

# High Density Visualizations

Chuck Boeheim

SLAC

# Clusters and Scaling

A cluster is a very large error amplifier

You pay for capacity and bandwidth separately

Cheap hardware isn't

GUIs don't scale

# Typical Display Problems

Diffuse information


Few variables displayed at once

Context and detail mutually exclusive

Excessive white space and decoration

# LSF Status Display

unavail	ok	unavail	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
adam	bali0001	bali0002	bali0003	bali0004	bali0005	bali0006	bali0007	bali0008	bali0009	bali0010	bali0011	bali0012	bali0013	bali0014
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0015	bali0016	bali0017	bali0018	bali0019	bali0020	bali0021	bali0022	bali0023	bali0024	bali0025	bali0026	bali0027	bali0028	bali0029
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0030	bali0031	bali0032	bali0033	bali0034	bali0035	bali0036	bali0037	bali0038	bali0039	bali0040	bali0041	bali0042	bali0043	bali0044
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0045	bali0046	bali0047	bali0048	bali0049	bali0050	bali0051	bali0052	bali0053	bali0054	bali0055	bali0056	bali0057	bali0058	bali0059
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0060	bali0061	bali0062	bali0063	bali0064	bali0065	bali0066	bali0067	bali0068	bali0069	bali0070	bali0071	bali0072	bali0073	bali0074
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0075	bali0076	bali0077	bali0078	bali0079	bali0080	bali0081	bali0082	bali0083	bali0084	bali0085	bali0086	bali0087	bali0088	bali0089
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	unavail	ok
bali0090	bali0091	bali0092	bali0093	bali0094	bali0095	bali0096	bali0097	bali0098	bali0099	bali0100	bali0101	bali0102	bali0103	bali0104
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0105	bali0106	bali0107	bali0108	bali0109	bali0110	bali0111	bali0112	bali0113	bali0114	bali0115	bali0116	bali0117	bali0118	bali0119
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0120	bali0121	bali0122	bali0123	bali0124	bali0125	bali0126	bali0127	bali0128	bali0129	bali0130	bali0131	bali0132	bali0133	bali0134
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0135	bali0136	bali0137	bali0138	bali0139	bali0140	bali0141	bali0142	bali0143	bali0144	bali0145	bali0146	bali0147	bali0148	bali0149
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0150	bali0151	bali0152	bali0153	bali0154	bali0155	bali0156	bali0157	bali0158	bali0159	bali0160	bali0161	bali0162	bali0163	bali0164
ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
bali0165	bali0166	bali0167	bali0168	bali0169	bali0170	bali0171	bali0172	bali0173	bali0174	bali0175	bali0176	bali0177	bali0178	bali0179



