

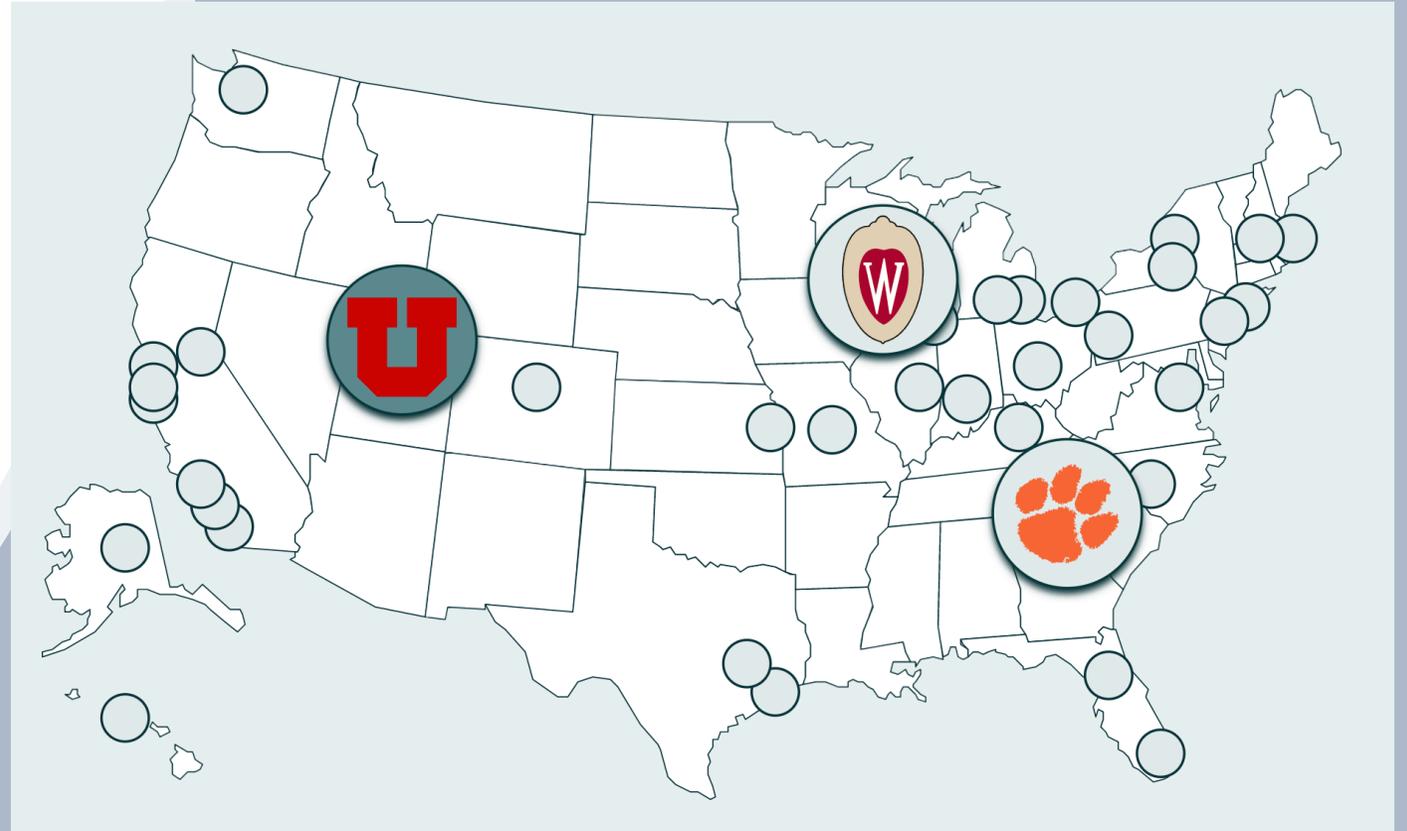
The Design and Operation of CloudLab

Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar, and Prabodh Mishra



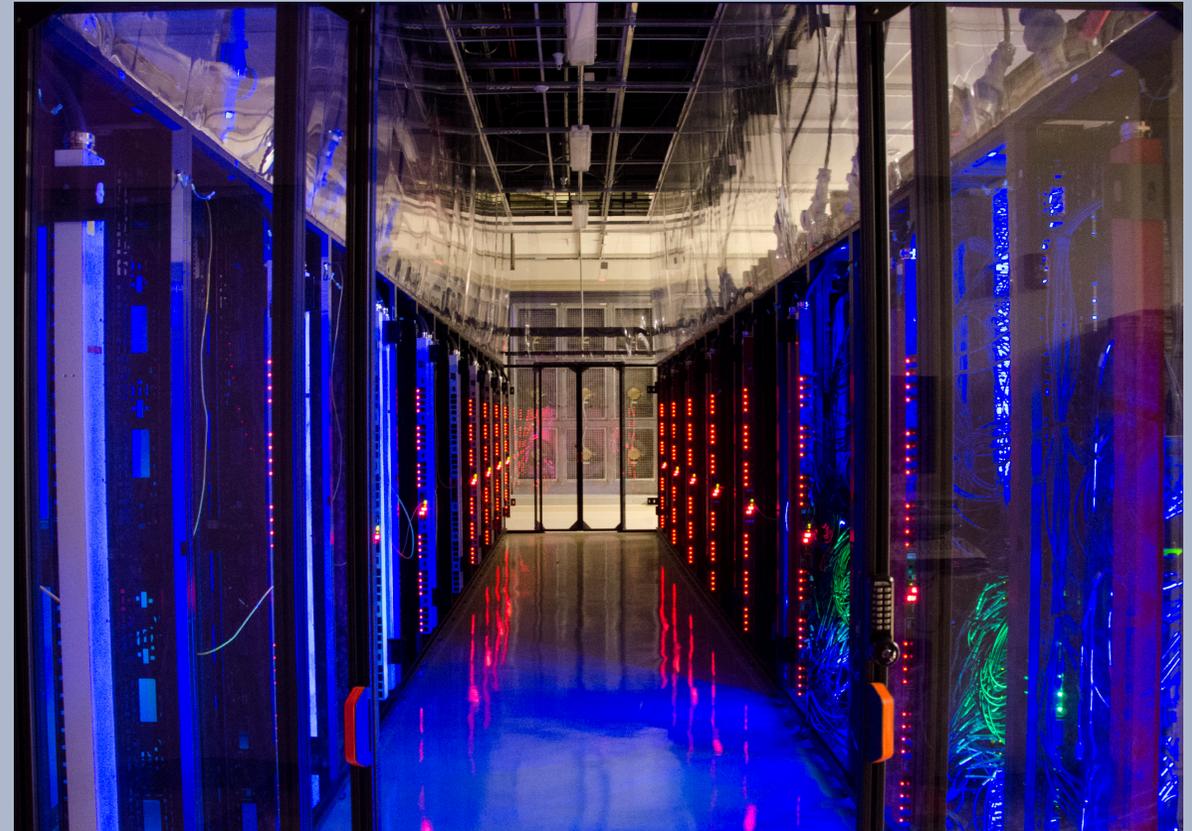
What is CloudLab ?

Distributed testbed for
cloud computing
and systems research



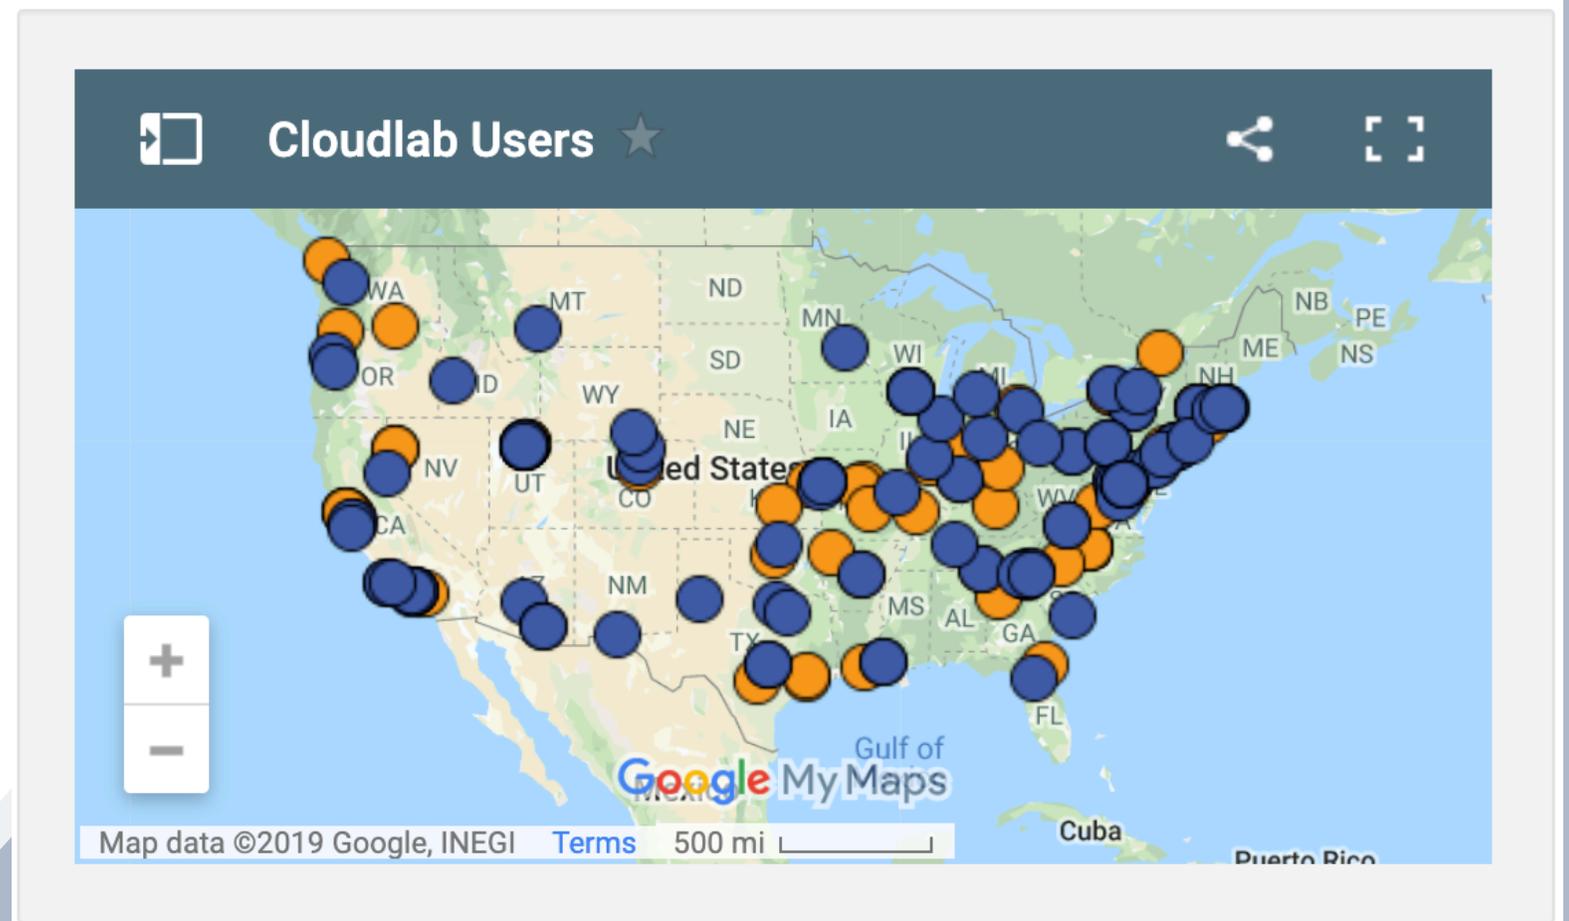
What is CloudLab ?

