The Advanced Profile-Driven Testbed

Robert Ricci

June 23, 2014

A platform for sharing research artifacts and environments

A facility for building testbeds tailored to specific domains

A choice...


WENGER

Role for Domain Experts

Profiles: Packaged experiments

Packaged testbeds

http://aptlab.net/p/tbres/nsdi14

AptLab – Instantiate a Prof

https://www.aptlab.net/instantiate.php?project=tbres&profile=nsdi14

Home | Manual | apt | Sign Up | Login

What is Apt?

Run an Experiment using profile "nsdi14"

Fill out the form below to run an experiment using this profile:

> A VM containing the data set and analysis tools for the paper "Operational Experiences with Disk Imaging in a Multi-Tenant Datacenter", which appeared in NSDI '14. Everything necessary for reproducing all figures in the paper is included.

When you click the "Create" button, the virtual or physical machines described in the profile will be booted on Apt's hardware

Pick a user name

Your email address

Optional: your SSH public key (upload a file or paste it in the text box)

Choose File No file chosen

Paste in your ssh public key.

Create!

AptLab - Experiment Statu ×

https://www.aptlab.net/status.php?uuid=e7a29346-faa4-11e3-8053-001143e453fe

Home Manual

apt

Sign Up Login

Please wait while we get your experiment ready

URN: urn:publicid:IDN+emulab.net+slice+aptrob11-QV433

State: created

Profile: nsdi14

Expires: Today at 3:06 AM (in 3 hours)

Extend Terminate

AptLab – Experiment Statu ×

https://www.aptlab.net/status.php?uuid=e7a29346-faa4-11e3-8053-001143e453fe

Home | Manual | apt | Sign Up | Login

Profile Expires: Today at 3:06 AM (in 3 hours)

Profile Instructions

Topology View | List View | Manifest | host

```
Imaging in a Multi-Tenant Datacenter" (NSDI '14).

The data set described in the paper is available in a SQL database:
    mysql image_usage

The tools used for analysis of the data are available under /home/nsdi14.
You can reproduce the statistics cited in the body of the paper, the
figures, and the tables, with:
    cd /home/nsdi14
    ./all-stats
    ./all-figures
    ./all-tables

You can view the generated figures by visiting the URL:
    http://host.aptrob11-qv433.emulab-net.utahddc.geniracks.net/

For further information, please see the file:

    /home/nsdi14/README
------------------------------------------------------------------------------
aptrob11@host:~$
```

https://www.aptlab.net/status.php?uuid=e7a29346-faa4-11e3-8053-001143e453fe#profile_status_collapse ? Join the Help Forum

AptLab - Experiment Statu × | Operational Experiences w ×

host.aptrob11-qv433.emulab-net.utahddc.geniracks.net

Operational Experiences with Disk Imaging in a Multi-Tenant Datacenter

Welcome to an APT profile instance for the paper "Operational Experiences with Disk Imaging in a Multi-Tenant Datacenter" (NSDI '14). This page is being served from a VM set up to replicate the experiment described in that paper. If you brought up this VM, you can log in to it via the profile page. If you did not bring it up, but would like to, you can instantiate one yourself.

Figures from the paper are available:

1. Number of requests for facility and user images
2. Lifespans of user images
3. Usage gap for user images
4. Histogram of similarity between images
5. Variation of facility image popularity over time
6. Daily working set size
7. Weekly working set size
8. Number of images used, for subsamples of the userbase
9. Daily working set size, for subsamples of the userbase
10. Image request distribution compared to exponential
11. Numbers of images users request
12. Profile of image types used by heaviest users
13. Percentage of requests for default, facility, and user images
14. Ratio of satisfied requests when free pool size is varied
15. Ratio of satisfied requests when reload rate is varied
16. Network traffic required to transfer image deltas

Those figures may be inspected, regenerated, or modified by logging in to the VM (as described above).