# Batch Domain Report for Fri, 26 Oct 2007 14:54:19 -0700

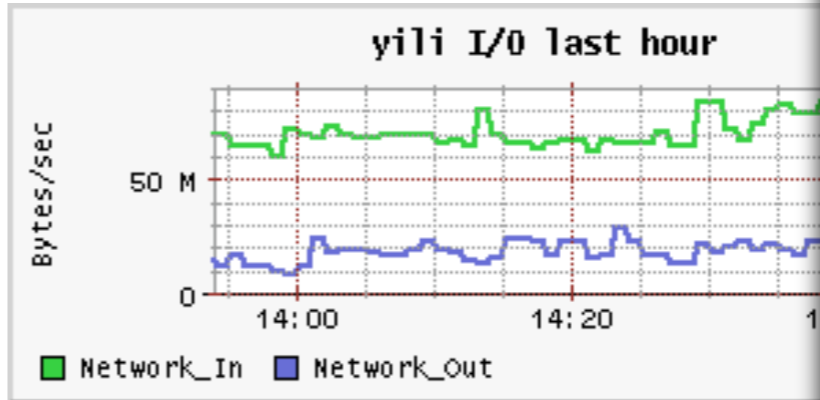

Last  Sorted

Batch Domain >

## yili (physical view)

CPU's Total: **612**  
Hosts up: **153**  
Hosts down: **0**


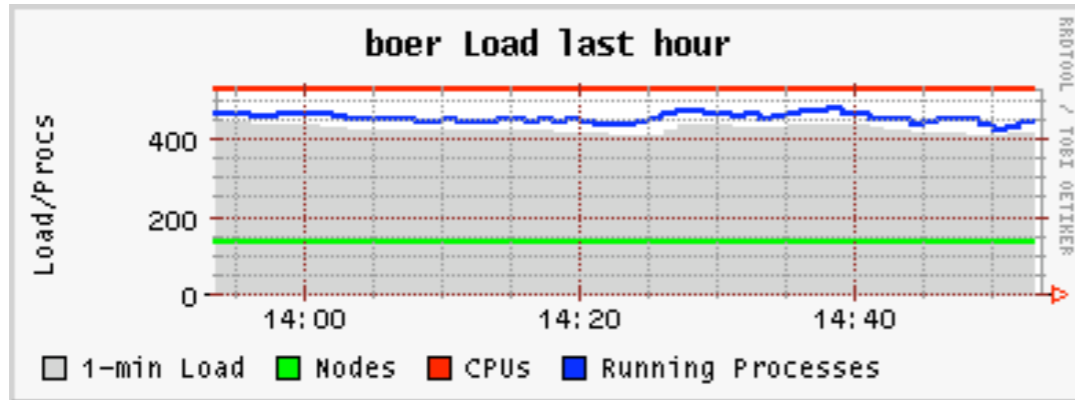
Avg Load (15, 5, 1m):  
79%, 80%, 82%  
Localtime:  
2007-10-26 14:53



## boer (physical view)



CPU's Total: **528**  
Hosts up: **132**  
Hosts down: **0**

Avg Load (15, 5, 1m):  
79%, 78%, 78%  
Localtime:  
2007-10-26 14:53



## bali (physical view)

CPU's



Batch Domain Report for Fri, 26 Oct 2007 14:54:19 -0700

Ge Batch Domain >

Snapshot of the Batch Domain Legend

Legend:

- yili
- boer
- hali
- tori
- noma
- fell
- cob
- don
- morab
- orlov

Footer: Gmetad Web Frontend version 2.5.7 Check for Updates. Gmetad Web Backend (gmetad) version 3.0.4 Check for Updates. Downloading and parsing ganglia's XML tree took 0.0009s. Images created with RRDtool.

# Good Practices

Allow the eye to find patterns

Provide the detail in context

Use all the space

Think beyond bar charts

Correlate variables

Allow comparisons

Provide different views

Allow exploration

# A Data Exploration Tool

Magnaview ([www.magnaview.nl](http://www.magnaview.nl))

Commercialization of Sequoiaview from  
University of Eindhoven

Can visualize 1 million observations

Support has been excellent

Also evaluated freeware packages, did not find  
one with as many features


# Display types in Magnaview

- Treemaps
- Colormaps
- Pixelmaps
- Grids
- Slices
- Bar charts
- Pie charts
- Line charts

Each type can contain any other type

Each project can have multiple views of data



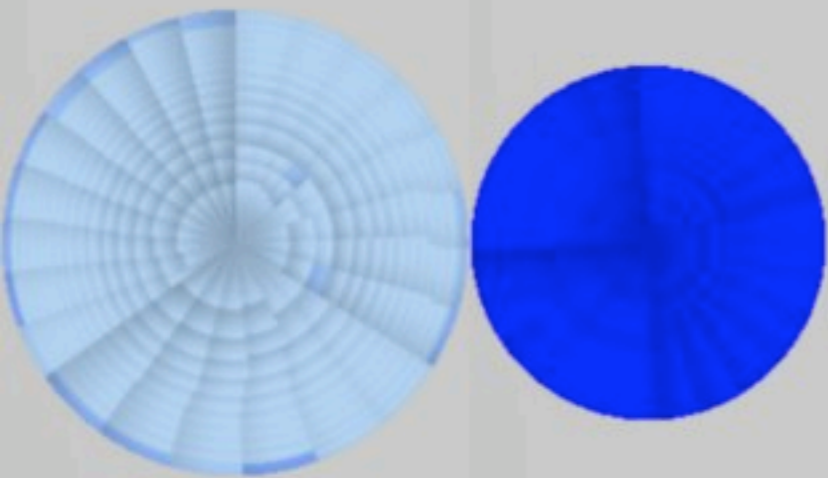


Status of batch system  
 Size of node  
 Size of cluster  
 Job distribution  
 Efficiency of job  
 Status of node

jobs=306

jobs=254

jobs=2036



jobs=77

simuq

kipac



jo...