where experiments are run on
real, physical hardware
(not Virtual Machines)



What is CloudLab?

with >4K users
from many U.S. institutions



What is CloudLab ?

**A place to
experiment with
your own clouds
and distributed
systems**



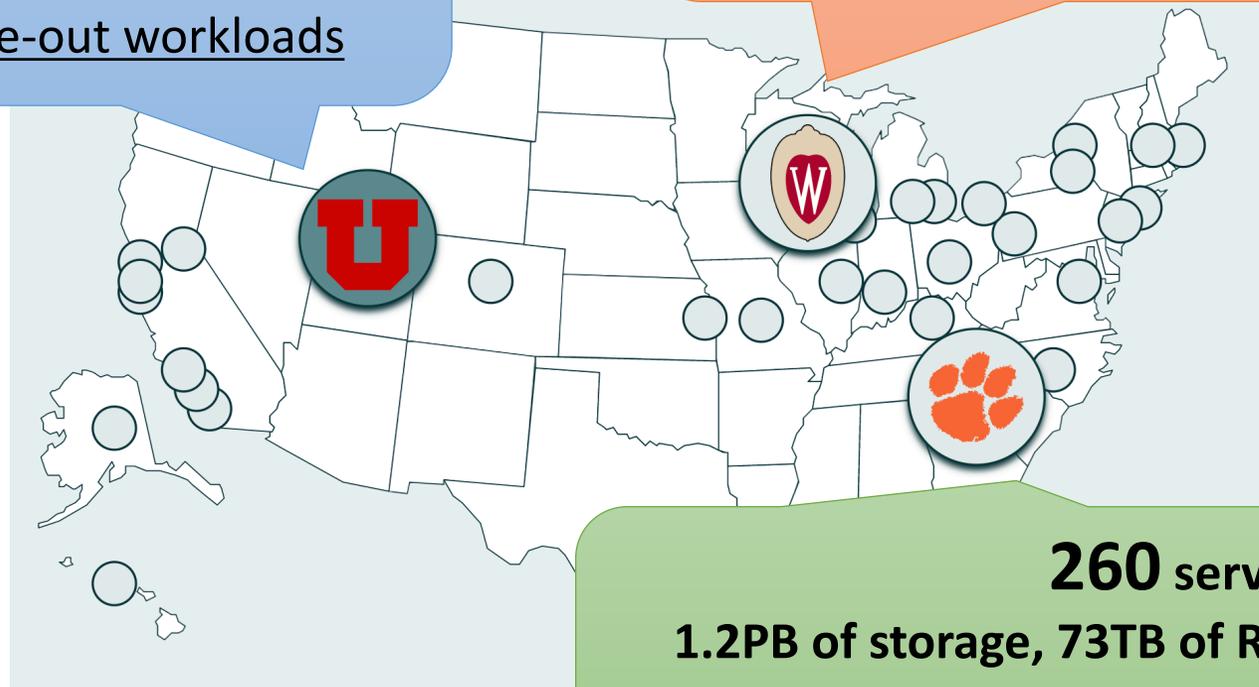
CloudLab Hardware

977 servers
10.6K Cores

Focus: scale-out workloads

527 servers
1PB of storage, SSD on every server

Focus: networking and storage work



260 servers

1.2PB of storage, 73TB of RAM, QDR Infiniband

Focus: analytics and high-performance workloads

CloudLab Hardware

977 servers
10.6K Cores

Focus: scale-out workloads

527 servers

1PB of storage, SSD on every server

Focus: networking and storage work

**19 Hardware
Types**

260 servers

1.2PB of storage, 73TB of RAM, QDR Infiniband

Focus: analytics and high-performance workloads

CloudLab Hardware

More at: <https://cloudlab.us>

What do researchers use CloudLab for?

Networking	30%
Security	16%
Storage	11%
Applications	10%
Computing	9%
Virtualization	8%
Databases	7%
Middleware	4%
Energy & Power	2%
Other	15%

Low-level access to hardware

Specific features

Performance isolation

* Based on 93 papers from 2017-2018

Why study CloudLab ?

To better understand how well it serves **diverse researchers' needs**

To discern the **impact** of design decisions and associated **tradeoffs**

To offer the insights that would benefit the **design of other testbeds and IaaS facilities**

I need 5 servers
now



I need 10 servers of
type **c240g5** with **GPUs**



I need servers
connected to a
programmable switch



I need **100** servers
next time I can get them



CloudLab needs to:

Satisfy **diverse user needs**
(scale, time, and features)

Help users select
feasible configurations

Return **meaningful errors** when
request or facility issues occur

I need 5 servers
now



I need 10 servers of
type **c240g5** with **GPUs**



I need servers
connected to a
programmable switch



I need **100** servers
next time I can get them



This study's questions:

Does CloudLab fulfill these goals?

How does CloudLab achieve them?

How can we generalize the lessons learned?

User Perspective

Satisfy **diverse user needs**
(scale, time, and features)

Help users select
feasible configurations

Return **meaningful errors** when
request or facility issues occur

“Under the hood”