AptLab - Experiment Statu × host.aptrob11-qv433.emu ×

host.aptrob11-qv433.emulab-net.utahddc.geniracks.net/figure-12.pdf

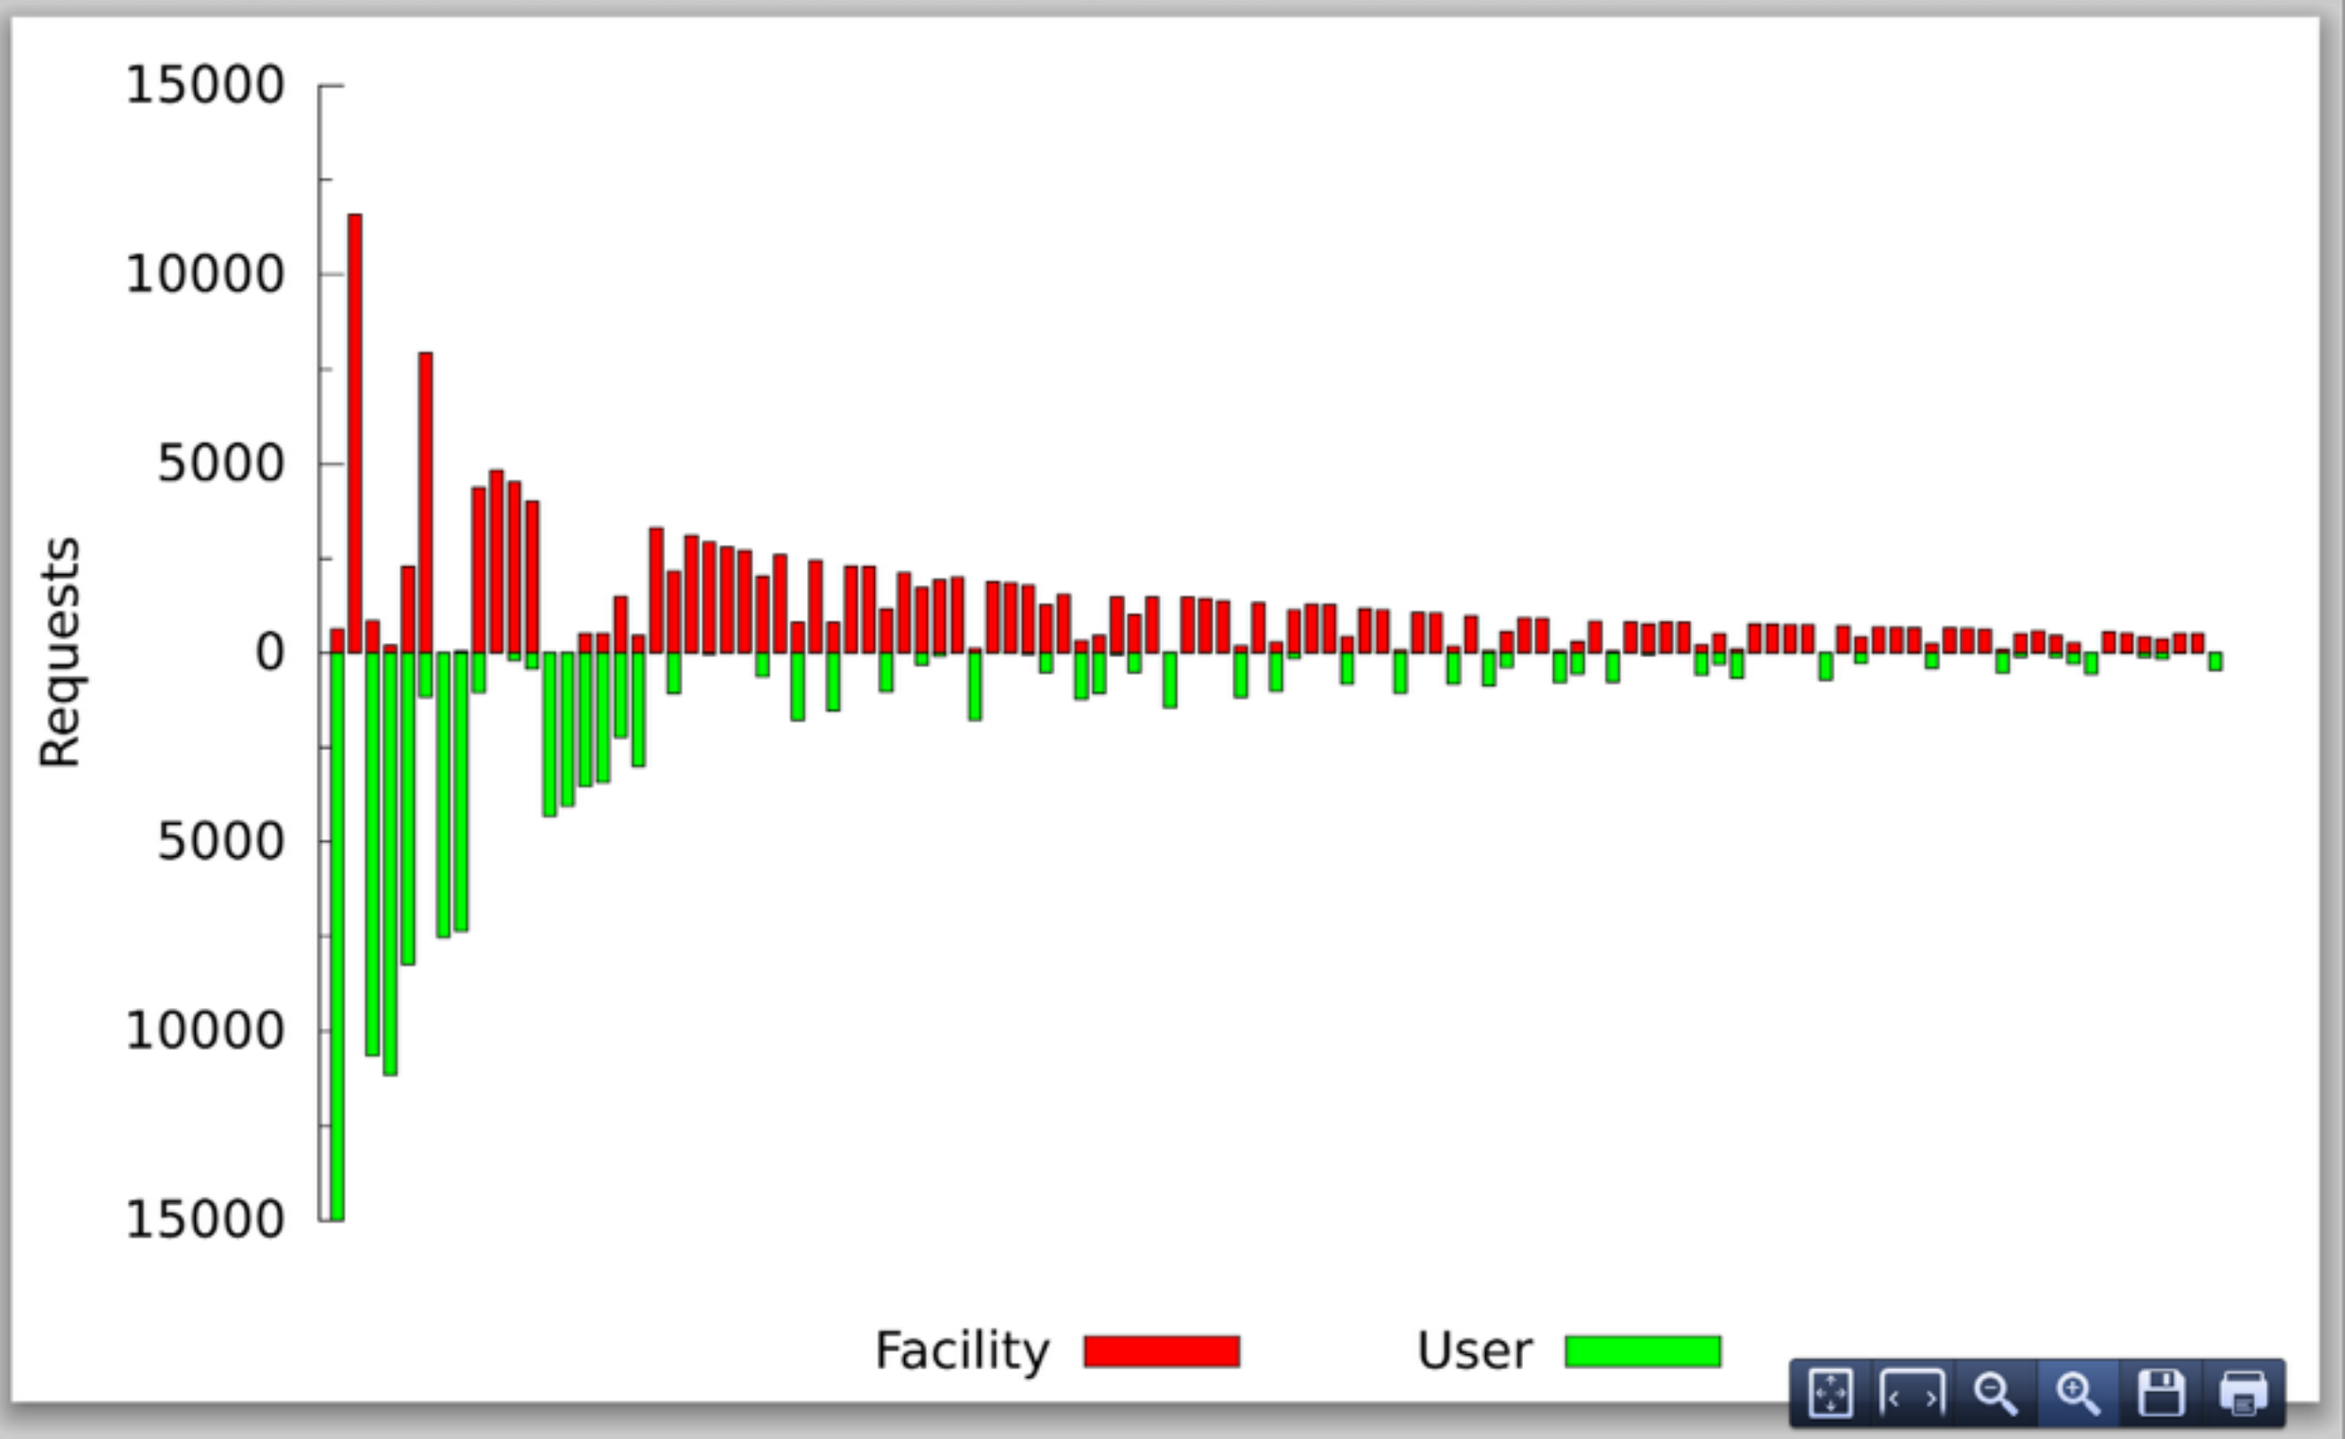

It's Just GENI*

Profile= RSpec

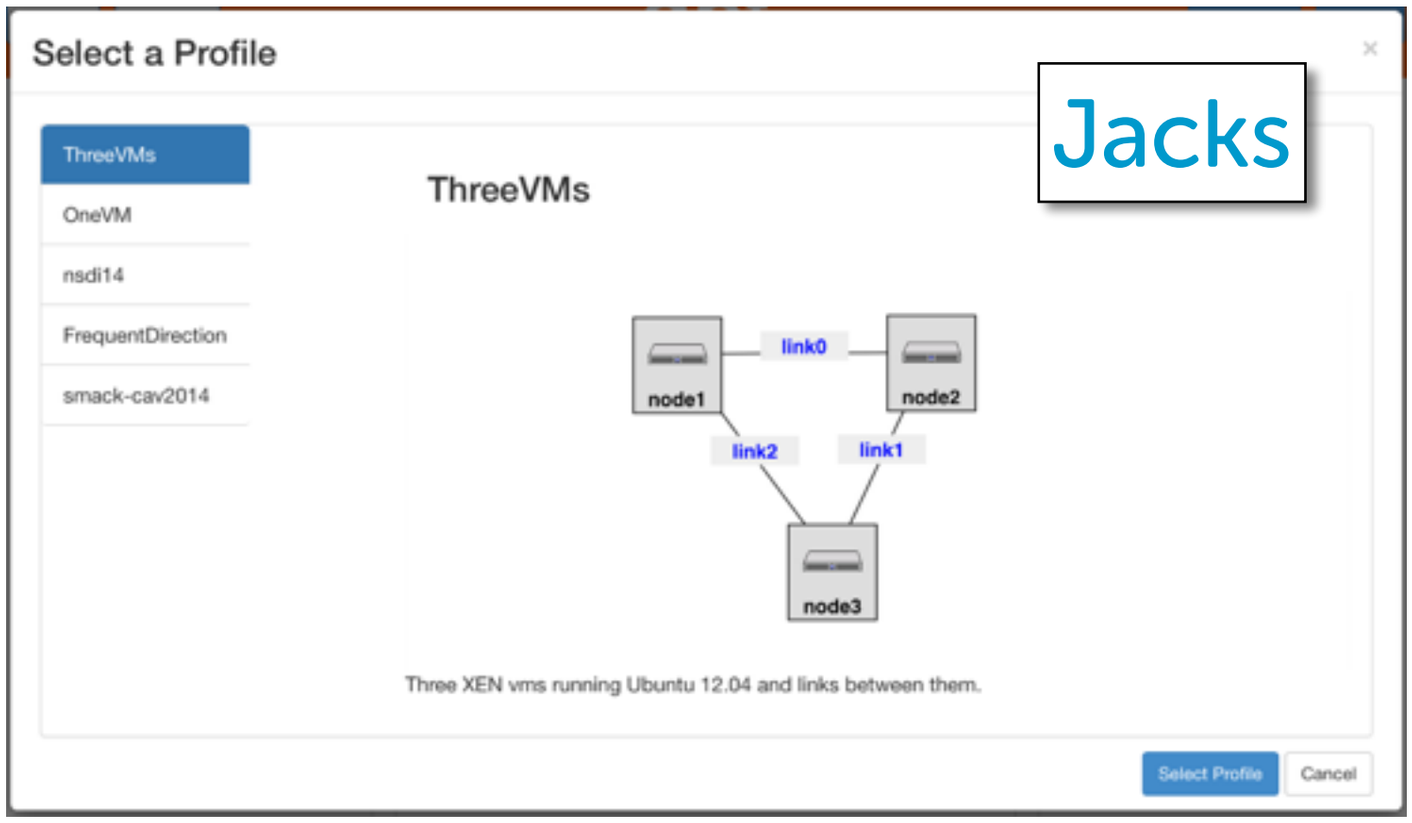

Jacks

Member Authorities

GENI Portal

Emulab

Apt Guests

*with some PGENI-specific features


NCSA
IBM

56 Gbps Ethernet/IB

Early Users

SMACK: Decoupling Source Language Details from Verifier Implementations*

Zvonimir Rakamarić[1] and Michael Emmi[2]

[1] School of Computing, University of Utah, USA
zvonimir@cs.utah.edu
[2] IMDEA Software Institute, Spain
michael.emmi@imdea.org

Abstract. A major obstacle to putting software verification research into practice is the high cost of developing the infrastructure enabling the application of verification algorithms to actual production code, in all of its complexity. Handling an entire programming language is a huge endeavor that few researchers are willing to undertake; even fewer could invest the effort to implement a verification algorithm for many source languages. To decouple the implementations of verification algorithms

Formal Software Verification

interpretation. Our initial experience in verifying C-language programs is encouraging: SMACK is competitive in SV-COMP benchmarks, is able to translate large programs (100 KLOC), and is being used in several verification research prototypes.