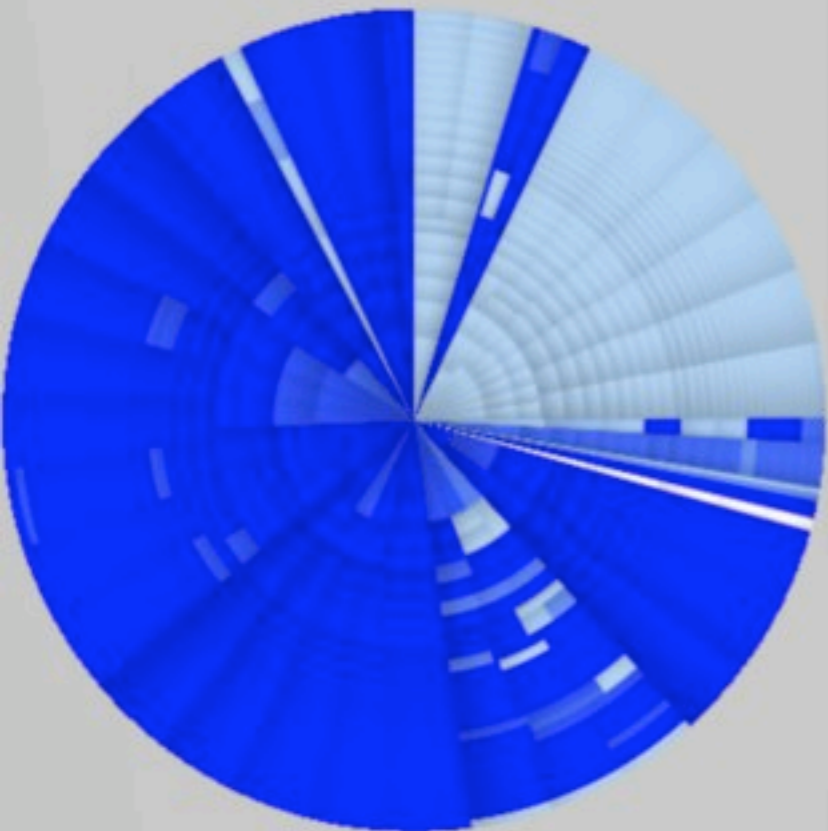
jobs=...

minitestq

s...