Reservation System



Constraints System



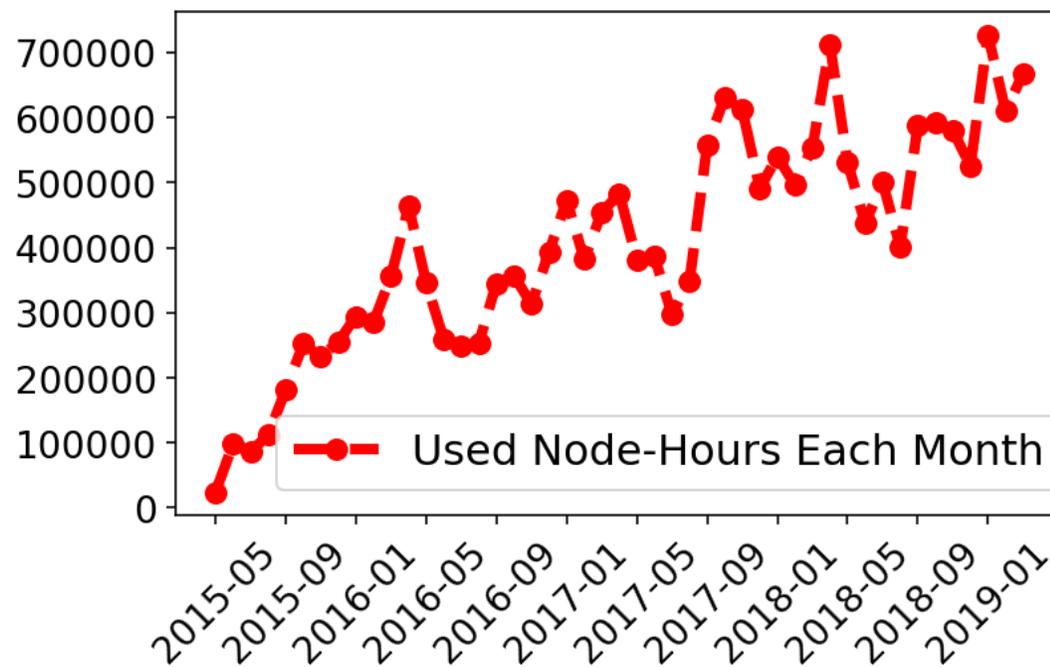
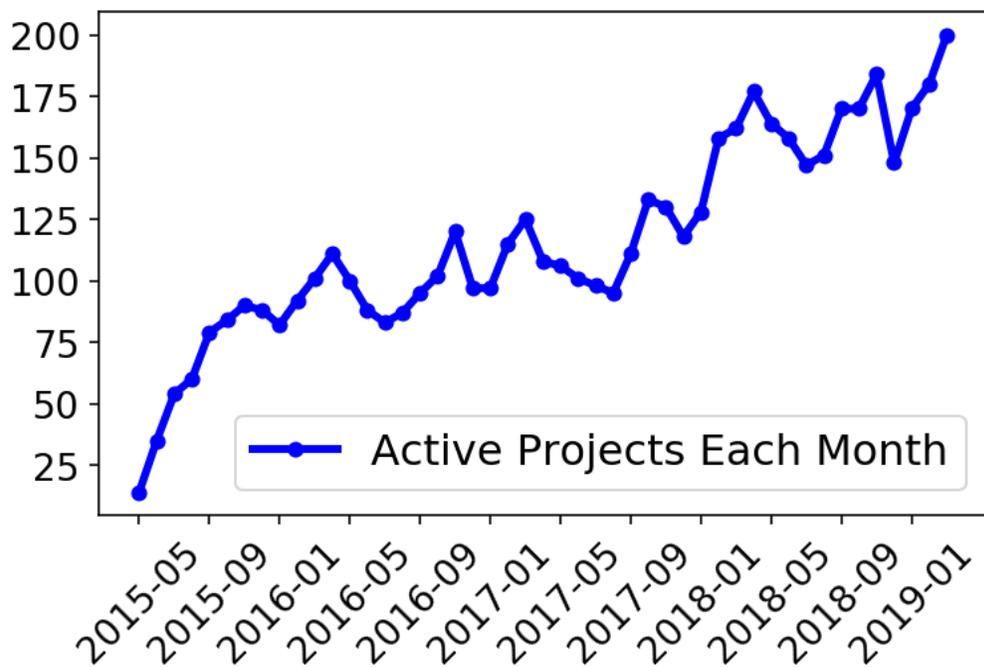
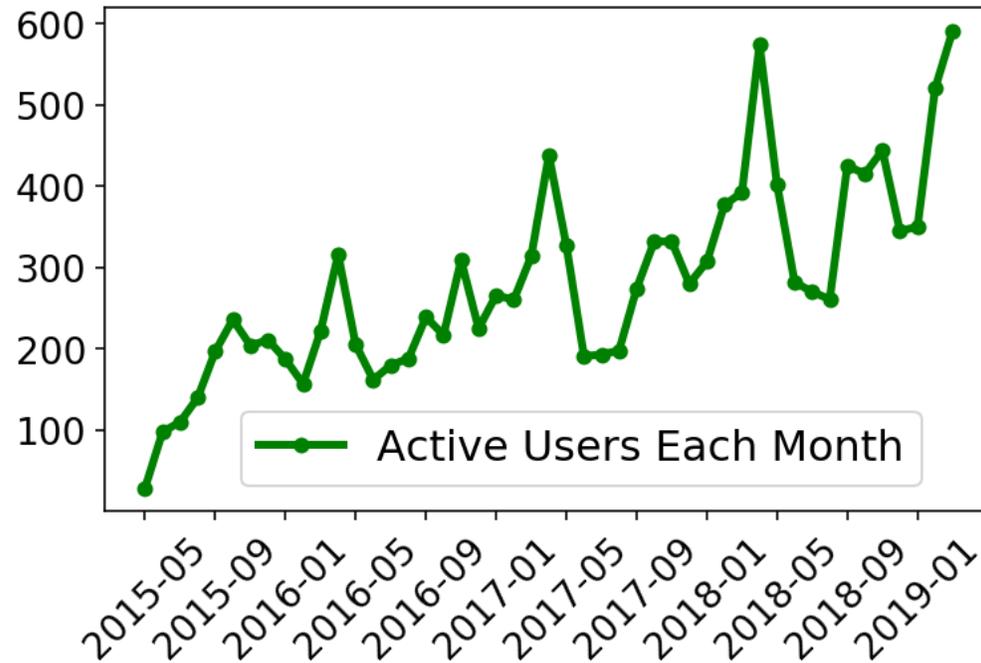
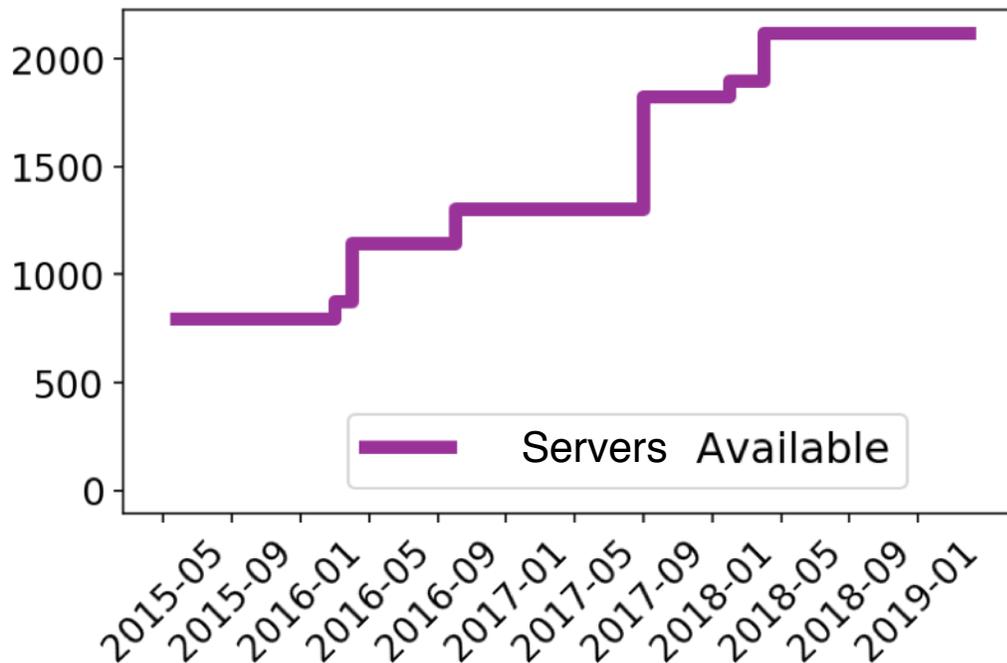
Error Reporting

 Reservations

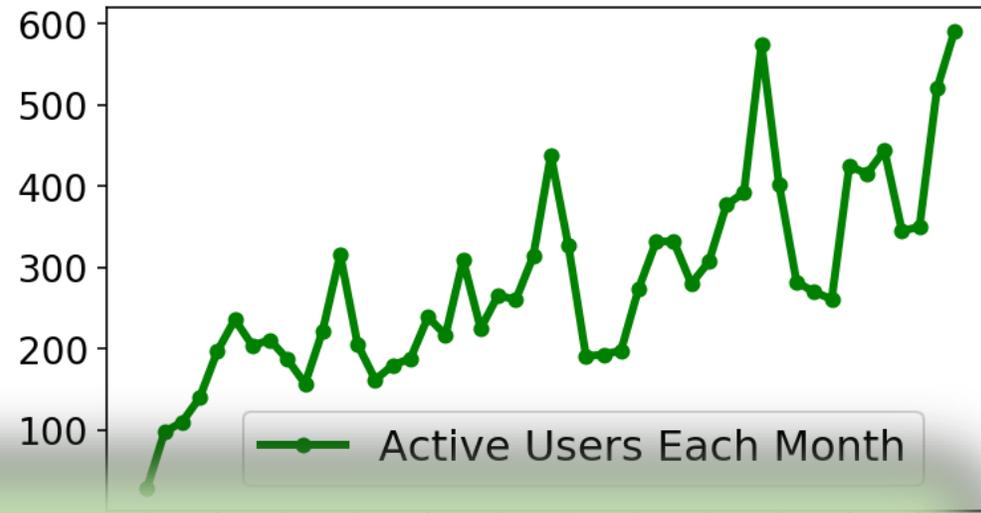
 Constraints  Errors

What can we learn from the historic data?

