


An Experimentation Workbench for Replayable Networking Research



emulab

Eric Eide, Leigh Stoller,
and Jay Lepreau

University of Utah,
School of Computing
NSDI 2007 / April 12, 2007



Repeated Research

An Experimentation Workbench for Replayable Networking Research

Eric Eidson
University of
{eeide, stoll}

Abstract

The networked and distributed system communities have an increasing need for research, but our current experimental methods are short of satisfying this need. Replayable experiments that can be re-executed, either forward or backward, yielding new results that can be compared to previous ones. Replayability requires capturing experiment processes and data, of course, and requires facilities that allow those processes to be examined, repeated, modified, and reused.

We are now evolving Emulab, our popular network testbed management system, to be the basis of a new experimentation workbench in support of realistic, large-scale, replayable research. We have implemented a new model of testbed-based experiments that allows people to move forward and backward through their experimentation processes. Integrated tools help researchers manage their activities (both planned and unplanned), software artifacts, data, and analyses. We present the workbench, describe its implementation, and report how it has been used by early adopters. Our initial case studies highlight both the utility of the current workbench and additional usability challenges that must be addressed.

1 Introduction

In the networking and operating system communities, there is an increasing awareness of the benefits of repeated research [5, 14]. A scientific community advances when its experiments are published, subjected to scrutiny, and repeated to determine the veracity of results. Repeated research not only helps to validate the conclusions of studies, but also to expand on previous conclusions and suggest new directions for research.

To repeat a piece of research, one first needs access to the complete records of the experiment that is to be redone. This obviously includes the results of the experiment—not only the final data products, but also the “raw” data products that are the bases for analysis. Data sets like those being collected in the networking community [2, 3, 6, 22] allow researchers to repeat analyses,

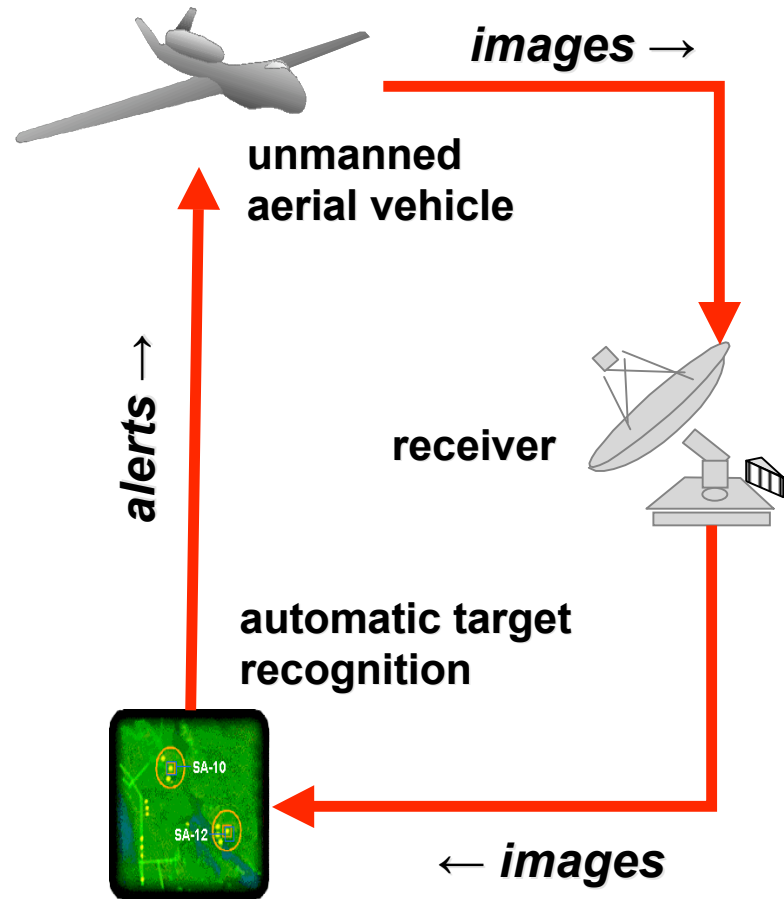
“A scientific community advances when its experiments are repeated...”

