

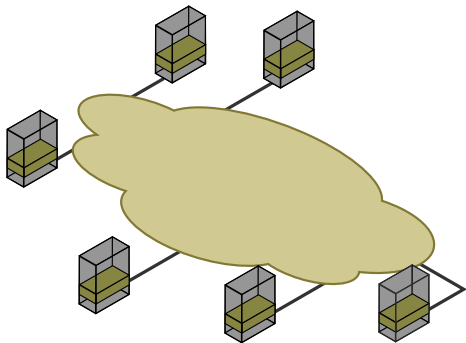
Overview of TopHat: Interconnecting the OneLab measurement infrastructures

Jordan Augé, Timur Friedman, Thomas Bourgeau (UPMC)


ISMA'2010 - 2nd AIMS workshop – February 8-10, San Diego, CA

- 1 Presentation of TopHat
- 2 Supporting PlanetLab applications
- 3 Future plans


- 1 Presentation of TopHat
- 2 Supporting PlanetLab applications
- 3 Future plans



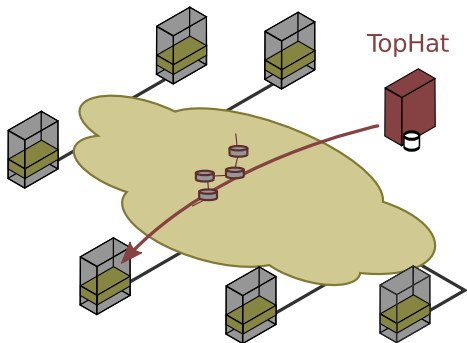
- ▶ PlanetLab testbed allows creation of overlay applications
 - ▶ P2P, CDN, etc.




- ▶ PlanetLab testbed allows creation of overlay applications
 - ▶ P2P, CDN, etc.
- ▶ the underlay is unknown
 - ▶ topology, delays, etc.
 - ▶ *and their evolution...*



- ▶ PlanetLab testbed allows creation of overlay applications
 - ▶ P2P, CDN, etc.
- ▶ the underlay is unknown
 - ▶ topology, delays, etc.
 - ▶ *and their evolution...*
- ▶ measurements are needed...



- ▶ PlanetLab testbed allows creation of overlay applications
 - ▶ P2P, CDN, etc.
 - ▶ the underlay is unknown
 - ▶ topology, delays, etc.
 - ▶ *and their evolution...*
 - ▶ measurements are needed...
-
- ▶ ... but a **measurement service** is better:
 - ▶ users can focus on developing the overlay instead of writing monitoring code
 - ▶ it allows the use of methods that reduce strain on the network,
 - ▶ it generally provides more efficient, correct and accurate results.



- ▷ There are many measurement systems
 - ▶ Ark/Archipelago, iPlane, ...
- ▷ Some are specifically designed for testbeds

- ▷ TopHat supports **the PlanetLab testbed experiment lifecycle**
 - ▶ from setup through completion
 - ▶ provides live measurements to the application
 - ▶ callbacks are used to communicate information to the application
- ▷ TopHat interconnects with OneLab partner systems
 - ▶ DIMES extends scope
 - ▶ ETOMIC provides high-precision measurements

- 1 Presentation of TopHat
- 2 Supporting PlanetLab applications
- 3 Future plans

- ▶ **four broad services**, following the **experiment lifecycle**.

Setup : help the user choose nodes before launching his/her experiment

Live : provide real-time information about the underlying network

Rewind : give access to historical data

Viz : allow visualization of experimental data

- ▶ two interfaces:
 - ▶ Web interface
 - ▶ an XML-RPC API + command-line tool



- ▶ TopHat provides topological information for the choice of nodes
- ▶ Leveraging PlanetLab topological and geographical diversity

Sample query:

Give me twenty relatively unloaded, reliable nodes that are each **at least five traceroute hops away** from each other with **stable routes** and **no load balancer** on the paths


- ▷ help researchers deploy and manage their experiments
 - ▶ on PlanetLab and future federated facilities
- ▷ uniform access to testbed data
 - ▶ system information from CoMon (reliability, load, etc.)
 - ▶ topological information for TopHat and interconnected systems
- ▷ for example, AS-level information in the default interface





PlanetLab Europe

- o Home
- News
- o About
- Join us
- Support
- Security Notice
- ▼ Documentation
 - o AUP
 - Guides
 - o API
 - Tutorials

Syndicate



My slice upmc_tophat

[Slice nodes](#)
[Site](#)
[Delete](#)



Details

1 Users







889 Nodes

889 nodes currently in upmc_tophat

« ‹ 4 5 6 7 **8** 9 10 11 12 › »

20 items/page

Search and

PEER	HOSTNAME	ST	R	L	AS	
PLC	pl1.rcc.uottawa.ca	boot	82	1.39	25826	
PLE	pl1.rennes.supelec.fr	failboot	79	0.22	2200	
PLC	pl1.ucs.indiana.edu	boot	82	0.58	87	
PLC	pl1.urm.edu	boot	0	n/a	3388	
PLC	pl1.snu.koren21.net	boot	82	0.54	9270	
PLC	pl2.higashi.ics.es.osaka-u.ac.jp	boot	82	5.26	4730	

jordan.auge@lip6.fr

- ▼ Logout
- ▼ My Account
- ▼ Sites
 - o My Site
- ▼ Nodes
 - o My Site Nodes
- ▼ My Slices
 - o Sirius
- ▼ About MyPLC
 - o PLCAPI doc
 - o IIMAPI doc