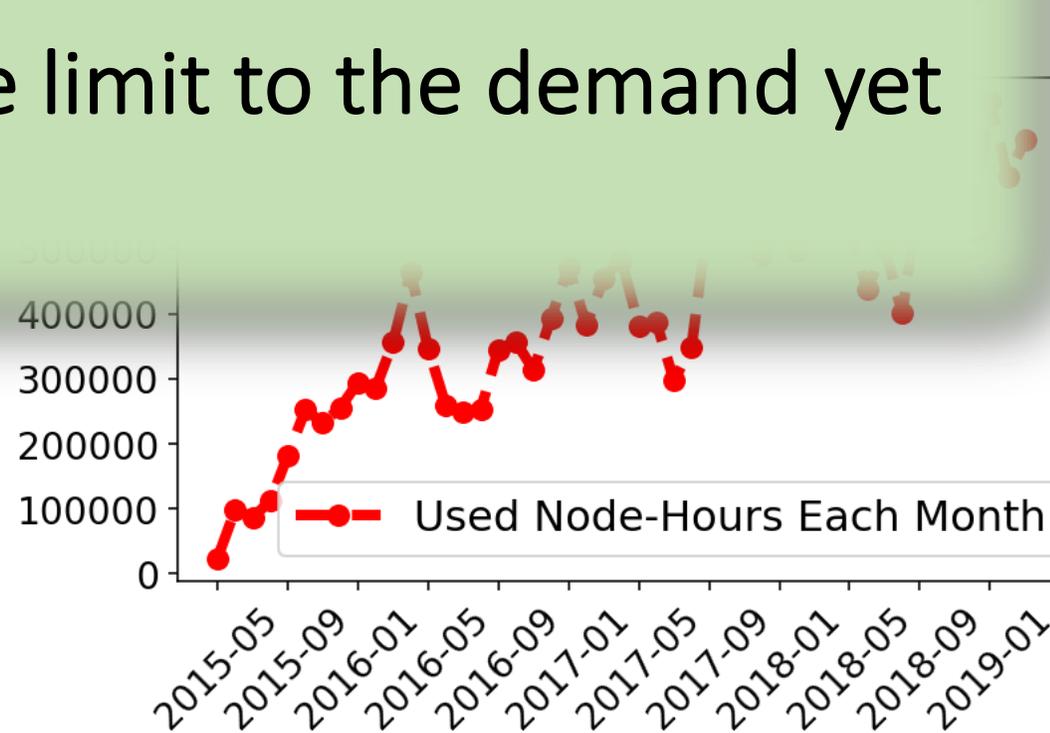
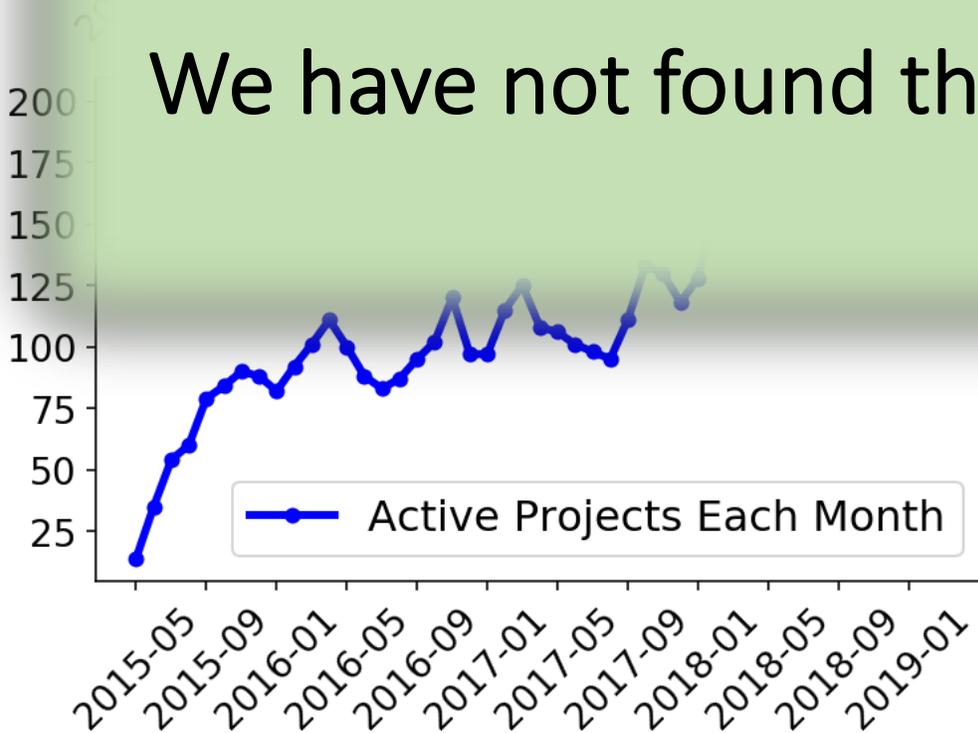
Growth and Usage



Growth and Usage



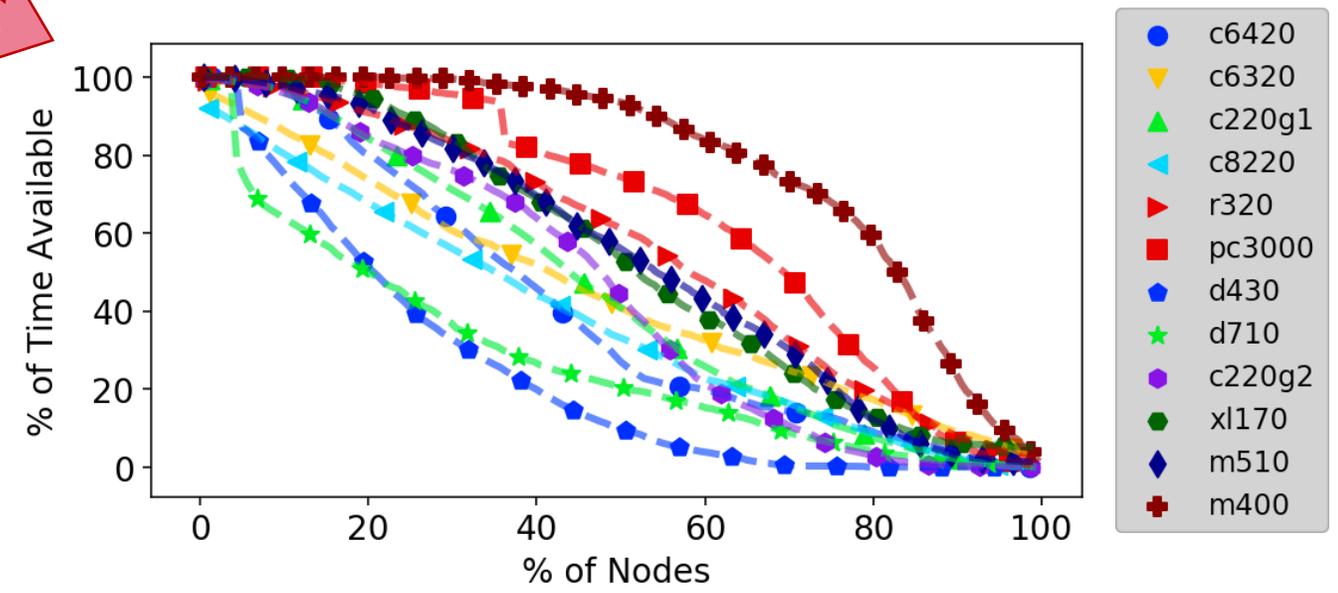
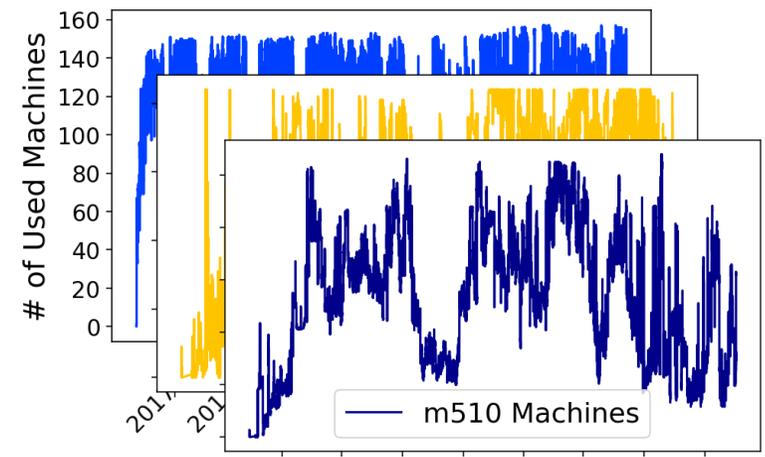
We have not found the limit to the demand yet



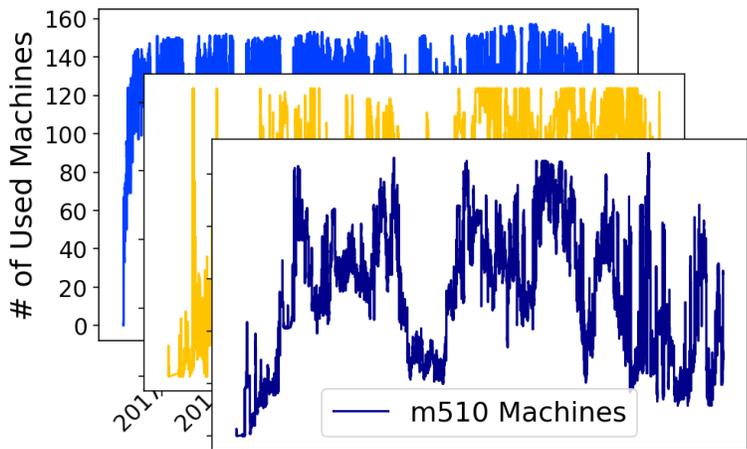
Growth
and
Usage

Is there enough hardware for everyone?

