun Cloud API* [Ken], a public Application Programming Interface (API) for Cloud services that most closely follows the spirit of the REST architectural style.

### 4.1 Common Features

This section specifies constraints that apply to all requests and responses in the CTF API. It includes the common headers, with special focus on the fields enabling efficient caching, and the shared status codes according to the HTTP /1.1 specification (RFC 2616) [FGM<sup>+</sup>99].

#### 4.1.1 HTTP Request Headers

Requests issued to the CTF APIs should include the following HTTP headers.

Header	Description
Accept	Specifies which media type(s) are acceptable for the response.
Authorization	HTTP Basic Authentication credentials, in case of access to protected resources by the user or by any other authorized entity.
Cache-Control	Specifies the caching directive(s) that must be followed in the request-response chain.
Content-Length	Indicates the size of the request body.
Content-Type	Indicated the media type of the request body.
Host	Specifies the Internet host and port number of the requested resource.
If-None-Match	Specifies a condition on the Etag for the request to be processed: the request should be forwarded to the server only if the cached response has been different ETag (GET, HEAD).
If-Match	Specifies a condition on the Etag for the request to be processed: the request should be forwarded to the server only if the cached response has not been modified (PUT, DELETE).
If-Modified-Since	Specifies a condition on the time for caching: the request should be forwarded to the server only if the cached response has been modified (GET, HEAD).
If-Unmodified-Since	Specifies a condition on the time for caching: the request should be forwarded to the server only if the cached response has not been modified (PUT, DELETE).

---

### 4.1.2 HTTP Response Headers

Responses returned by the CTF APIs should include the following HTTP headers.

---

Header	Description
Allow	Lists the set of methods supported by the resource identified by the Request-URI.
Cache-Control	Specifies the caching directive(s) that must be followed in the request-response chain.
Content-Length	Indicates the size of the response body.
Content-Type	Indicated the media type of the response body.
ETag	Provides the current value of the entity tag for the requested variant.
Expires	Gives the date/time after which the response is considered stale unless it is first validated with the origin server.
Last-Modified	Indicates the date and time at which the origin server believes the variant was last modified.
Location	This field is used to redirect the recipient to a location other than the Request-URI for completion of the request or identification of a new resource.
WWW-Authenticate	This field must be included in 401 (Unauthorized) response messages to invite the client to submit credentials for accessing a protected resource.

---

### 4.1.3 HTTP Status Codes

The CTF APIs will return HTTP status codes as described in the following table.

HTTP Status	Description
200 OK	The request was successfully completed.
201 Created	A request that created a new resource was completed, and no response body containing a representation of the new resource is being returned. A Location header containing the canonical URI for the newly created resource should also be returned.
202 Accepted	The request has been accepted for processing, but the processing has not been completed. The response will include an indication of the request's current status, and either a pointer to a status monitor or some estimate of when the user can expect the request to be fulfilled.
204 No Content	The server fulfilled the request, but does not need to return a response message body.
304 Not Modified	If the client has performed a conditional GET request and access is allowed, but the document has not been modified, the server will respond with this status code.
400 Bad Request	The request could not be processed because it contains missing or invalid information (such as validation error on an input field, a missing required value, and so on).
401 Unauthorized	The authentication credentials included with this request are missing or invalid.
403 Forbidden	The server recognized your credentials, but you do not possess authorization to perform this request.
404 Not Found	The request specified a URI of a resource that does not exist.
405 Method Not Allowed	The HTTP verb specified in the request (DELETE, GET, HEAD, POST, PUT) is not supported for this request URI.
406 Not Acceptable	The server is refusing to service the request because the Request-URI is longer than the server is willing to interpret.
409 Conflict	A creation or update request could not be completed, because it would cause a conflict in the current state of the resources supported by the server (for example, an attempt to create a new resource with a unique identifier already assigned to some existing resource).
410 Gone	The requested resource is no longer available at the server and no forwarding address is known. This condition is expected to be considered permanent. Clients with link editing capabilities SHOULD delete references to the Request-URI after user approval.
412 Precondition Failed	The precondition given in one or more of the request-header fields evaluated to false when it was tested on the server.
415 Unsupported Media Type	The resource identified by this request is not capable of generating a representation corresponding to one of the media types in the Accept header of the request.
500 Internal Server Error	The server encountered an unexpected condition which prevented it from fulfilling the request.
501 Not Implemented	The server does not (currently) support the functionality required to fulfill the request.
503 Service Unavailable	The server is currently unable to handle the request due to temporary overloading or maintenance of the server.

#### 4.1.4 Common Resources

##### Status [application/ctf.Status+json]

A Status represents a report on the progress of an asynchronous request for a change in the state of a specific resource. Such requests receive an HTTP status code of 202 Accepted, with the response body being a Status representation.

Status resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name describing the request whose status this represents.
media_type	String	1	application/ctf.Status+json
num_done	Integer	1	Represents the progress towards completion of the request. A value equal to <i>num_tot</i> indicates that the request has completed.
num_tot	Integer	1	Represents the value that <i>num_done</i> reaches when the request had completed.
target	URI	1	Represents the resource upon which the request is acting.
message	String	0..1	Brief message describing the completed operation (if successful) or an error message (if not successful).

##### Index [application/ctf.Index+json]

An Index represents a collection of resources. It allows browsing and creation of new resource by submitting an HTTP POST request to this resource.

The Index resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Index+json
items	Resource[]	1	An array of URIs to resources belonging to this collection.