Translation: “I have trouble managing my own experiments.”



Example From My Past

- a distributed, real-time application
- evaluate improvements to real-time middleware
 - *vs. CPU load*
 - *vs. network load*
- **4 research groups**
- **x 19 experiments**
- **x 56 metrics**
- **use Emulab**





A Laboratory Is Not Enough

- testbeds give you lots of resources...
- ...but offer little help in *using* those resources

- *package / distribute / configure / instrument / init / execute / monitor / stop / collect / analyze / archive / revise / repeat*





What's Missing: Workflow

- current network testbeds
 - ...*manage the “laboratory”*
 - ...*not the experimentation process*
 - *i.e., scientific workflow*
- → a big problem for large-scale activities



Need

- my experiment needs...

- *encapsulation*
- *automation*
- *instrumentation*
- *preservation*

package / distribute /
configure / instrument /
init / execute / monitor /
stop / collect / analyze /
archive / revise / repeat

- benefits

- *verify previous results*
- *establish base for new research*
- *my own, or someone else's*

***repeatable
research***





Opportunity

- get the lab manager to help us out!
 - *integrated support for experimental procedures*
 - *resources + encapsulation + automation*
 - *framework: rapid start & common basis*
- *manage scientific workflow, but also manage lab*




Experimentation Workbench

- an environment for “replayable research”
 - *experiment management + experiment execution*
 - *(but really: help me manage my work)*
 - *all Emulab-managed devices, incl. PlanetLab slivers, ...*
- initial design, implementation, and evaluation
 - *new model of testbed-based experiments*
 - *prototype implementation*
 - *case studies*
 - *lessons learned*



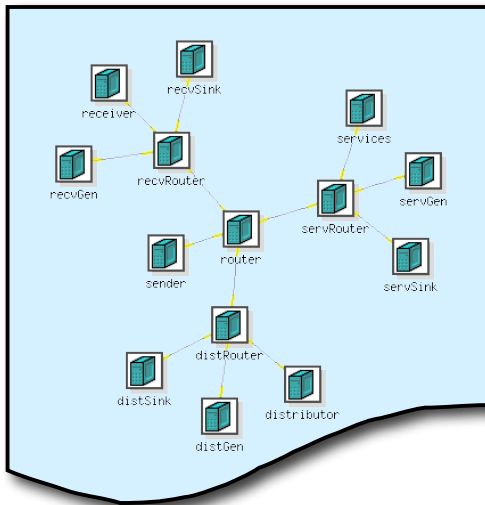
Workbench




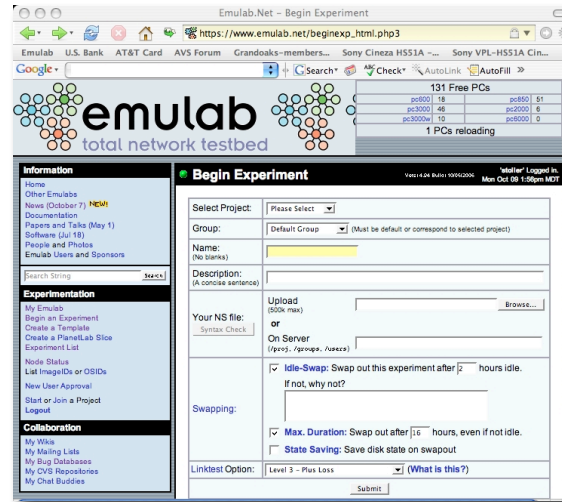
emulab



Classic “Experiments”



*topology +
SW (by reference) +
events*





Problems

- definition versus instance
- related experiments
- multiple trials per session
- data management
 - *instrumentation, collection, archiving, analyzing*
- ecosystem
 - *topology, software, config, input data, ...*
- evolution over time



New Model

- template



- instance



- run



- activity



- record



- *divide and conquer*

- *separate the roles that an experiment plays*

- *evolve the new abstractions*

- *build on what testbed users already know and do*



Template

- *template*



- instance



- run




- activity



- record



- *a “repository”*



- *definition role of the classic “experiment” notion*

- resources by value

- *files, scripts, DB tables, disk images, ...*

- resources by reference

- prototype: implemented with Subversion (*user-hidden*)