1 Introduction

This is a reformatted version of the paper that appears in SIGCOMM's proceedings

Using RDMA Efficiently for Key-Value Services

Anuj Kalia Michael Kaminsky† David G. Andersen
Carnegie Mellon University †Intel Labs
{akalia,dga}@cs.cmu.edu michael.e.kaminsky@intel.com

ABSTRACT

This paper describes the design and implementation of HERD, a key-value system designed to make the best use of an RDMA network. Unlike prior RDMA-based key-value systems, HERD focuses its design on reducing network round trips while using efficient RDMA primitives; the result is substantially lower latency, and throughput that saturates modern, commodity RDMA hardware.

HERD has two unconventional decisions: First, it does not

supports up to 26 million key-value operations per second with 5 µs average latency. Notably, for small key-value items, our full system throughput is similar to native RDMA read throughput and is over 2X higher than recent RDMA-based key-value systems. We believe that HERD further serves as an effective template for the construction of RDMA-based datacenter services.

table access? To answer this question, we first evaluate the performance that, with sufficient attention to engineering, can be achieved by each of the RDMA communication primitives. Using this understanding, we show how to use an unexpected combination of methods and system architectures to achieve the maximum performance possible on a high-performance RDMA network.

Our work is motivated by the seeming contrast between the fundamental time requirements for cross-node traffic vs. CPU-to-memory lookups, and the designs that have recently

RDMA read bypasses many potential sources of overhead, such as servicing interrupts and initiating control transfers, which involve the host CPU. In this paper, we show that there is a better path to taking advantage of RDMA to achieve high-throughput, low-latency key-value storage.

A challenge for both our and prior work lies in the lack of richness of RDMA operations. An RDMA operation can only

Big Data Computer Systems — Hari Sundar

Hari Sundar

home research teaching blog code

Big Data Computer Systems

Fall 2014

Mon,Wed 1:25pm-2:45pm
MEB 3147

Catalog number: CS 5965/6995

Overview

The exponential increase in the quantity and quality of measurements and data holds tremendous promise for data-driven scientific discovery. However, much of this data remains

real world. Big data is a broad concept that covers many aspects of computer science. This course will focus on the computer systems aspect—for instance, how various parts of a big software, and applications) are put together? What high performance, scalability, and reliability in

system from the Flux research group.

course but CS 3505 or equivalent programming structor if you are not sure whether you possess me projects might require knowledge of Numerical nts are not required to have this background, but

Big Data Class

Hadoop

Ethernet

AptLab – Manage Profile

Please wait while we make your disk image

Preparing	Imaging	Finishing	Ready

Node Status	TBSETUP
Image Size	608.17 MB

Instructions

Log into your VM and poke around. You have root access via `sudo`. Any work you do in the VM will be lost when it terminates.

Steps

Type	ID	Description
This Grid Is Empty		

List on the home page for anyone to view.

Who can instantiate your profile?

- *Anyone* on the internet (guest users)
- Only registered users of the APT website
- Only members of your project

AptLab – Instantiate a Profile

Select a Profile

- ThreeVMs
- OneVM
- nsdi14
- FrequentDirection
- smack-cav2014

FrequentDirection

In our paper 'Improve Practical Matrix Sketching with Guarantees' we introduce three algorithms Parametrized Frequent Direction, Space Saving Direction, Compensative Frequent Direction. We provide code and datasets for our new algorithms, as well as existing algorithms for Frequent Direction, ISVD, Random Projections, Hashing, Sampling. We encourage other researchers to reproduce our results and any other variation they would like to perform.

Code + Data in /local/FreqDir

Select Profile | Cancel

Direction, ISVD, Random Projections, Hashing, Sampling. We encourage other researchers to reproduce our results and any other variation they would like to perform.

aptlab.net