## 4.2 Resource Model

According to the REST style, this section specifies the resources comprising the CTF platform. The following table provides an overview of the resources and their association with the TF API and the TA API:

Part of TF API	Part of TA API and TF API
Federation	Testbed
Project	User
Experiment	Job
PropertySet	NodeGroup
VirtualNodeGroup	Image
VirtualNode	Task
VirtualTask	Trace
	Log
	Socket
	Node
	Platform
	Interface
	Radio
	Sensor
	Actuator
	Mobility
	ImageFormat

### 4.2.1 Testbed Federation API

#### Federation [application/ctf.Federation+json]

A Federation represents a user's view of all accessible Testbeds, Users and Projects in the Federation of Testbeds. It is also the starting point for discovery of all other resources within the Testbed Federation API. Furthermore it includes common resource types (Platforms, Sensors, Actuators, Interfaces) that should be universally used when describing the properties of the testbed resources.

The Federation resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Federation+json
users	User[]	0..1	The list of Users currently registered to the CONET Testbed Federation.
projects	Project[]	0..1	The list of Projects currently stored in the CONET Testbed Federation.
testbeds	Testbed[]	0..1	The list of Testbeds currently registered to the CONET Testbed Federation.
platforms	Platform[]	0..1	The list of platforms adopted for nodes within the CONET Testbed Federation (testbeds supporting platforms must refer to this resource).
radio_technologies	RadioTechnology[]	0..1	The list of radio technologies adopted for nodes within the CONET Testbed Federation (nodes with radio devices must refer to this resource).
sensors	Sensor[]	0..1	The list of sensors adopted for nodes within the CONET Testbed Federation (nodes with sensors must refer to this resource).
actuators	Actuator[]	0..1	The list of actuators adopted for nodes within the CONET Testbed Federation (nodes with actuators must refer to this resource).
interfaces	Interfaces[]	0..1	The list of interfaces defined for communication with nodes within the CONET Testbed Federation (nodes implementing interfaces must refer to this resource).
image_formats	ImageFormat[]	0..1	The list of software image file formats within the CONET Testbed Federation (nodes allowing burning image files must refer to this resource).

#### Project [application/ctf.Project+json]

A Project represents a grouping of Experiments used in its context. It also holds references to Users participating in the project by running, designing and evaluating experiments.

Project resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Project+json
description	String	1	A short description of the Project.
users	Users[]	1	A list of Users registered to this Project.
testbeds	Testbed[]	0..1	A list of Testbeds the Project acts on.
experiments	Experiment[]	0..1	A list of Experiments used.
jobs	Job[]	0..1	A list of Jobs that implement project-related experiments on the federated testbeds.

### Experiment [application/ctf.Experiment+json]

The Experiment resource represents an experiment in terms of owner (User), documentation (Title, Description), the list of Jobs which implement the Experiment on different Testbeds, and finally the output data collected from all related Jobs. Experiments have three different sharing modes: *public*, *protected* and *private*. *Public* experiments are accessible by all users in the Federation, *protected* experiments are accessible by all users within the same Project and *private* experiments are only accessible to the Experiment's owner (the User). All experiment resources contain a reference to the project they are allocated to.

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Experiment+json
description	String	1	A short description of the Experiment.
owner	User	1	User owning the experiment.
project	Project	1	Project the experiment is allocated to.
testbeds	Testbed[]	0..1	A list of testbeds the experiment has been implemented for.
image_files	Image[]	0..1	A list of image files needed for running the experiment.
property_sets	PropertySet[]	0..1	Array of all property sets that define the kind of devices needed for experiment execution.
virtual_tasks	VirtualTasks[]	0..1	Array of all virtual tasks, that are needed for experiment execution.
virtual_node_groups	VirtualNodeGroup[]	0..1	Array of all virtual node groups needed for executing virtual tasks.
virtual_nodes	VirtualNode[]	0..1	Array of all virtual nodes needed for executing virtual tasks.
jobs	Job[]	0..1	A list of job that have been scheduled for the experiment.
traces	Trace[]	0..1	A list of Traces created when running the experiment-related jobs.
sharing	String	1	String indicating the sharing mode of the experiment. Either one of: <i>public</i> , <i>protected</i> and <i>private</i> .

### PropertySet [application/ctf.PropertySet+json]

The PropertySet resource represents a grouping of all properties required by a node for an experiment. It contains information about the needed platform type, interface types, sensor and actuator

types, radio technology and the number of nodes needed with these properties. This information is used during the resource discovery process to locate testbeds capable of executing the experiment.

The PropertySet resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	ctf.PropertySet+json
description	String	1	A short description for this resource.
owner	User	1	The user owning this property set.
experiment	Experiment	1	Experiment that uses this property set.
platform	Platform	0..1	The platform type needed.
radios	Radio[]	0..1	The type of radio technology required by the experiment.
interfaces	Interface[]	0..1	The Interface types that are required by the experiment.
sensors	Sensor[]	0..1	Sensors that are required by the experiment.
actuators	Actuator[]	0..1	Actuators that are required by the experiment.
mobility	Mobility	0..1	The supported type of mobility
node_count	Integer	1	The number of nodes needed of this type.

### VirtualNodeGroup [application/ctf.VirtualNodeGroup+json]

The VirtualNodeSet resource represents a frozen set of nodes that can be created for different purposes. It is a subset of the nodes participating in an experiment. A virtual node set is composed of all virtual nodes that take part in one task, e.g. switch off power supply to all nodes in this virtual node set. Virtual node groups are not necessarily disjoint.

The VirtualNodeSet resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.VirtualNodeGroup+json
virtual_nodes	VirtualNode[]	1	A list of virtual nodes.

### VirtualNode [application/ctf.VirtualNode+json]

The VirtualNode resource allows naming of each node participating in the experiment. Naming is necessary for defining which virtual tasks have to act on which dedicated virtual nodes. Furthermore the virtual node references the property set that defines its features.

The VirtualNode resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.VirtualNode+json
property_set	PropertySet	1	The property set defining the properties of this virtual node.