Templates vs. Experiments

- a template is like a classic Emulab experiment, but a template has...
 - *datastore (file repository)*
 - *parameters*
 - *multiple instances*
 - *metadata*
 - *history*



Template History


Template BT-V12 (10080/13)

Vers: 4.104 Build: 02/15/2007

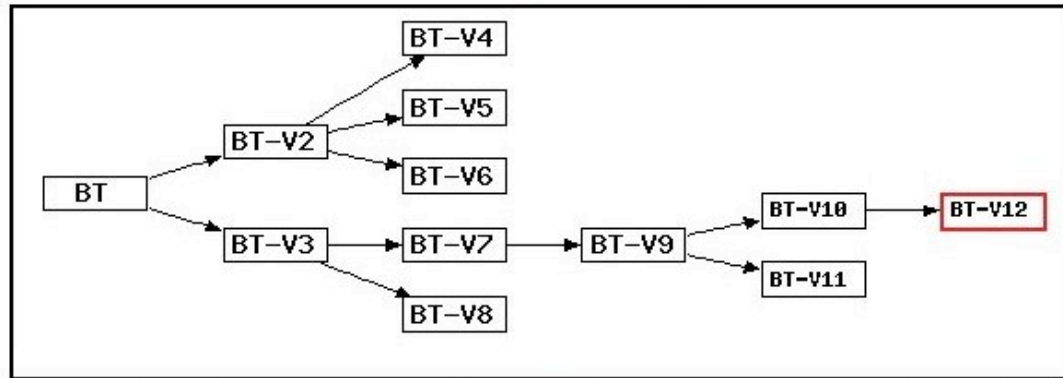
'stoller' Logged in.
Thu Feb 15 2:08pm MST

Template Options

- Activate Template
- Modify Template
- Instantiate Template
- Add Metadata
- Browse Datastore
- View Records



Topology **NS File** History



Hide Template Recursive? Show Hidden Templates Zoom Out Zoom In

Details

GUID:	10080/13
ID:	BT-V12
Project:	testbed
Group:	testbed
Creator:	stoller
Created:	2007-02-15 12:51:40

Parameters

Name	Default Value	Description
BSD	FBSD61-STD	FreeBSD OS version
MaxNodes	6	Maximum number of nodes
RHL	RHL90-STD	RedHat OS version



Instantiating a Template

parameters

ID: (alphanumeric, no blanks)	<input type="text" value="TemplateTest"/>
Swapping:	<input checked="" type="checkbox"/> Idle-Swap: Swap out this experiment after <input type="text" value="2"/> hours idle. If not, why not? <input type="text"/>
	<input checked="" type="checkbox"/> Max. Duration: Swap out after <input type="text" value="16"/> hours, even if not idle.
	<input type="checkbox"/> State Saving: Save disk state on swapout
Formal Parameters:	CLIENT_COUNT <input type="text" value="2"/> DURATION <input type="text" value="60"/> HWTYPE <input type="text" value="pc850"/> or XML file: On Server <input type="text"/> (/proj, /groups, /users)
Use this text area for an (optional) description: <input type="text"/>	
<input type="checkbox"/> Batch Mode Instantiation (See Tutorial for more information)	
Linktest Option:	<input type="text" value="Level 3 -"/> (What is this?)
<input type="button" value="Instantiate"/>	

template instances can also be created programmatically



Template Instance

- template



- *instance*



- run



- activity



- record



- *a container of testbed resources*
- *resource-owner role of classic “experiment” notion*
- a transitory object
 - *created and destroyed by users and activities*
- nodes & network
 - *files from the datastore*
- *database for user*



Run & Activity

- template



- instance



- *run*



- *activity*



- record



- *run: a container of a user-defined “unit of work”*
 - defines a context
 - a “trial”
 - one / serial / parallel
- *activity: a process, script, workflow, ... within a run*
 - events & data come from activities in a run
- runs and activities can be scripted or interactive
- prototype: implemented via agents & events