PlanetLab Europe

- Home
- News
- About
- Join us
- Support
- Security Notice
- Documentation
 - AUP
 - Guides
 - API
 - Tutorials

Syndicate

My slice upmc_tophat

Slice nodes Site Delete

Details

1 Users

889 Nodes

889 nodes currently in upmc_tophat

Search and

ST = status of the node
 R = avg. reliability (% uptime) over the last week
 L = avg. 5min. load over the last week
 AS = Autonomous System Number

PEER	HOSTNAME	ST	R	L	AS
PLC	pl1rcc.uottawa.ca	boot	82	1.39	25826
PLE	pl1rennes.supelec.fr	failboot	79	0.22	2200
PLC	pl1lucs.indiana.edu	boot	82	0.58	87
PLC	pl1lum.edu	boot	0	n/a	3388
PLC	pl1arku.koren21.net	boot	82	0.54	9270
PLC	pl2-nigashi.ics.es.osaka-u.ac.jp	boot	82	5.26	4730

Functions to get measurements (Get), manage the list of callbacks, etc.

Prototype of the Get function

RET = Get(Auth, Method, Timestamp, Input, Output, Callback)

Auth : authentication token (login/pass, session, etc.)

Method : eg. *traceroute, delay, active bandwidth*

Timestamp : date, interval, description (*now, today, latest*, etc.)

Input : a node/a set of nodes, a path/a set of paths

Output : subset of available fields, depends on Method

Callback : used for periodic measurements

Sample query:

```
list = ['planet2.elte.hu', 'planetlab-europe-02.ipv6.lip6.fr']  
Get(auth, 'latest', 'nodeinfo', {'hostname': list},  
    ['hostname', 'prefix', 'asn', 'as_name', 'platform_name'])
```

Result:

```
[{'hostname': 'planet2.elte.hu', 'prefix': '132.227.0.0/16',  
  'asn': '1307', 'as_name': 'FR-U-JUSSIEU-PARIS',  
  'platform_name': 'Team Cymru'},  
 {'hostname': 'planetlab-europe-02.ipv6.lip6.fr', 'prefix': '157.181.0.0',  
  'asn': '2012', 'as_name': 'ELTENET ELTENET',  
  'platform_name': 'Team Cymru'}]
```


- ▷ classical set of topological queries provided directly via the API



Akihiro Nakao, Larry Peterson, Andy Bavier,

A routing underlay for overlay networks,
in proceedings of SIGCOMM'03, pp.11–18, 2003.

- ▷ callbacks are provided for events, periodic and async. measurements
 - ▶ XML-RPC call, email, RSS, etc.

Example: alert me when the delay between two nodes changes by more than 20%

Sample query: traceroute on two different platforms:

```
path_list = [('planet2.elte.hu', 'planetlab-europe-02.ipv6.lip6.fr'),  
            ('ape.onelab.elte.hu', 'planetlab-europe-02.ipv6.lip6.fr')]  
Get(auth, 'now', 'traceroute', path_list,  
    ['src_ip', 'dst_ip', 'hops.ttl', 'hops.ip', 'hops.hostname', 'platform_name'])
```

Result:

```
[{'src_ip': '157.181.175.248', 'dst_ip': '132.227.62.19',  
  'hops': [  
    {'ttl': '1', 'ip': '157.181.175.254', 'hostname': None},  
    {'ttl': '2', 'ip': '157.181.126.45', 'hostname': 'taurus.taurus-leo.elte.hu'},  
    ...],  
  'platform_name': 'TDMI'},  
  
 {'src_ip': '157.181.175.247', 'dst_ip': '132.227.62.19',  
  'hops': [ ... ],  
  'platform_name': 'SONoMA'}  
]
```

Clients



Researchers



Applications

Services

Setup

Live

Rewind

Viz

Aggregators

TopHat

**Measurement
infrastructures**

TDMI

DIMES

ETOMIC

Team Cymru

Ongoing interconnection with OneLab partners and others

DIMES

- ▷ large-scale, measurements from outside of PlanetLab

ETOMIC, SONoMA

- ▷ high precision synchronized measurements

TDMI

- ▷ our own measurement agent, running in a slice

Team Cymru

- ▷ AS-level information

- 1 Presentation of TopHat
- 2 Supporting PlanetLab applications
- 3 Future plans

Accounting

- ▷ We need to keep track of who uses the system, when, why
 - ▶ to understand usage of the system and for engineering purposes and improvements
 - ▶ we are funded for the service and need to report about the usage of the system
 - ▶ there are scientific reasons to study users' behaviour itself
- ▷ Interconnecting multiple systems makes the accounting problem harder

Authentication

- ▷ Make it easy for users to access the different platforms. . .
- ▷ . . . but there are some constraints
 - ▶ access to sensitive data, ability to perform measurements
 - ▶ some issues have been solved for individual systems (eg. PlanetLab AUP, traceability)
 - ▶ more challenging in an interconnected environment

References



Thomas Bourgeau, Jordan Augé, Timur Friedman,
**TopHat: A topology information service to support applications
in a future-internet testbed,**
accepted to TridentCom'2010, 18-20 May 2010, Berlin, Germany.

Contact

`<firstname.lastname@lip6.fr>`

Website

<http://www.top-hat.info>

The service will be unveiled for TridentCom

We are interested in interconnecting with your platform !