**VirtualTask [application/ctf.VirtualTask+json]**

A virtual task describes one step of the experiment execution. A Task can be loading an image to a single node or a whole node group, erasing an image from a single node or a whole node group, starting or stopping tracing, powering a node or node group on or off. Information about which of the named generic tasks should be executed is defined by the triple method, target\_uri and payload. The virtual task differs from the task, as it only operates on virtual node sets.

Virtual Task resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	ctf.VirtualTask+json
description	String	1	Short description of the virtual task.
eta	Integer	1	Integer value representing time of execution of task, relative to start time of the job.
method	String	1	Verb defining type of HTTP request.
headers	String	1	Headers defining media type/representation of reply and request, cache options, authorization options, definition of host.
payload	String	1	JSON representation of the needed input information for this operation, may contain references to other resources of fixed name-value pairs involved in the operation.
target	String	1	The URI of the virtual node group resource on which this request should act on.

## 4.2.2 TA API and TF API

### Testbed [application/ctf.Testbed+json]

The Testbed resource is the representation of a real testbed member of the CONET Testbed Federation.

Testbed resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Testbed+json
organization	String	1	The name of the organization running this testbed.
sockets	Socket[]	0..1	The list of sockets that may host nodes in this testbed.
nodes	Node[]	0..1	The list of nodes connected to the testbed.
platforms	Platform[]	0..1	Platforms currently supported by this testbed.
sensors	Sensor[]	0..1	The list of Sensors supported for Nodes within the testbed.
actuators	Actuator[]	0..1	The list of Actuators supported for Nodes within the testbed.
interfaces	Interfaces[]	0..1	The list of Interfaces supported for communication with Nodes within the testbed.
jobs	Job[]	0..1	Jobs scheduled on the testbed.

### User [application/ctf.User+json]

A User represents the user registered to the CONET Testbed Federation. The User resource model references all projects and experiments owned by the user. Furthermore it contains references to all experiments run by the user in form of references to jobs and traces.

User resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.User+json
openid	String	1	OpenID URL belonging this User.
email	String	1	E-mail address for the User.
organization	String	1	Organization which the User is part of.
projects	Project[]	0..1	Projects the User is involved.
experiments	Experiment[]	0..1	The Experiments owned by the User.
jobs	Job[]	0..1	Jobs owned by the User.

### Job [application/ctf.Job+json]

In case of *batch mode*, the user must specify every single Task to be automatically performed, including boot, configuration, control, tracing and shut-down. In case of *interactive mode*, the tasks attribute includes only the common Tasks needed for the boot and the shut-down. All other Task are left to the user.

Job resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Job+json]
description	String	1	Short description for this resource.
owner	User	1	The user running this job.
project	Project	1	The Project this job is related to.
experiment	Experiment	1	The Experiment that was implemented by this job.
testbed	Testbed	1	The testbed on which this job is to be run.
platform	Platform	1	The platform which this job will use.
datetime_from	Timestamp	1	The time when the Job starts.
datetime_to	Timestamp	1	The time when the Job shuts down.
duration	Integer	1	The duration of this job expressed in minutes.
nodes	Node[]	1	The list of nodes involved in this job.
node_groups	NodeGroup[]	0..1	Frozen sets of nodes involved in this job.
tasks	Task[]	0..1	The list of operations that has to be performed during this job.
traces	Trace[]	0..1	The list of trace files created during the execution of the job.

### NodeGroup [application/ctf.NodeGroup+json]

A NodeGroup is a frozen set of Nodes that can be created for different purposes.

NodeGroup resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.NodeGroup+json
nodes	Node[]	1	A list of Nodes.

### Image [application/ctf.Image+json]

The Image resource provides an image ready to be burned on a node. Furthermore it defines its own format and symbols that allow generic variables in the image file to be replaced with values that are provided only at runtime.

Image resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	Image [application/ctf.Image+json]
description	String	1	A short description for this Image.
owner	User	1	The user that owns this Image.
image_format	ImageFormat	1	Image format as defined by the ImageFormat resource.
content	Binary	1	The raw content of the Image file.
mapping_keys	String[]	0..1	A list of symbols in the Image file allowing generic variables to be assigned at runtime.
mapping_values	String[]	0..1	A list of values to assign at runtime to respective symbols.

**Task [application/ctf.Task+json]**

A Task describes one step of the experiment/job execution. A Task can be loading an image to a single node or a whole node group, erasing an image from a single node or a whole node group, starting or stopping tracing, powering a node or node group on or off. Information about which of the named tasks should be executed is defined by the triple method, target\_uri and payload. The Field relative\_time defines the time to execute the task relative to the beginning of a Experiment/Job start time.

Task resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Task+json
description	String	1	Short description of the task.
eta	Integer	1	Integer value representing time of execution of task, relative to start time of the job.
method	String	1	Verb defining type of HTTP request.
headers	String	1	Headers defining media type/representation of reply and request, cache options, authorization options, definition of host.
payload	String	1	JSON representation of the needed input information for this operation, may contain references to other resources of fixed name-value pairs involved in the operation.
target	String	1	The URI of the node group resource on which this request should act.
status	Status	0..1	The URI of the status resource containing information about the execution state of the task request.

**Trace [application/ctf.Trace+json]**

The Trace resource is a resource that contains the results of running an experiment on a dedicated testbed. The Trace resource holds references to the job that produced the result, the testbed where it was run as well as the experiment that was implemented and the project it is related to. The Content field contains the actual output that was collected when the Job was run.

Trace resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Trace+json
description	String	1	Short description of the trace.
user	User	1	User owning the trace.
project	Project	1	Project in which context the trace was created.
experiment	Experiment	1	The Experiment that was implemented, run an that produced the trace.
testbed	Testbed	1	Testbed on which the trace was generated.
job	Job	1	Job that created the trace file.
content	String[]	0..1	Each String is composed of the timestamp when the message was issued, the node it was issued by, the actual message payload and the number of bytes of the message payload.