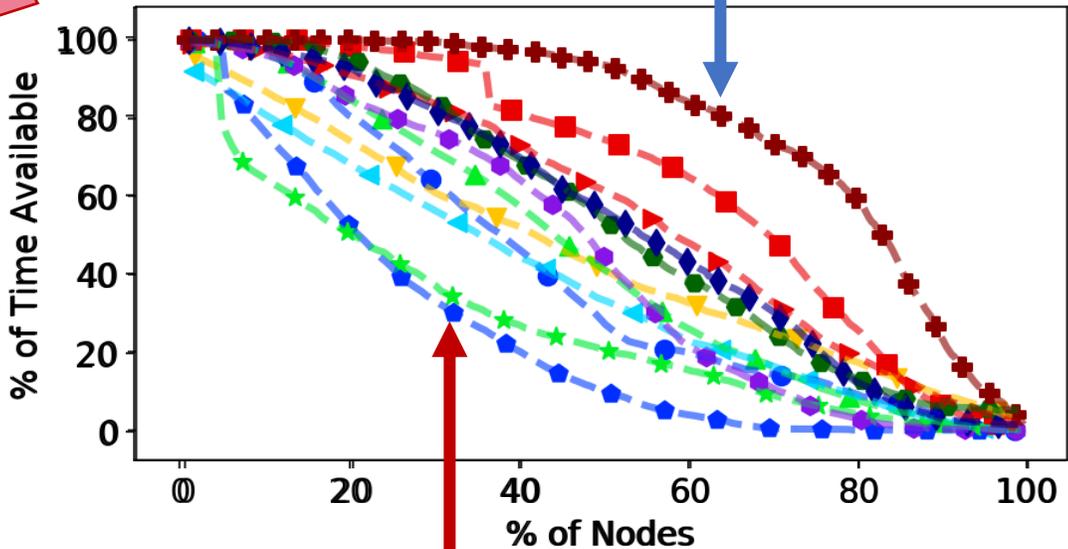
Growth and Usage



Growth and Usage



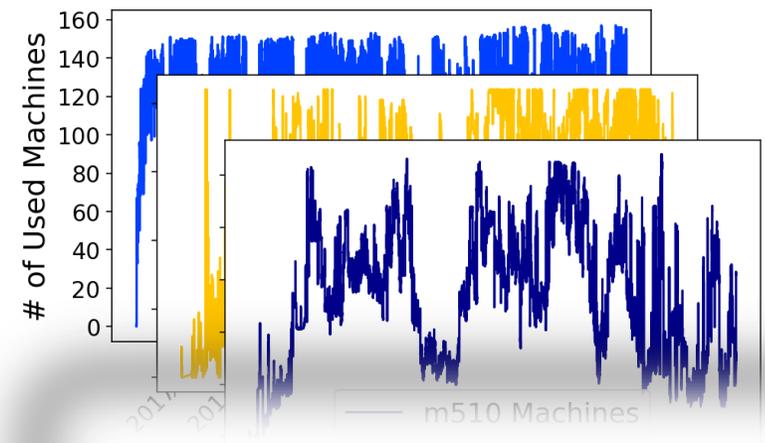
Wow, I can get **200** machines almost **80%** of the time!



I can get **50** machines around **30%** of the time 😞

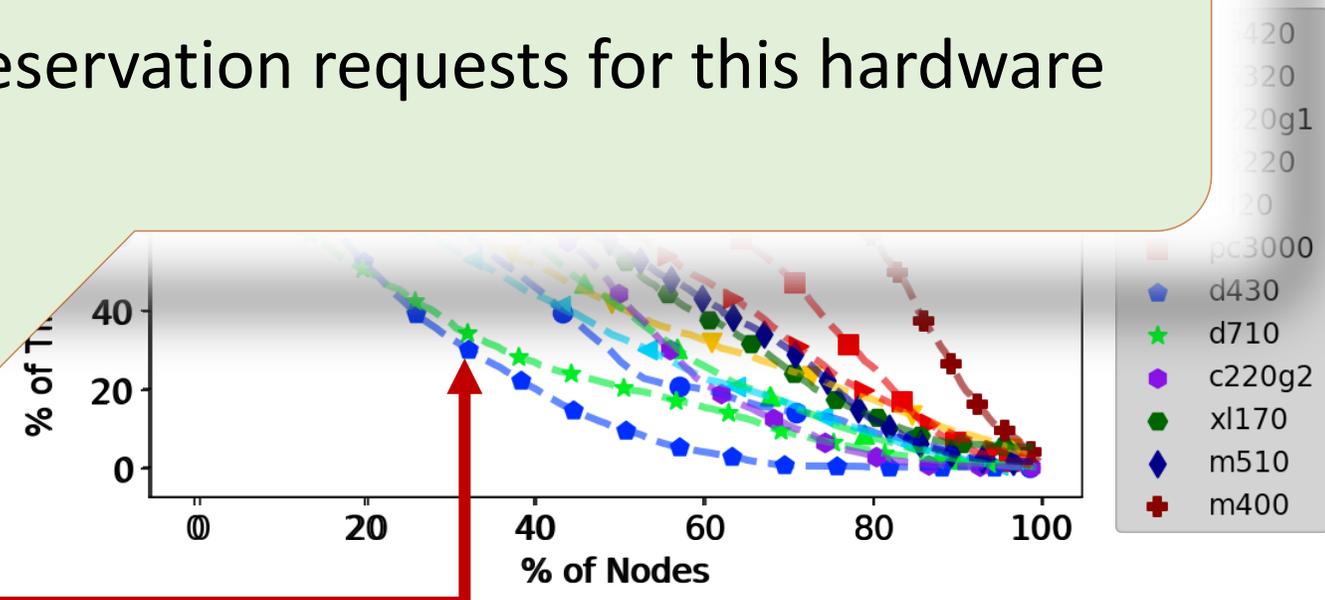


Growth and Usage

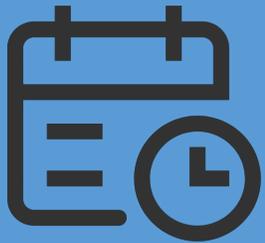


Wow, I can get **200** machines almost **80%** of the time!

I should submit reservation requests for this hardware



- pc3000
- d430
- d710
- c220g2
- xl170
- m510
- m400



Reservations

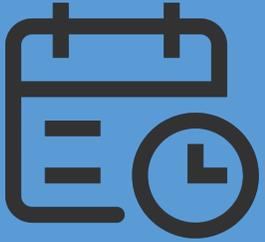
I request **100** machines
of type **d430**
between **08/01/19** and **08/15/19**



Submitted **per project** and **per hardware type**

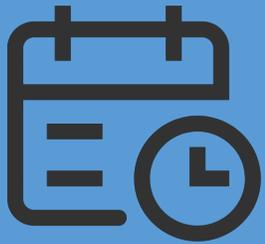
Subject to validation checks

Do not automatically launch experiments –
– only enforce availability

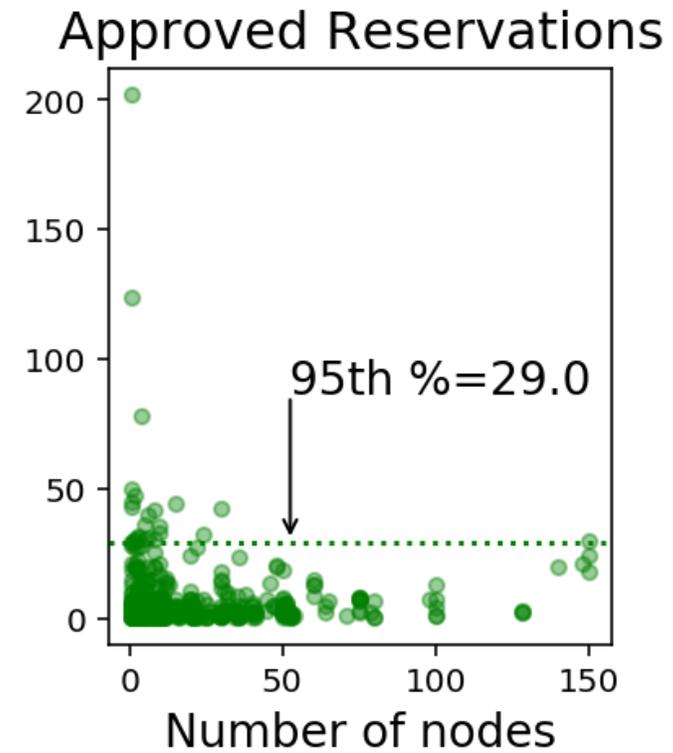
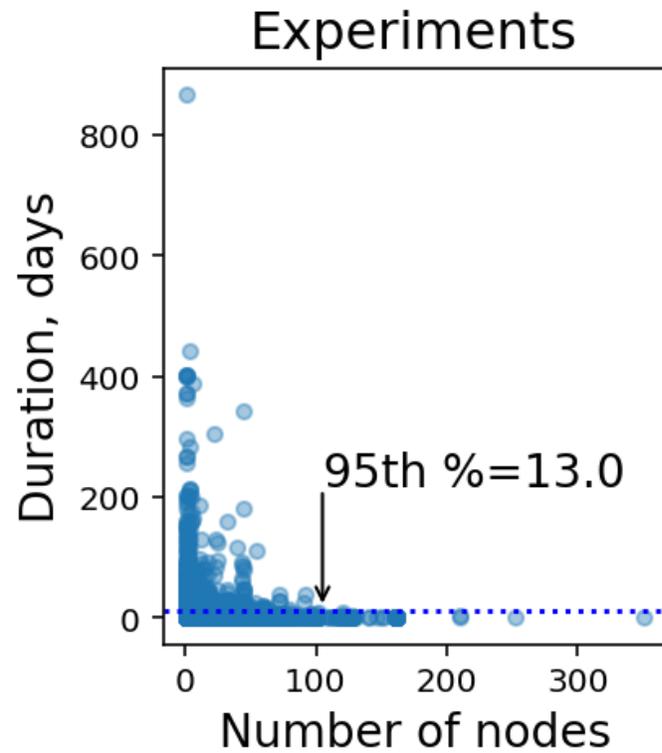
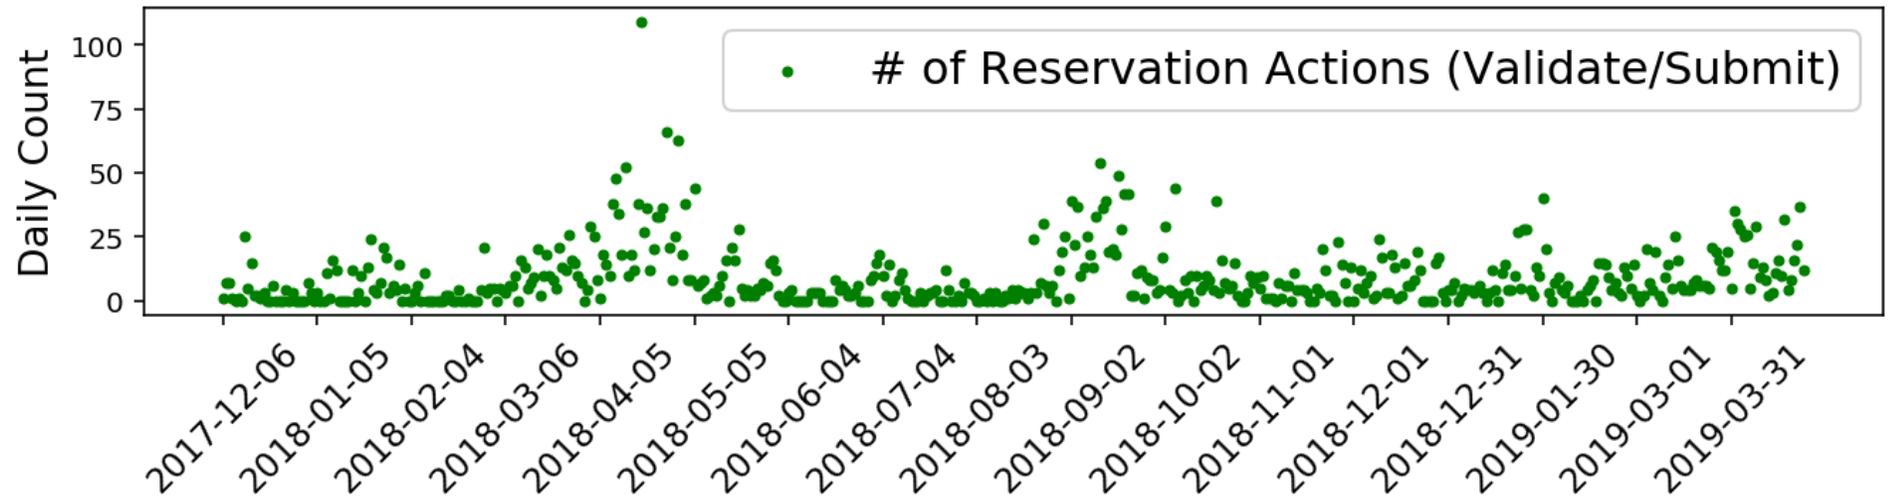