Record

- template



- instance



- run



- activity



- *record*



- *the “flight recorder” of a run*
 - *parameter values*
 - *input & output files, DBs*
 - *raw data & derived data*
 - *template’s by-reference resources*
 - *dynamically recorded events*





Record Repository

view, export, and replay

Template History 'stoller' Logged in.
Ver: 1.4.24 Build: 10/02/2006
Mon Oct 09 3:29pm MDT


Experiment Template **10000/3**

Template History (Swapins)

Expand	EID	UID	Start Time	Stop Time	Show	Archive	Export	Replay	Load DBs
▶	TT	stoller	2006-08-01 12:28:12	2006-08-01 14:04:01	●	●	●	●	●
▶	TT	stoller	2006-08-07 14:09:03	2006-08-08 06:10:03	●	●	●	●	●
▼	TT	stoller	2006-08-08 07:09:38	2006-08-08 07:24:51	●	●	●	●	●

Show	Archive	RunID	ID	Start Time	Stop Time	Description
●	●	TT	1	2006-08-08 07:09:38	2006-08-08 07:20:14	
●	●	TT-R2	2	2006-08-08 07:20:32	2006-08-08 07:23:13	
●	●	TT-R3	3	2006-08-08 07:23:51	2006-08-08 07:24:51	

Evaluation and Lessons Learned



emulab



Evaluation

- how to evaluate?
 - *new capabilities* → *user studies*
- **goal:** early feedback about design & impl.
- **approach:** three case studies
- **outcome:** specific & general lessons learned



Study 1: Flexlab Development

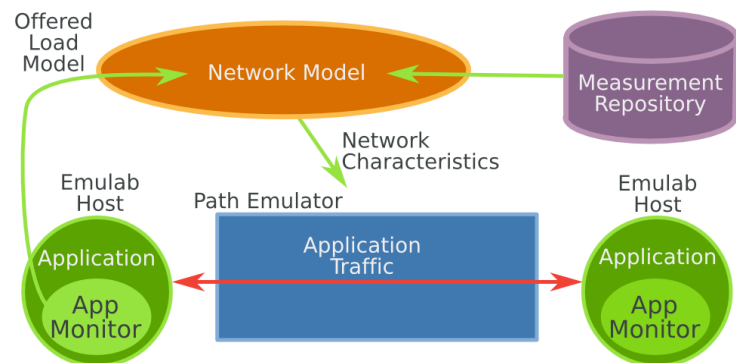
- replace ad hoc experiment management

- originally:

- *a configurable ns file*
- *start/stop trial scripts*
- *“scaffold” in CVS*
- *manual archiving*
- *destructive modification*

- now:

- *templates & params*
- *runs, start/stop hooks*
- *scaffold & results in WB*
- *automatic archiving*
- *preserved history*