**Log [application/ctf.Log+json]**

The Log resource links information about a change of the state of a component of the testbed, by referencing a task, its target and the status and the time of finished execution.

Log resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Log+json]
job	Job	1	The job that caused this Log.
task	Task	1	Task that was logged.
status	Status	1	Status resource that reported the finishing of the task.
target	URI	1	URI of the resource changed.
timestamp	Integer	1	Time when the event logged occurred.

**Socket [application/ctf.Socket+json]**

The Socket represents a testbed socket which serves as a container for information about a testbed receptacle that is independent from the node.

Socket resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Socket+json
x_coord	Number	1	Location for this node.
y_coord	Number	1	Location for this node.
z_coord	Number	1	Location for this node.
node	Node	1	A link to the node currently plugged to this socket.

**Node [application/ctf.Node+json]**

The Node describes properties (platform, interfaces, sensors, actuators, radio technology, mobility) and state (power, coordinates) of a testbed node.

Node resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Node+json
platform	Platform	1	The Platform represent the hardware characteristics of this node.
interfaces	Interface[]	1	Interfaces this node supports.
sensors	Sensor[]	0..1	A list of sensors featured by this Node.
actuators	Actuator[]	0..1	A list of actuators featured by this Node.
radio	Radio[]	0..1	The radio technology adopted by this node.
mobility	Mobility	0..1	If present indicates the kind of mobility featured by this socket.
power	Boolean	1	If <i>true</i> means power-on, if <i>false</i> power off.
image	Image	1	The Image file that is to be loaded in this node.
x_coord	Number	1	Location for this node (inherited by the socket this node is plugged to)
y_coord	Number	1	Location for this node (inherited by the socket this node is plugged to)
z_coord	Number	1	Location for this node (inherited by the socket this node is plugged to)

### Platform [application/ctf.Platform+json]

The Platform resource represents a set of hardware components of a Node.

Platform resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Platform+json
cpu	String	1	An identifier representing the model of CPU featured by this platform.
memory	Number	1	The size of the memory featured by this platform.
image_formats	ImageFormat[]	1	A list of supported image formats for this platform.
bandwidth_lower_bound	Integer	1	The lower bound in MHz for this platform (used to calculate the collision domain).
bandwidth_upper_bound	Integer	1	The upper bound in MHz for this platform (used to calculate the collision domain).

### Interface [application/ctf.Interface+json]

The Interface defines a communication interface supported by nodes for communication between the testbed infrastructure and the SUT. Since different kinds of interface may be supported by nodes it is important to have this resource universally defined within the CTF so that every member testbed may refer to this resource in the description of a given node.

Interface resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Interface+json

**Radio [application/ctf.Radio+json]**

The Radio resource standardizes radio technologies adopted in the CTF. Since different kinds of radio technologies may be supported by nodes it is important to have this resource universally defined within the CTF so that every member testbed may refer to this resource in the description of a given node.

Radio resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Radio+json

**Sensor [application/ctf.Sensor+json]**

The Sensor resource represents a device dedicated to measuring a particular physical parameter like temperature, pressure, humidity, light or sound. Since different kinds of sensor may be supported by nodes it is important to have this resource universally defined within the CTF so that every member testbed may refer to this resource in the description of a given node.

Sensor resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Sensor+json]
description	String	1	A human-readable description of this Sensor.
values	Number[]	1	The measured values of this Sensor.
units	String[]	1	The units of measurement by which values are expressed.

**Actuator [application/ctf.Actuator+json]**

The Actuator resource represents a device which can operate a particular action in the physical world, i.e. sound speakers, LED or robots. Since different kinds of actuator may be supported by nodes it is important to have this resource universally defined within the CTF so that every member testbed may refer to this resource in the description of a given node.

Actuator resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Actuator+json
description	String	1	A human-readable description of this Actuator.
values	Number[]	1	The values this Actuator should assume.
units	String[]	1	The units of measurement by which values are expressed.

**Mobility [application/ctf.Mobility+json]**

The Mobility resource defines a mobility feature for a node. Since different kinds of mobility may be supported by nodes it is important to have this resource universally defined within the CTF so that every member testbed may refer to this resource in the description of a given node. For example nodes may move either in a 2-dimensional or 3-dimensional space.

Mobility resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.Mobility+json
description	String	1	A description for this resource.

**ImageFormat [application/ctf.ImageFormat+json]**

The ImageFormat resource defines a format for the Image file that can be burned on a given node. Since different platforms may support different image formats it is important to have this resource universally defined within the CTF.

ImageFormat resource model contains the following fields:

Field Name	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the representation of this resource.
name	String	1	A given name for this resource.
media_type	String	1	application/ctf.ImageFormat+json]

### 4.3 Requests and Responses

In the following tables we present a description for the application of the four HTTP methods to the resources introduced in the previous section. In the *Status Codes* column, the following codes 400, 401, 403, 404, 406 and 500 are intended to be always implicit since they may refer to any kind of HTTP Request with the same meaning, while we specify only those HTTP Status Codes that have a different meaning according to the different HTTP Methods.