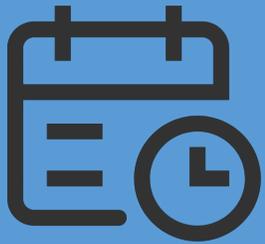
Reservations

How well do they work in practice?

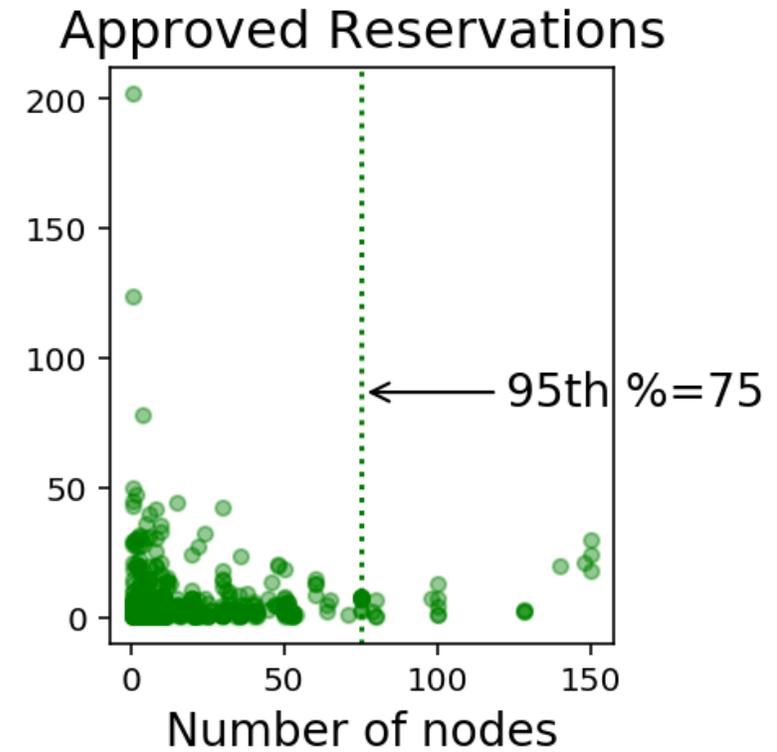
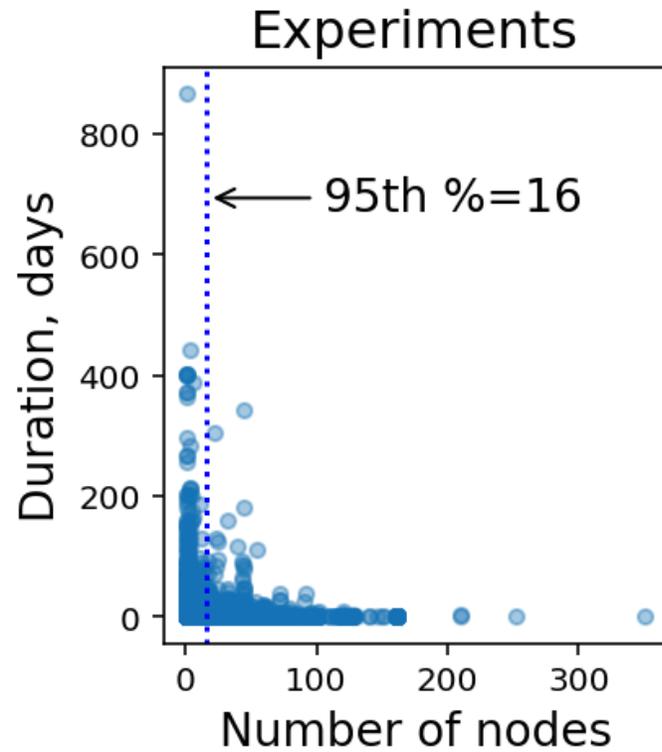
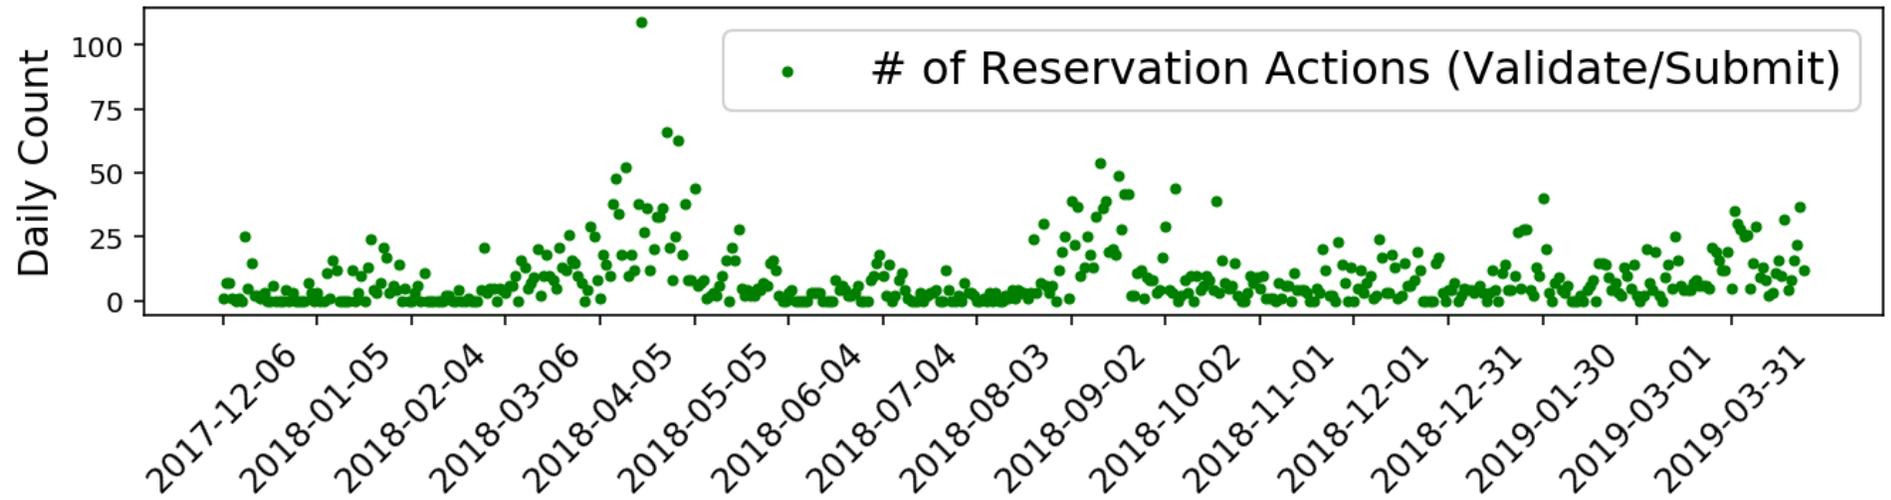