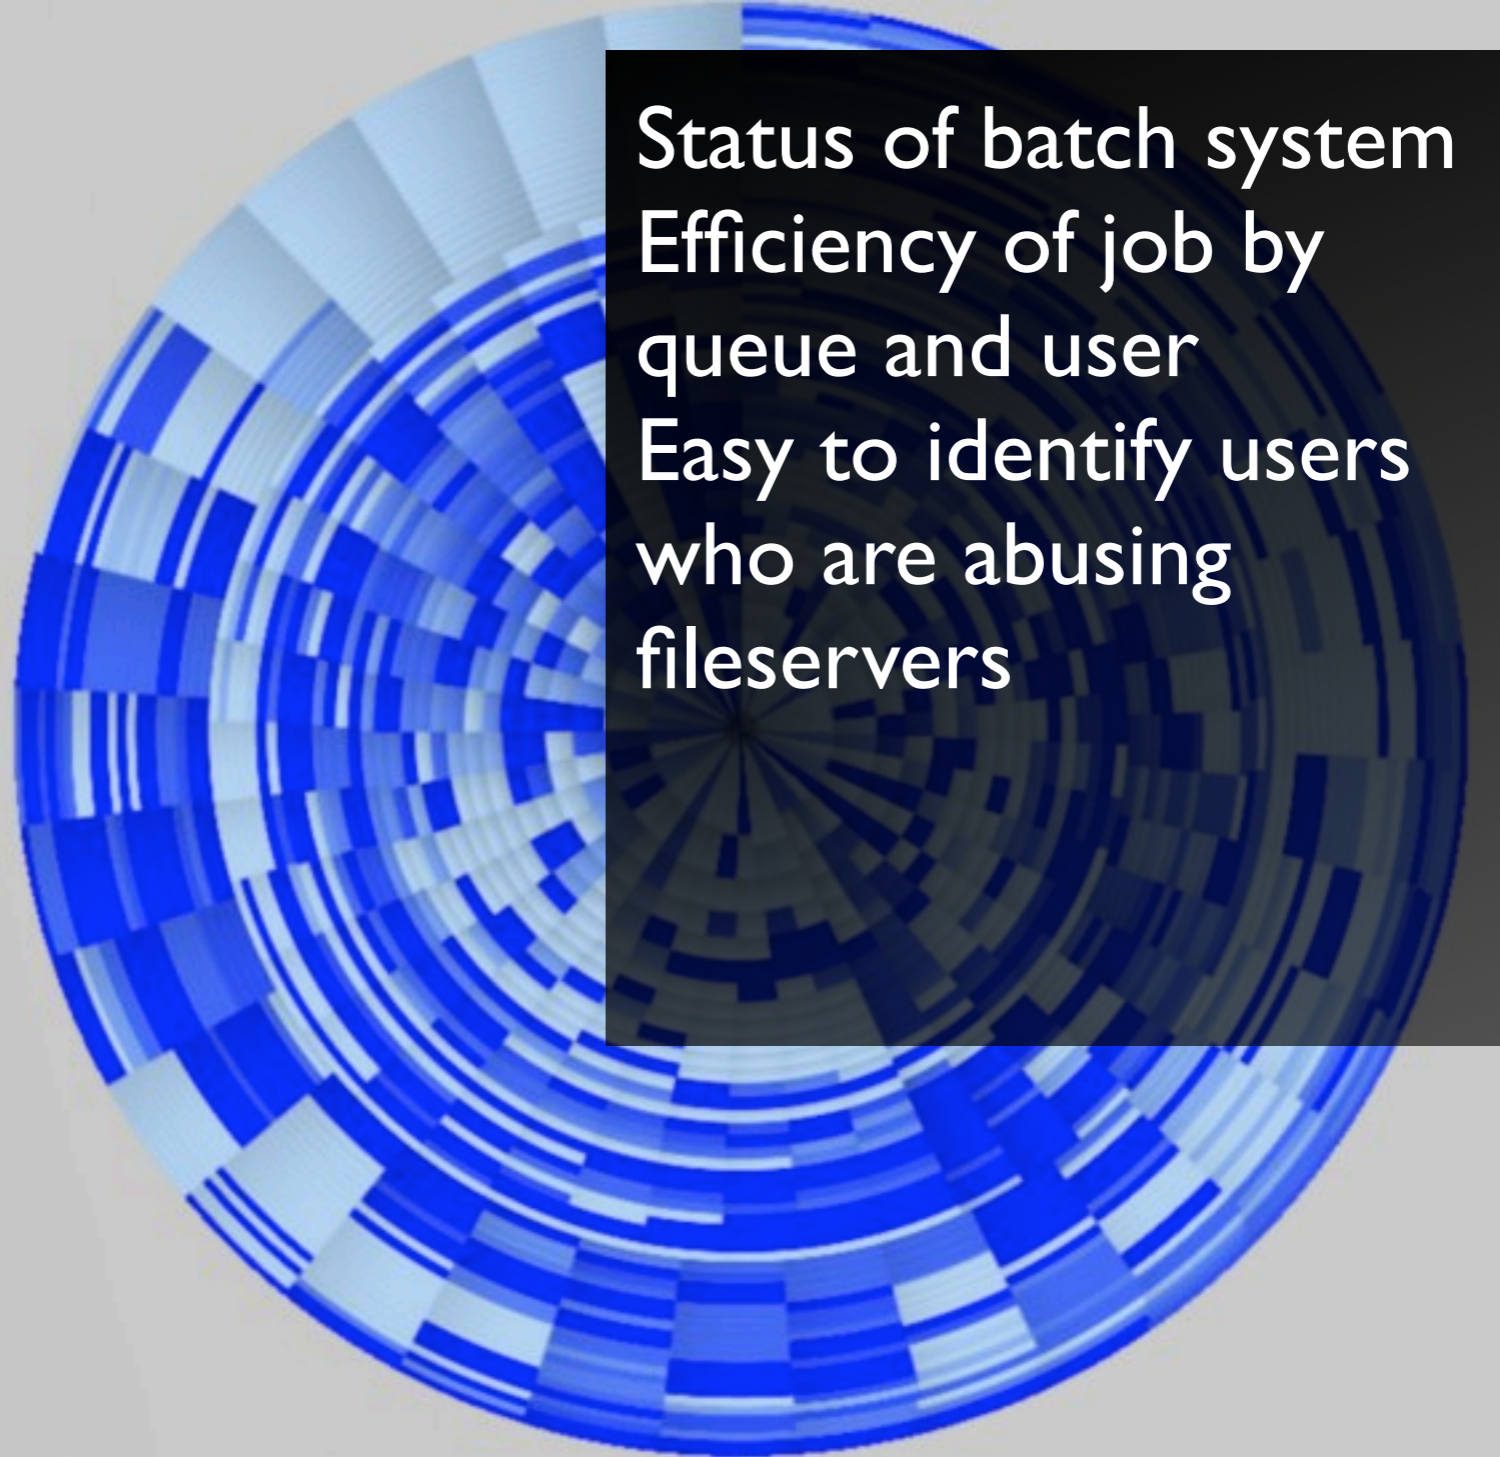
mp...