#### 4.3.1 Testbed Federation API

##### Project [application/ctf.Project+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the Project resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the Project including Users, Experiments, Testbeds or Jobs.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

##### Experiment [application/ctf.Experiment+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the Experiment resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the Experiment including PropertySets, VirtualNodes, VirtualTasks or Images.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

##### PropertySet [application/ctf.PropertySet+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the PropertySet resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the PropertySet including User, Experiment, Platforms, Radios, Interfaces, Sensors, Actuators, Mobility and node number.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

**VirtualNodeGroup [application/ctf.VirtualNodeGroup+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the VirtualNodeGroup resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the VirtualNodeGroup including references to VirtualNodes contained.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

**VirtualNode [application/ctf.VirtualNode+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the VirtualNode resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

**VirtualTask [application/ctf.VirtualTask+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the VirtualTask resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the VirtualTask including values of method, headers, payload or target fields.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

### 4.3.2 TA API and TF API

#### NodeGroup [application/ctf.NodeGroup+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the NodeGroup resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the NodeGroup including references to Nodes contained.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

#### Image [application/ctf.Image+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the Image resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the Image updating mapping_keys, mapping_values, binary or image format resource reference.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

#### Task [application/ctf.Task+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the Task resource.	200, 204, 304, 410
PUT	Modifying JSON representation for the Task, should only affect the Status resource reference.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

#### Trace [application/ctf.Trace+json]

Method	Description	Status Codes
GET	Returns the JSON representation of the Trace resource.	200, 204, 304, 410
POST	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

**Log [application/ctf.Log+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Log resource.	200, 204, 304, 410
POST	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

**Socket [application/ctf.Socket+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Socket resource.	200, 204, 304, 410
PUT	Upload a new JSON representation for a given Socket.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Delete the resource specified by the URI	200, 409, 412

**Node [application/ctf.Node+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Node resource.	200, 204, 304, 410
PUT	Upload a new JSON representation for a given Node.	200, 202, 409, 412
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**Platform [application/ctf.Platform+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Platform resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**Interface [application/ctf.Interface+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Interface resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**Radio [application/ctf.Radio+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Radio resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**Sensor [application/ctf.Sensor+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Sensor resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**Actuator [application/ctf.Actuator+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Actuator resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**Mobility [application/ctf.Mobility+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the Mobility resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

**ImageFormat [application/ctf.ImageFormat+json]**

Method	Description	Status Codes
GET	Returns the JSON representation of the ImageFormat resource.	200, 204, 304, 410
PUT	Method not allowed.	405
POST	Method not allowed.	405
DELETE	Method not allowed.	405

## 4.4 Invocation Examples

### 4.4.1 User creates a new Project

```
POST /projects/
Host: federation.com
Authorization: Basic XXXXXXXXXXXX
Accept: application/ctf.Project+json
Content-Length: nnnn
Content-Type: application/ctf.Project+json

{
  "name" : "Test Project",
  "description" : "A simple test project",
  "users" :
  [
    {
      "uri" : "http://federation.com/users/123",
      "name" : "John Smith",
      "media_type" : "application/ctf.User+json"
    }
  ]
}

HTTP/1.1 201 Created
Location: http://federation.com/projects/456
```

### 4.4.2 User creates a new Experiment

```
POST /projects/456/experiments
Host: federation.com
Authorization: Basic XXXXXXXXXXXX
Accept: application/ctf.Experiment+json
Content-Length: nnnn
Content-Type: application/ctf.Experiment+json

{
  "name" : "Test Experiment",
  "description" : "A simple test experiment",
  "user" :
  {
    "uri" : "http://federation.com/users/123",
    "name" : "John Smith",
    "media_type" : "application/ctf.User+json"
  },
  "project" :
  {
    "uri" : "http://federation.com/projects/456",
    "name" : "Test Project",
    "media_type" : "application/ctf.Project+json"
  },
}
```

```
    "sharing" : "public"  
  }  
}
```

HTTP/1.1 201 Created

Location: <http://federation.com/projects/456/experiments/789>

#### 4.4.3 User creates two PropertySets

POST /projects/456/experiments/789/property\_sets

Host: federation.com

Authorization: Basic XXXXXXXXXX

Accept: application/ctf.PropertySet+json

Content-Length: nnnn

Content-Type: application/ctf.PropertySet+json

```
{  
  "name" : "Test Property Set 1",  
  "description" : "Data collection property set",  
  "node_count" : "10",  
  "platform_type" :  
  {  
    "uri" : "http://federation.com/platforms/165",  
    "name" : "TelosB",  
    "media_type" : "application/ctf.Platform+json"  
  },  
  "interfaces" :  
  [  
    {  
      "uri" : "http://federation.com/interfaces/769",  
      "name" : "USB",  
      "media_type" : "application/ctf.Interface+json"  
    }  
  ],  
  "sensors" :  
  [  
    {  
      "uri" : "http://federation.com/sensors/587",  
      "name" : "Temperature",  
      "media_type" : "application/ctf.Sensors+json"  
    },  
    {  
      "uri" : "http://federation.com/sensors/639",  
      "name" : "Light",  
      "media_type" : "application/ctf.Sensors+json"  
    }  
  ],  
  "experiment" :  
  {  
    "uri" : "http://federation.com/projects/456/experiments/789",  
    "name" : "Test Experiment",  
    "media_type" : "application/ctf.Experiment+json"  
  }  
}
```

```
}
}
```

HTTP/1.1 201 Created

Location: [http://federation.com/projects/456/experiments/789/property\\_sets/659](http://federation.com/projects/456/experiments/789/property_sets/659)

POST /projects/456/experiments/789/property\_sets

Host: federation.com

Authorization: Basic XXXXXXXXXX

Accept: application/ctf.PropertySet+json

Content-Length: nnnn

Content-Type: application/ctf.PropertySet+json

```
{
  "name" : "Test Property Set 2",
  "description" : "Data sink",
  "node_count" : "1",
  "platform_type" :
  {
    "uri" : "http://federation.com/platforms/165",
    "name" : "TelosB",
    "media_type" : "application/ctf.Platform+json"
  },
  "interfaces" :
  [
    {
      "uri" : "http://federation.com/interfaces/769",
      "name" : "USB",
      "media_type" : "application/ctf.Interface+json"
    }
  ],
  "mobility" :
  {
    "uri" : "http://federation.com/mobilities/295",
    "name" : "2D",
    "media_type" : "application/ctf.Mobility+json"
  }
  "experiment" :
  {
    "uri" : "http://federation.com/projects/456/experiments/789",
    "name" : "Test Experiment",
    "media_type" : "application/ctf.Experiment+json"
  }
}
```