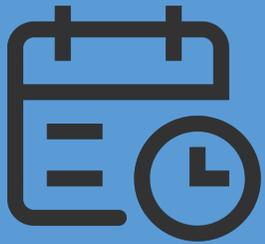
Reservations



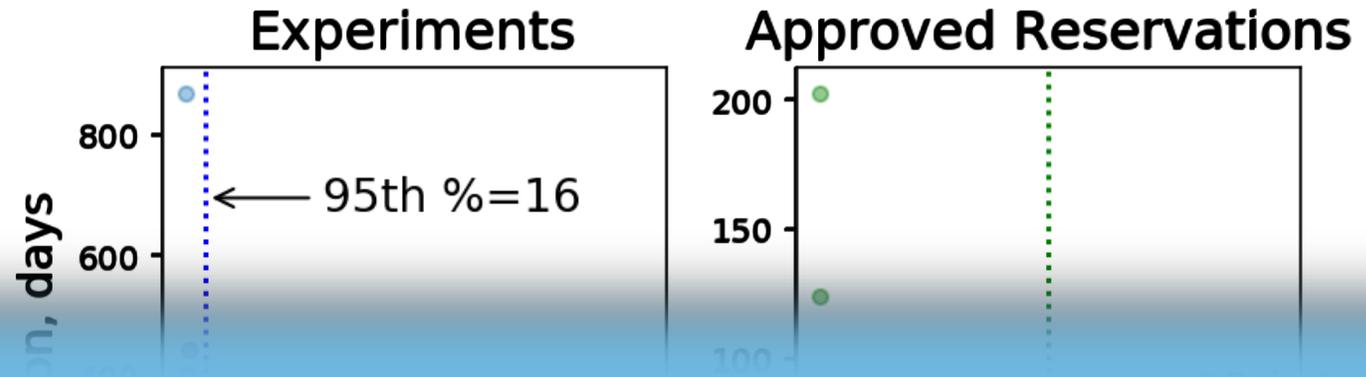
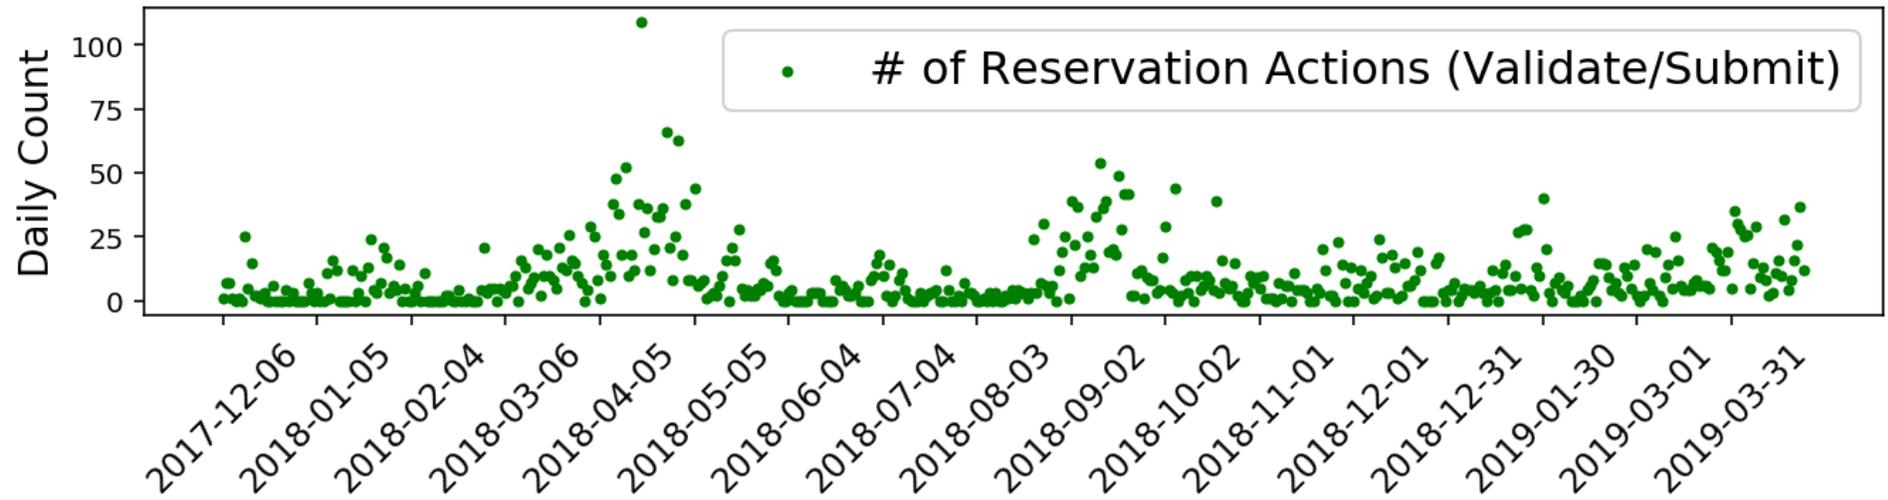


Reservations

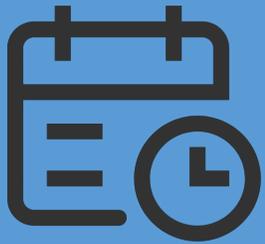




Reservations

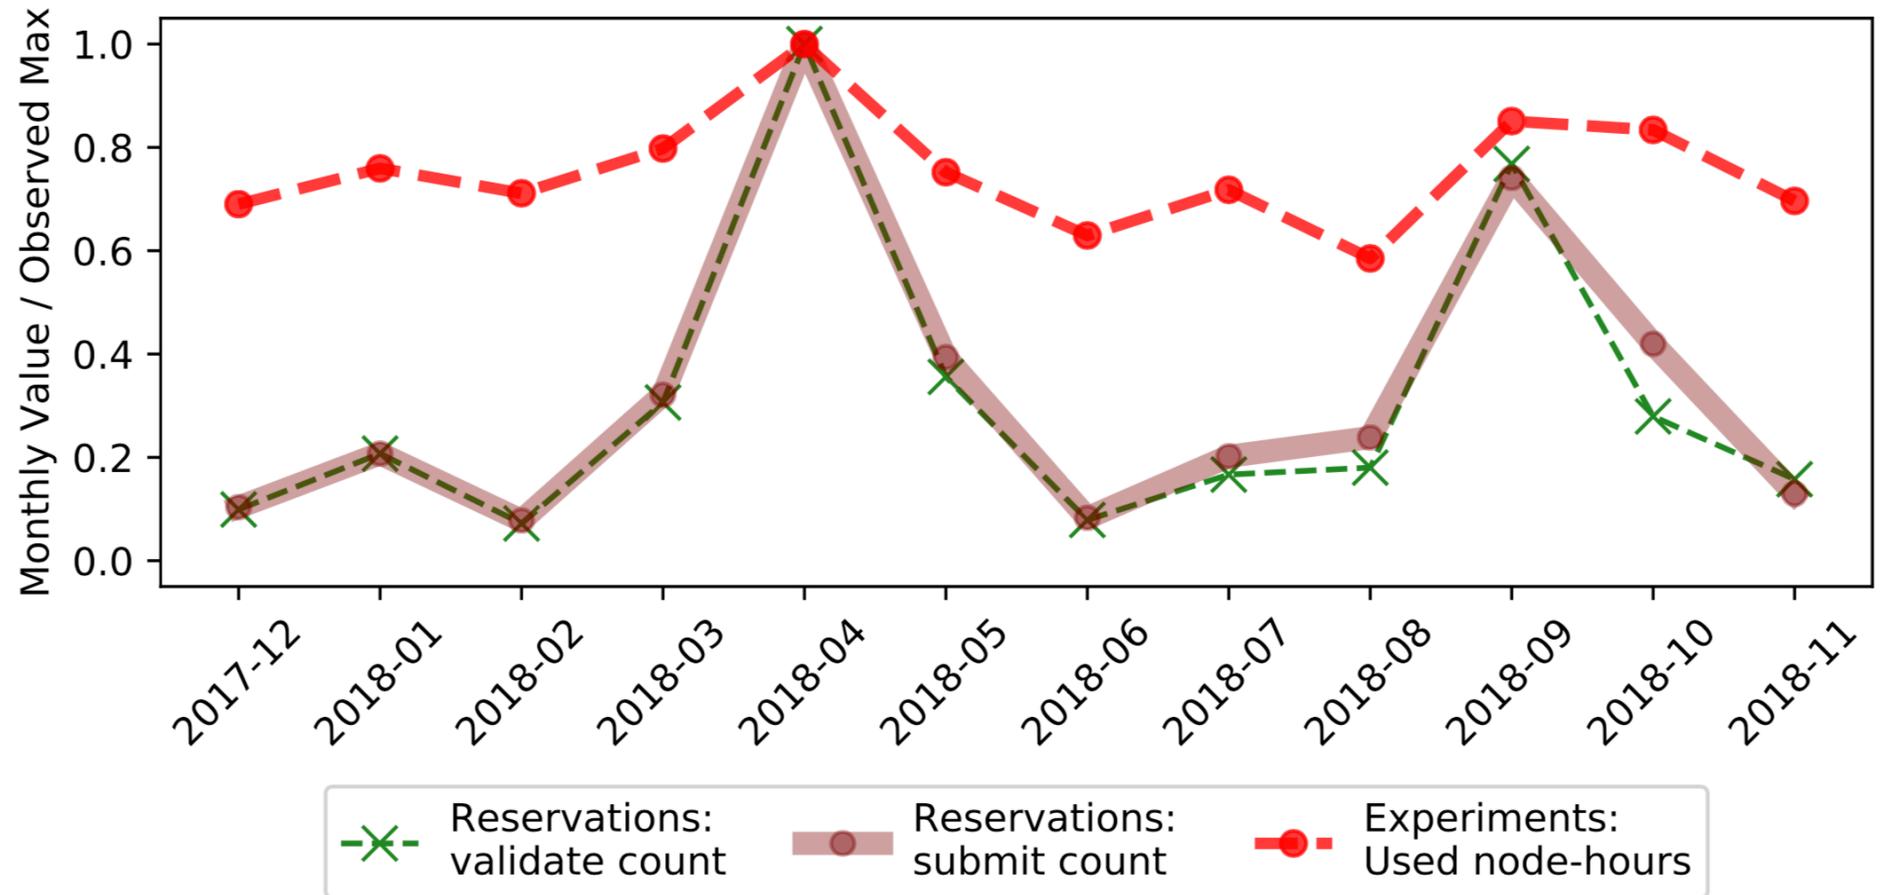


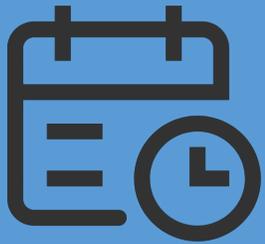
Reservations allow users to run longer and larger experiments