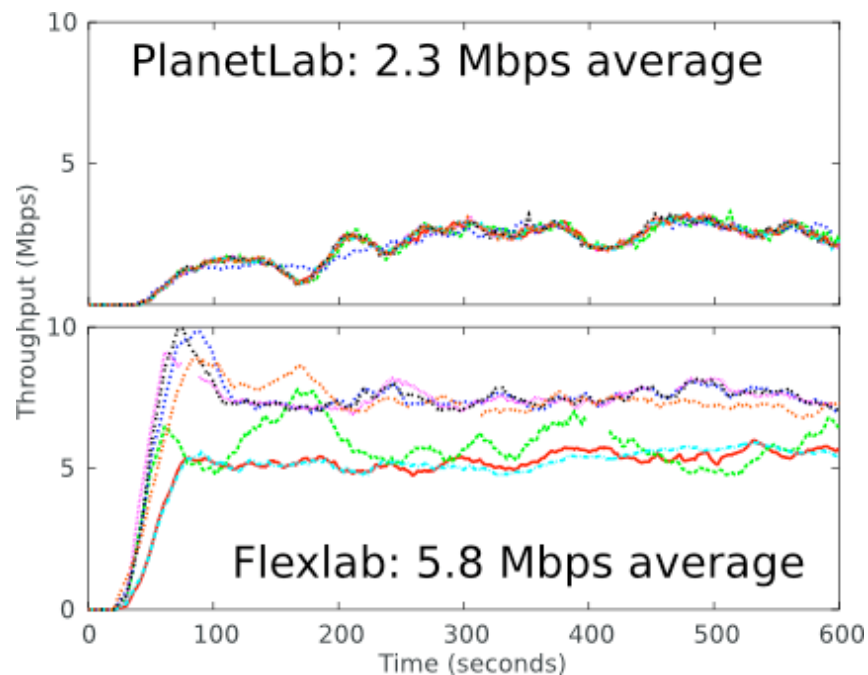


Conclusion: the new model “fits” developers’ model



Study 2: Flexlab Use

- study BitTorrent on Flexlab and PlanetLab
- outcome:
 - *parameterization*
 - *utilized per-run database*
 - *team communication*
 - *results for publication*
- *stress point: latency*
- *stress point: node failure*





Lessons: Storage

- initial philosophy: “store everything”
 - *templates + results + history + metadata + ...*
- space efficiency + group commits
 - → *Subversion*
- cognitive overload
 - → *careful UI design*



Space and Time

	Record size (MB)	Stored in repo. (MB)	Elapsed time (m)
BitTorrent	31.0	18.9	7.0
GHETE	70.6	19.8	14.4

- **solution:** pipeline record-making with user activities
- **new problem:** isolation
- **new approach:** branching file systems



What Users Want



- deletion!
 - *not a space concern*
 - *cognitive clutter — “junk”*
 - *privacy — “mistakes”*
- a range of options is required
- “true deletion” is a new requirement



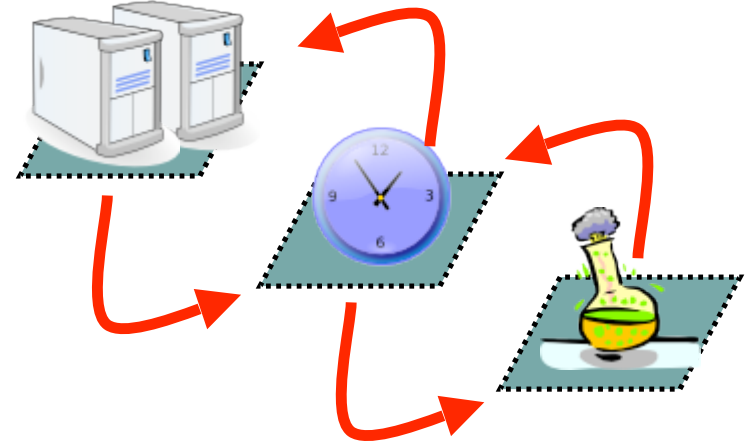
Lessons: The Model

- initial philosophy: “divide and conquer”
 - *more kinds of entities*
 - *describe notions and relationships*
- experience:
 - *new model does map to users' abstractions*
 - *captures separations and connections...*
 - *...but not “life cycle” concerns*



“Life Cycle” Concerns

- multiple levels of abstraction
 - *instance: “the lab”*
 - *run & activity: “the work”*
- intertwined & concurrent
 - *workbench must manage experiments and the lab*
 - *a key difference with “ordinary” scientific workflow systems*
- approach: further refine and enhance our model
 - *e.g., adopt features of Plush [Albrecht et al., OSR 40(1)] or SmartFrog [Sabharwal, ICNS ‘06]*





Summary

- **goal: better research** ← **better process tools**
- **experiment management + experiment execution**
- **prototype** builds on existing testbed infrastructure
 - *model maps pretty well to user notions*
- **experience:** strong integration is required
 - *...for overlapping activities safely*
 - *...for lab management + experiment management*
 - *...for making it user-friendly in practice*




Status

- “alpha test” stage
- internal users
- select external users...
 - *mail to testbed-ops@emulab.net*



<http://www.emulab.net/>



emulab

**Thank you!
Questions?**