HTTP/1.1 201 Created

Location: [http://federation.com/projects/456/experiments/789/property\\_sets/825](http://federation.com/projects/456/experiments/789/property_sets/825)

#### 4.4.4 User uploads two Image files

POST /projects/456/experiments/789/images

```
Host: federation.com
Authorization: Basic XXXXXXXXXXXX
Accept: application/ctf.Image+json
Content-Length: nnnn
Content-Type: application/ctf.Image+json
```

```
{
  "name" : "Test Image 1",
  "description" : "Image for data collection",
  "image_format" :
  {
    "uri" : "http://federation.com/image_formats/658",
    "name" : "IHEX",
    "media_type" : "application/ctf.ImageFormat+json"
  },
  "content" : "YYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYY",
  "mapping_keys" : ["NODE_ID"],
  "mapping_values" : ["http://federation.com/testbeds/{testbed}/nodes/{node}/id"]
}
```

HTTP/1.1 201 Created

Location: http://federation.com/projects/456/experiments/789/images/648

POST /projects/456/experiments/789/images

```
Host: federation.com
Authorization: Basic XXXXXXXXXXXX
Accept: application/ctf.Image+json
Content-Length: nnnn
Content-Type: application/ctf.Image+json
```

```
{
  "name" : "Test Image 2",
  "description" : "Image for data sink",
  "image_format" :
  {
    "uri" : "http://federation.com/image_formats/658",
    "name" : "IHEX",
    "media_type" : "application/ctf.ImageFormat+json"
  },
  "content" : "XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX",
  "mapping_keys" : ["NODE_ID"],
  "mapping_values" : ["http://federation.com/testbeds/{testbed}/nodes/{node}/id"]
}
```

HTTP/1.1 201 Created

Location: http://federation.com/projects/456/experiments/789/images/346

#### 4.4.5 User creates two Virtual Tasks

POST /projects/456/experiments/789/virtual\_tasks

Host: federation.com

```
Authorization: Basic XXXXXXXXXXXX
Accept: application/ctf.VirtualTask+json
Content-Length: nnnn
Content-Type: application/ctf.VirtualTask+json
```

```
{
  "name" : "Test Virtual Task",
  "description" : "Burning Image 1 on Virtual NodeGroup 1",
  "eta" : "0",
  "method" : "PUT",
  "headers" : "
    Host: federation.com
    Authorization: Basic XXXXXXXXXXXX
    Accept: application/ctf.VirtualNodeGroup+json
    Content-Length: nnnn
    Content-Type: application/ctf.VirtualNodeGroup+json
  "
  "payload" : "
    "image" :
    {
      "uri" : "http://federation.com/projects/456/experiments/789/images/648",
      "name" : "Test Image 1",
      "media_type" : "application/ctf.Image+json"
    },
    "eta" : "0"
  "
  "target" : "/projects/456/experiments/789/virtual_node_groups/659"
}
```

HTTP/1.1 201 Created

Location: [http://federation.com/projects/456/experiments/789/virtual\\_tasks/729](http://federation.com/projects/456/experiments/789/virtual_tasks/729)

POST /projects/456/experiments/789/virtual\_tasks

```
Host: federation.com
Authorization: Basic XXXXXXXXXXXX
Accept: application/ctf.VirtualTask+json
Content-Length: nnnn
Content-Type: application/ctf.VirtualTask+json
```

```
{
  "name" : "Test Virtual Task",
  "description" : "Burning Image 2 on Virtual NodeGroup 2",
  "eta" : "0",
  "method" : "PUT",
  "headers" : "
    Host: federation.com
    Authorization: Basic XXXXXXXXXXXX
    Accept: application/ctf.VirtualNodeGroup+json
    Content-Length: nnnn
    Content-Type: application/ctf.VirtualNodeGroup+json
  ",
```

```

    "payload" : "
      "image" :
        {
          "uri" : "http://federation.com/projects/456/experiments/789/images/346",
          "name" : "Test Image 2",
          "media_type" : "application/ctf.Image+json"
        }
      ,
      "target" : "/projects/456/experiments/789/virtual_node_groups/825"
    }

```

HTTP/1.1 201 Created

Location: [http://federation.com/projects/456/experiments/789/virtual\\_tasks/628](http://federation.com/projects/456/experiments/789/virtual_tasks/628)

#### 4.4.6 User explores the federation by submitting the experiment definition

GET /testbeds?experiment=<http://federation.com/projects/456/experiments/789>

Host: federation.com

Authorization: Basic XXXXXXXXXXXX

Accept: application/ctf.Testbed+json

HTTP/1.1 200 OK

Content-Length: nnnn

Content-Type: application/ctf.Testbed+json

```

[
  {
    "uri" : "http://federation.com/testbeds/234",
    "name" : "Testbed A",
    "media_type" : "application/ctf.Testbed+json"
  },
  {
    "uri" : "http://federation.com/testbeds/872",
    "name" : "Testbed B",
    "media_type" : "application/ctf.Testbed+json"
  },
  {
    "uri" : "http://federation.com/testbeds/926",
    "name" : "Testbed C",
    "media_type" : "application/ctf.Testbed+json"
  },
]