Reservations

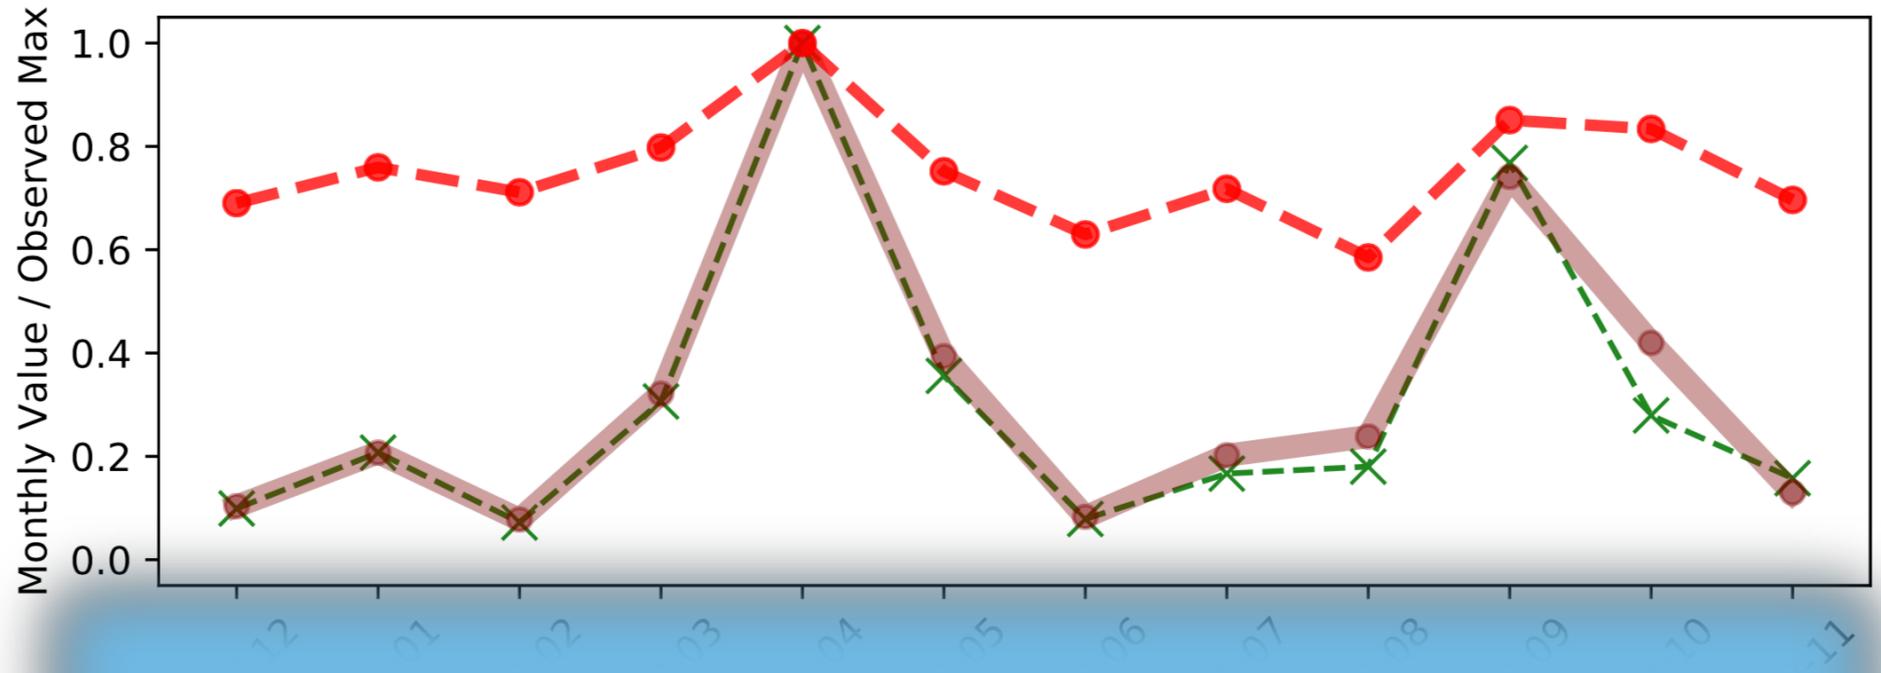
High testbed utilization ~ high use of reservations





Reservations

High testbed utilization \sim high use of reservations



Reservations allow users to meet deadlines

 Reservations Constraints

Errors

How can the testbed efficiently assign resources to users?

Approaches to resource mapping:

General algorithm:

Very few assumptions

Constraint-satisfaction problem
and optimization problem

Specialized algorithm:

Knowledge of the facility

More tailored and
actionable feedback

Resource
Mapping

Approaches to resource mapping:

General algorithm:

Very few assumptions

Constraint-satisfaction problem
and optimization problem

Specialized algorithm:

Knowledge of the facility

More tailored and
actionable feedback

CloudLab

Simulated annealing for solving graph
isomorphism problem (NP-hard)

+

Set of deterministic heuristics
as a wrapper for improved feedback

Results of Mapping Errors



Errors

Error Message

Helpful? %

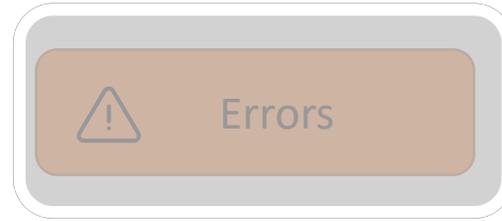
1. Resource reservation violation: X nodes of type HW requested, but only Y available	✓	27.79
2. X nodes of type HW requested, but only Y available nodes of type HW found	✓	21.86
3. No Possible Mapping for X: Too many links of type Y	✓	6.64
4. No Connection	✗	5.22
5. Insufficient Bandwidth	✗	4.88
6. No Possible Mapping for X: OS 'Y' does not run on this hardware type	✓	4.74
7. Not enough nodes because of policy restrictions or existing resource reservations	✓	4.37
8. No Possible Mapping for X: No physical nodes have feature Y	✓	3.54
9. Insufficient Nodes: Unexplained	✗	3.39
10. Fixed physical node X not available.	✓	2.56



Errors

Error	We have identified common error scenarios and addressed them using custom heuristics	%
1. Resource not found		17.79
2. Username or password incorrect		11.86
3. Network error		6.64
4. Network error		5.22
5. Invalid email address		4.88
6. Network error		4.74
7. Network error		4.37
8. Network error		3.54
9. Invalid email address		3.39
10. File not found		2.56

86.5% of last year's errors resulted in helpful error messages



How can the testbed help avoid some of these errors?



Constraints System

vs.



Error Reporting

Analogy

Feedback provided by
an IDE

Feedback provided by
a compiler



Constraints

Candidate: $x = \{\text{utah}, \text{m400}, \text{pc}, \text{ubuntu16-64-ARM}\}$

Groups – whitelists of acceptable combinations with
 ≥ 2 properties

Evaluation – Boolean *product of sums*

Group relating site, hardware, and type:	Group relating hardware and image:
$a_1(x) = \{\text{utah}, \text{m510}, \text{xen}\} \subseteq x$	$b_1(x) = \{\text{m400}, \text{ubuntu16-64-ARM}\} \subseteq x$
$a_2(x) = \{\text{utah}, \text{m400}, \text{pc}\} \subseteq x$	$b_2(x) = \{\text{m510}, \text{ubuntu16-64-STD}\} \subseteq x$
...	...
$a_n(x) = \{\text{wisconsin}, \text{c220g2}, \text{pc}\} \subseteq x$	$b_m(x) = \{\text{c220g2}, \text{fbbsd110-64-STD}\} \subseteq x$
$A(x) = a_1(x) \vee a_2(x) \vee \dots \vee a_n(x)$	$B(x) = b_1(x) \vee b_2(x) \vee \dots \vee b_m(x)$

$a_2(x) \wedge b_1(x) = 1 \rightarrow$ candidate passes the check

Candidate: `x={utah m400 nc ubuntu16-64-ARM}`

Used in Two Contexts

Interactive Topology Design:

Early feedback/warnings

More permissive

Cluster Selection:

Block instantiation if request is infeasible

Disable selection of incompatible clusters

More conservative

$a_2(x) \wedge b_1(x) = 1 \rightarrow$ candidate passes the check



Constraints

Candidate: `x={utah m400 nc ubuntu16-64-ARM}`



Constraints

Constraints checker running as a **lightweight** system in front of the **complex** mapper improves user experience

$a_2(x) \wedge b_1(x) = 1 \rightarrow$ candidate passes the check

In Conclusion



Reservation System

Reservations allow users to run **longer** and **larger** experiments and meet **deadlines**



Constraints System

Constraints checker running as a **lightweight** system in front of the mapper improves user experience



Error Reporting

The identified common and addressed error scenarios account for **86.5%** of errors; they lead to helpful messages

In Conclusion

CloudLab



Reservations



Constraints



Errors

Satisfies diverse researchers' needs, helps them select feasible configurations, and provides helpful feedback

Data & Code

Activity of 4K users
in 79K experiments
on over 2K servers
for over 4 years
and complete record of testbed
events, states, and errors

<https://gitlab.flux.utah.edu/emulab/cloudlab-usage>



Awards: 1419199 and 1743363



More about CloudLab

Sign up: <https://cloudlab.us>

Join the BoF later today: 7:30-8:30pm
in Seattle Room

Thank you!

Dmitry Duplyakin

dmdu@cs.utah.edu