```

#### 4.4.7 User creates a new Job for a given Experiment

POST /projects/456/experiments/789/jobs

Host: federation.com

Authorization: Basic XXXXXXXXXXXX

Accept: application/ctf.Job+json

Content-Length: nnnn

Content-Type: application/ctf.Job+json

```
{
  "name" : "Test Job",
  "description" : "A simple test job",
  "owner" :
  {
    "uri" : "http://federation.com/users/235",
    "name" : "John Smith",
    "media_type" : "application/ctf.User+json"
  },
  "project" :
  {
    "uri" : "http://federation.com/projects/623",
    "name" : "My sample project",
    "media_type" : "application/ctf.Project+json"
  },
  "experiment" :
  {
    "uri" : "http://federation.com/projects/623/experiments/812",
    "name" : "My sample experiments",
    "media_type" : "application/ctf.Experiment+json"
  },
  "testbed" :
  {
    "uri" : "http://federation.com/testbeds/234",
    "name" : "Testbed A",
    "media_type" : "application/ctf.Testbed+json"
  },
  "platform" :
  {
    "uri" : "http://federation.com/platforms/653",
    "name" : "TelosB",
    "media_type" : "application/ctf.Platform+json"
  },
  "datetime_from" : "2009-11-18 14:00:00",
  "datetime_to" : "2009-11-16 18:00:00",
}
```

HTTP/1.1 201 Created

Location: <http://federation.com/projects/456/experiments/789/jobs/124>

## Bibliography

- [AER<sup>+</sup>06] A. Arora, E. Ertin, R. Ramnath, M. Nesterenko, and W. Leal. Kansei: a high-fidelity sensing testbed. *Internet Computing, IEEE*, 10(2):35–47, March–April 2006. doi:10.1109/MIC.2006.37.
- [ARE<sup>+</sup>05] A. Arora, R. Ramnath, E. Ertin, P. Sinha, S. Bapat, V. Naik, V. Kulathumani, Hongwei Zhang, Hui Cao, M. Sridharan, S. Kumar, N. Seddon, C. Anderson, T. Herman, N. Trivedi, M. Nesterenko, R. Shah, S. Kulkarni, M. Aramugam, Limin Wang, M. Gouda, Young ri Choi, D. Culler, P. Dutta, C. Sharp, G. Tolle, M. Grimmer, B. Ferreira, and K. Parker. Exscal: Elements of an extreme scale wireless sensor network. In *Proceedings of the 11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA '05)*, pages 102–108, Aug. 2005. doi:10.1109/RTCSA.2005.47.
- [BABM05] John Bicket, Daniel Aguayo, Sanjit Biswas, and Robert Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 31–42, New York, NY, USA, 2005. ACM. doi:http://doi.acm.org/10.1145/1080829.1080833.
- [BLFM05] T. Berners-Lee, R. Fielding, and L. Masinter. Uniform resource identifiers (uri): Generic syntax, 2005.
- [cas] Casino Lab, Colorado School of Mines. <http://toilers.mines.edu/Public/>. [Online; accessed: 14.05.2009].
- [CBA<sup>+</sup>05] B.N. Chun, P. Buonadonna, A. AuYoung, Chaki Ng, D.C. Parkes, J. Shneidman, A.C. Snoeren, and A. Vahdat. Mirage: A microeconomic resource allocation system for sensornet testbeds. *Embedded Networked Sensors, 2005. EmNetS-II. The Second IEEE Workshop on*, pages 19–28, May 2005.
- [Cro06] D. Crockford. The application/json media type for javascript object notation (json), 2006.
- [Dat] Datalogisk Institut (DIKU), University of Copenhagen. Re-Mote Testbed Framework. <http://code.google.com/p/remote-testbed/>. [Online; accessed: 14.05.2009].
- [DBK<sup>+</sup>07] Matthias Dyer, Jan Beutel, Thomas Kalt, Patrice Oehen, Lothar Thiele, Kevin Martin, and Philipp Blum. Deployment Support Network - A Toolkit for the Development of WSNs. In *Proceedings of the 4th European Conference on Wireless Sensor Networks (EWSN)*, pages 195–211, Delft, The Netherlands, 2007.

- [dL] École Polytechnique Fédérale de Lausanne. Wasal testbed. <http://wasal.epfl.ch/>. [Online; accessed: 14.05.2009].
- [DRSC05] P. De, A. Raniwala, S. Sharma, and T. Chiueh. Mint: a miniaturized network testbed for mobile wireless research. *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, 4:2731–2742 vol. 4, March 2005. doi:10.1109/INFCOM.2005.1498556.
- [DVGIF02] L. Daigle, D.W. Van Gulik, R. Iannella, and P. Faltstrom. Uniform resource names (urn) namespace definition mechanisms, 2002.
- [FGM<sup>+</sup>99] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext transfer protocol – http/1.1, 1999.
- [Fie00] R.T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California, Irvine, 2000.
- [FN96] N. Freed and Borenstein N. Multipurpose internet mail extensions (mime) part two: Media types, 1996.
- [GBS08] M. Gunes, B. Blywis, and J. Schiller. A hybrid testbed for long-term WSN studies. In *Int. Workshop on Sensor Network Engineering (IWSNE)*, 2008.
- [GFJ<sup>+</sup>09] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection tree protocol. Technical Report SING-09-01, Stanford University, 2009.
- [Goo] Federated Login for Google Account Users. <http://code.google.com/apis/accounts/docs/OpenID.html>. [Online; accessed: 02.10.2009].
- [GSR<sup>+</sup>04] Lewis Girod, Thanos Stathopoulos, Nithya Ramanathan, Jeremy Elson, Deborah Estrin, Eric Osterweil, and Tom Schoellhammer. A system for simulation, emulation, and deployment of heterogeneous sensor networks. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 201–213, New York, NY, USA, 2004. ACM Press. doi:http://doi.acm.org/10.1145/1031495.1031519.
- [HHP<sup>+</sup>08] I. Haratcherev, G. Halkes, T. Parker, O. Visser, and K. Langendoen. PowerBench: A Scalable Testbed Infrastructure for Benchmarking Power Consumption. In *Int. Workshop on Sensor Network Engineering (IWSNE)*, 2008.
- [HKWW06] Vlado Handziski, Andreas Köpke, Andreas Willig, and Adam Wolisz. Twist: a scalable and reconfigurable testbed for wireless indoor experiments with sensor networks. In *REALMAN '06: Proceedings of the 2nd international workshop on Multi-hop ad hoc networks: from theory to reality*, pages 63–70, New York, NY, USA, 2006. ACM. doi:http://doi.acm.org/10.1145/1132983.1132995.
- [Hum] Humboldt Universität zu Berlin. Berlin Roof Net. <http://www.berlinroofnet.de/>. [Online; accessed: 14.05.2009].

- [IAN] IANA. <http://www.iana.org/assignments/urn-namespaces/>. [Online; accessed: 02.10.2009].
- [JSF<sup>+</sup>06] D. Johnson, T. Stack, R. Fish, D. M. Flickinger, L. Stoller, R. Ricci, and J. Lepreau. Mobile emulab: A robotic wireless and sensor network testbed. *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, pages 1–12, April 2006. doi:10.1109/INFOCOM.2006.182.
- [Ken] Sun Cloud API. <http://kenai.com/projects/suncloudapis/>. [Online; accessed: 02.10.2009].
- [Ker] Kerberos Authentication Protocol. <http://web.mit.edu/Kerberos/>. [Online; accessed: 02.10.2009].
- [LLN<sup>+</sup>02] H. Lundgren, D. Lundberg, J. Nielsen, E. Nordstrom, and C. Tschudin. A large-scale testbed for reproducible ad hoc protocol evaluations. *Wireless Communications and Networking Conference, 2002. WCNC2002. 2002 IEEE*, 1:412–418 vol.1, Mar 2002. doi:10.1109/WCNC.2002.993531.
- [mer] Meraki Website. <http://meraki.com>. [Online: 14.05.2009].
- [Met07] Chad Metcalf. TOSSIM live: Towards a testbed in a thread. Master's thesis, Colorado School of Mines, 2007.
- [MKJK99] Robert Morris, Eddie Kohler, John Jannotti, and M. Frans Kaashoek. The click modular router. *SIGOPS Oper. Syst. Rev.*, 33(5):217–231, 1999. doi:http://doi.acm.org/10.1145/319344.319166.
- [mot] Motescope Testbed, University of California, Berkeley. <https://www.millennium.berkeley.edu/sensornets>. [Online; accessed: 14.05.2009].
- [Net] NetEquality Portland Deployment. <http://www.netequality.org/map.php?id=hacienda>. [Online; accessed: 14.05.2009].
- [OAu] OAuth. <http://oauth.net/>. [Online; accessed: 02.10.2009].
- [ome] Omega Testbed, University of California, Berkeley. <http://omega.cs.berkeley.edu/>. [Online; accessed: 14.05.2009].
- [Ope] OpenID. <http://openid.net/>. [Online; accessed: 02.10.2009].
- [OSSS05] M. Ott, I. Seskar, R. Siraccusa, and M. Singh. Orbit testbed software architecture: supporting experiments as a service. *Testbeds and Research Infrastructures for the Development of Networks and Communities, 2005. Tridentcom 2005. First International Conference on*, pages 136–145, Feb. 2005. doi:10.1109/TRIDNT.2005.27.
- [Pro] ProtoGENI. <http://www.protogeni.net/>. [Online; accessed: 02.10.2009].

- [RHF<sup>+</sup>08] Olof Rensfelt, Frederik Hermans, Christofer Ferm, Per Gunningberg, and Lars-Åke Larzon. An interactive testbed for heterogeneous wireless sensor networks. In *Demo Proceedings of the 4th International Conference on Distributed Computing in Sensor Systems (DCOSS 2008)*, 2008.
- [SBBD03] S. Sanghani, T.X. Brown, S. Bhandare, and S. Doshi. Ewant: the emulated wireless ad hoc network testbed. *Wireless Communications and Networking, 2003. WCNC 2003. 2003 IEEE*, 3:1844–1849 vol.3, March 2003. doi:10.1109/WCNC.2003.1200667.
- [Shi] Shibboleth (R). <http://shibboleth.internet2.edu/>. [Online; accessed: 02.10.2009].
- [SMP<sup>+</sup>08] Jeff Slipp, Changning Ma, Nagesh Polu, James Nicholson, Martin Murillo, and Sajid Hussain. Winter: Architecture and applications of a wireless industrial sensor network testbed for radio-harsh environments. *Communication Networks and Services Research, Annual Conference on*, 0:422–431, 2008. doi:http://doi.ieeecomputersociety.org/10.1109/CNSR.2008.42.
- [tut] Tutornet Testbed, University of Southern California. <http://enl.usc.edu/projects/tutornet>. [Online; accessed: 14.05.2009].
- [UT] Harvard University and BBN Technologies. Citysense testbed. <http://www.citysense.net/>. [Online; accessed: 07.10.2009].
- [VBV<sup>+</sup>05] Nitin H. Vaidya, Jennifer Bernhard, V. V. Veeravalli, P. R. Kumar, and R. K. Iyer. Illinois wireless wind tunnel: a testbed for experimental evaluation of wireless networks. In *E-WIND '05: Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, pages 64–69, New York, NY, USA, 2005. ACM. doi:http://doi.acm.org/10.1145/1080148.1080162.
- [WASW05] Geoffrey Werner-Allen, Pat Swieskowski, and Matt Welsh. Motelab: A wireless sensor networks testbed. In *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks (IPSN'05), Special Track on Platform Tools and Design Methods for Network Embedded Sensors (SPOTS)*, 2005.
- [Wis] Wisebed. <http://www.wisebed.eu/>. [Online; accessed: 02.10.2009].
- [ZV08] Yuhao Zheng and Nitin H. Vaidya. Repeatability of illinois wireless wind tunnel. Technical report, University of Illinois at Urbana-Champaign, 2008.