jobs=1803



jobs=1918

general



Status of batch system  
 Efficiency of job by  
 queue and user  
 Easy to identify users  
 who are abusing  
 fileservers

miniq

kanga



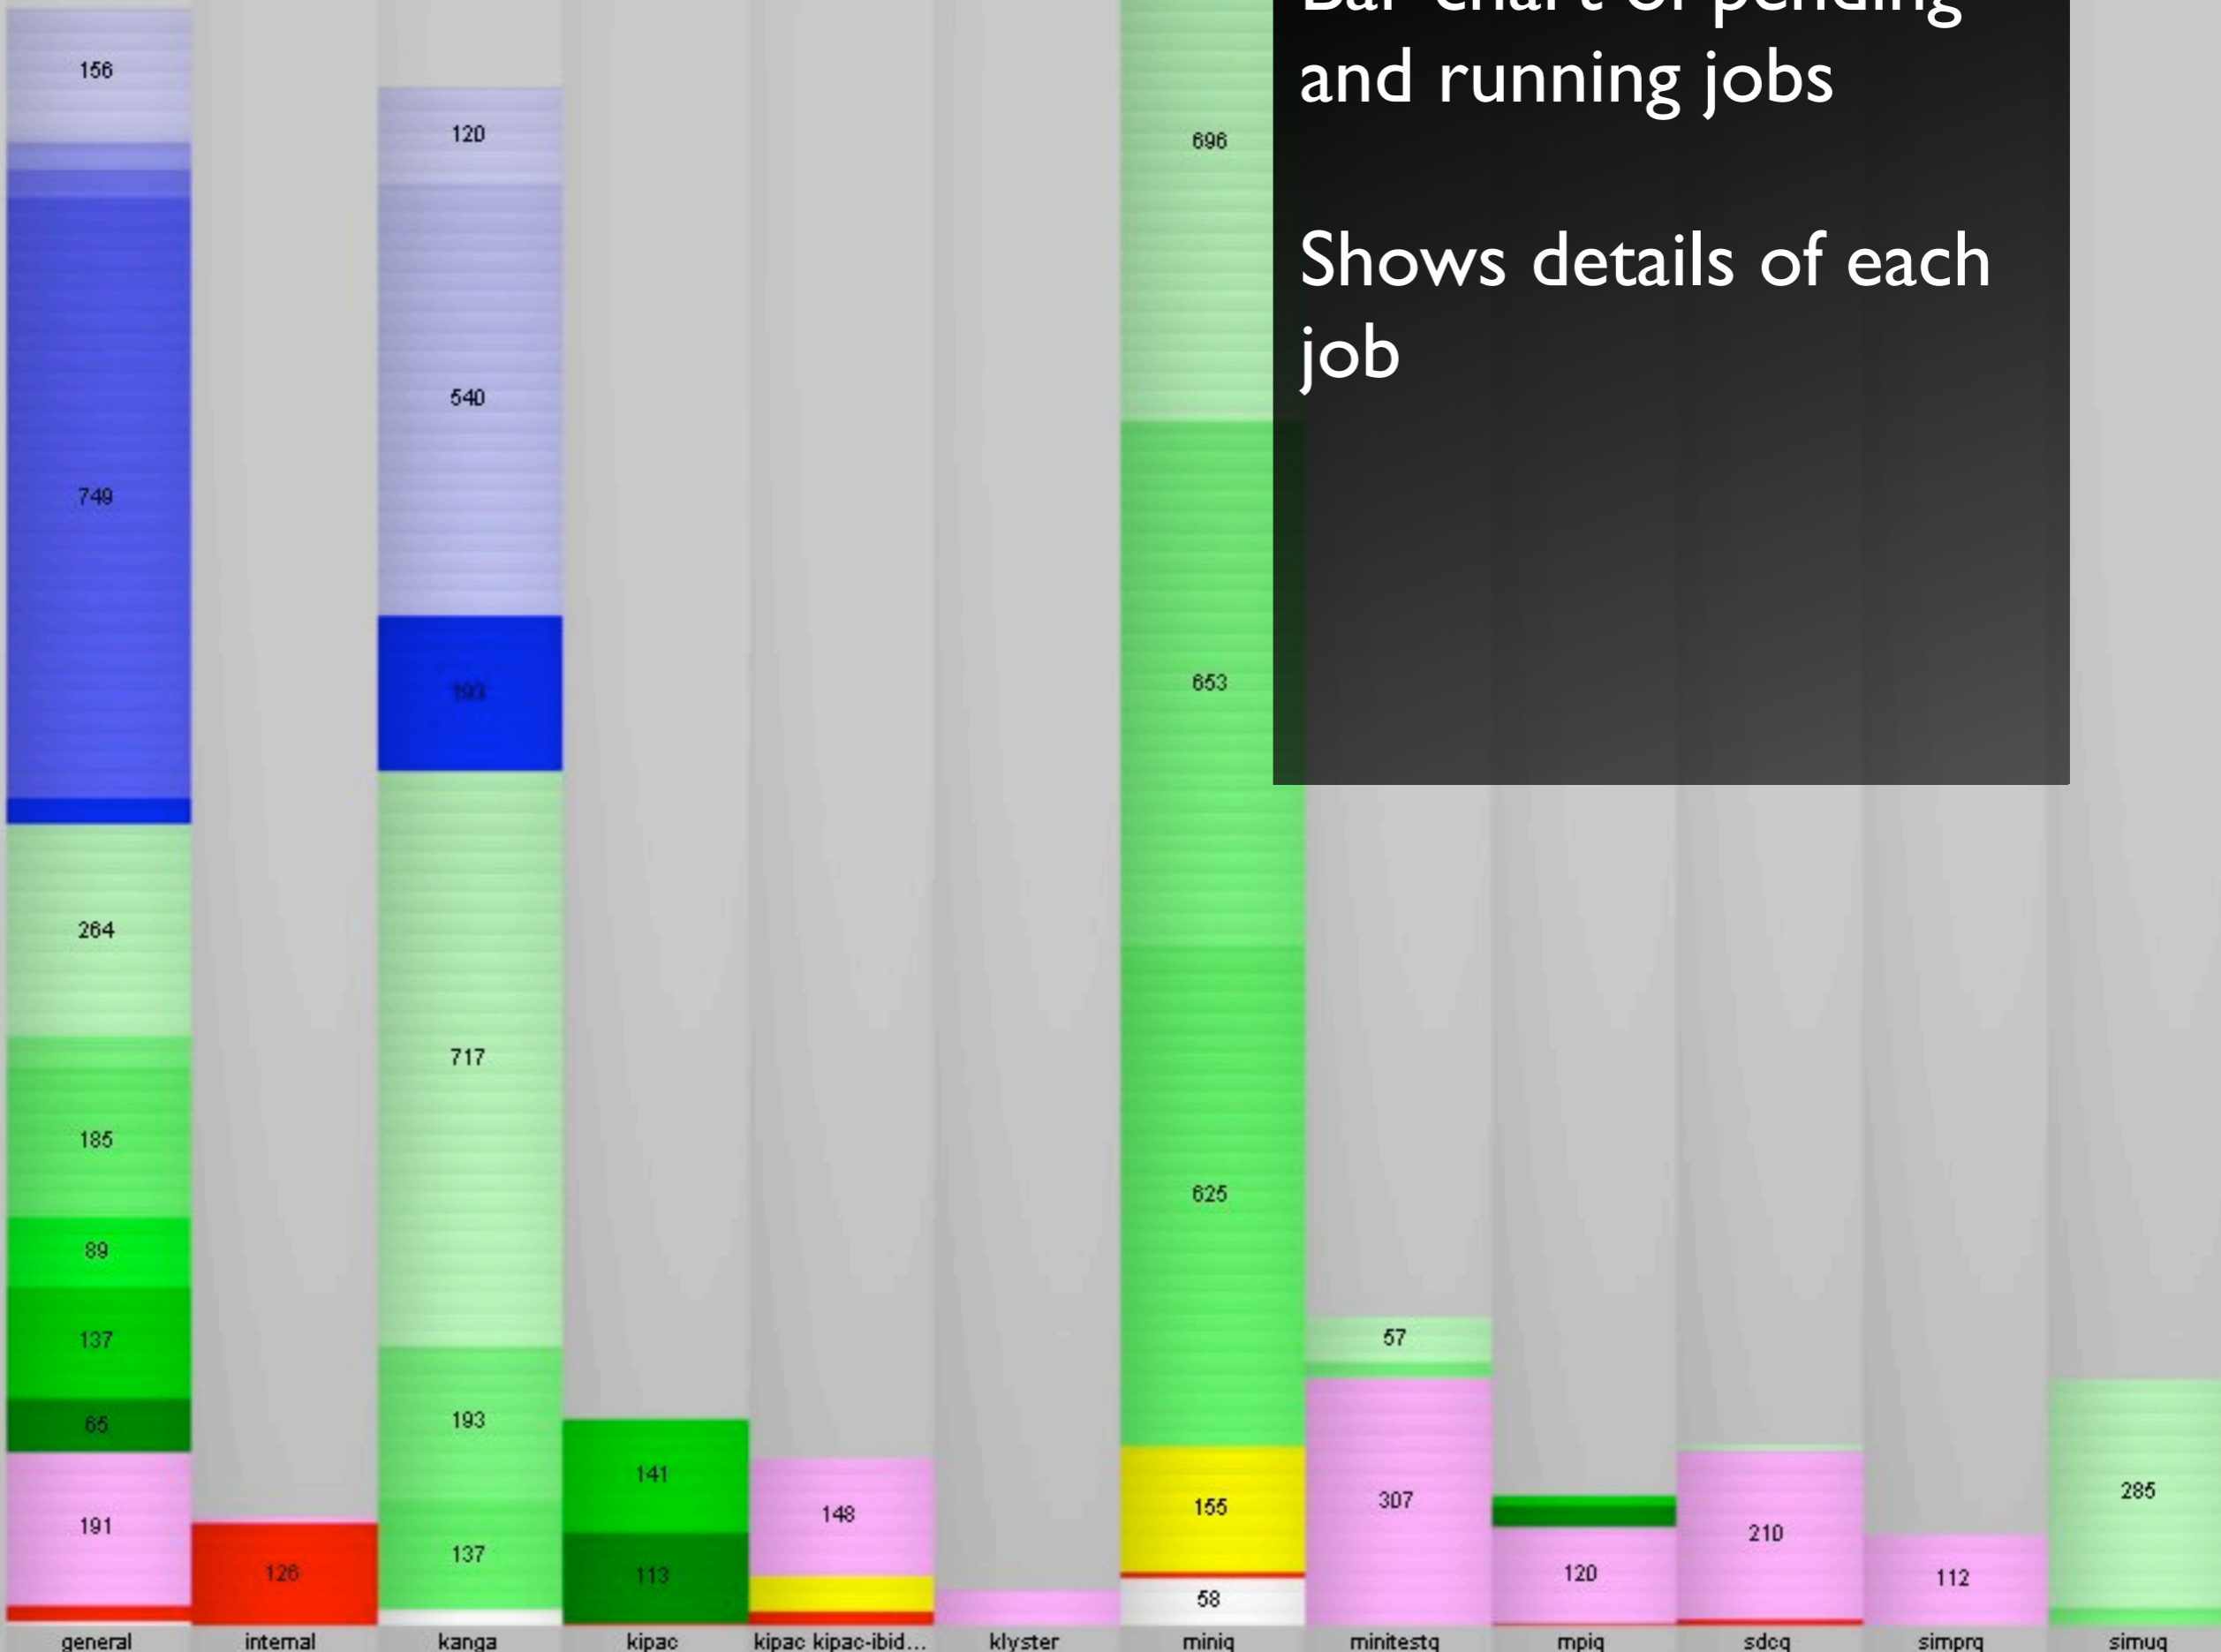
Tree map of users

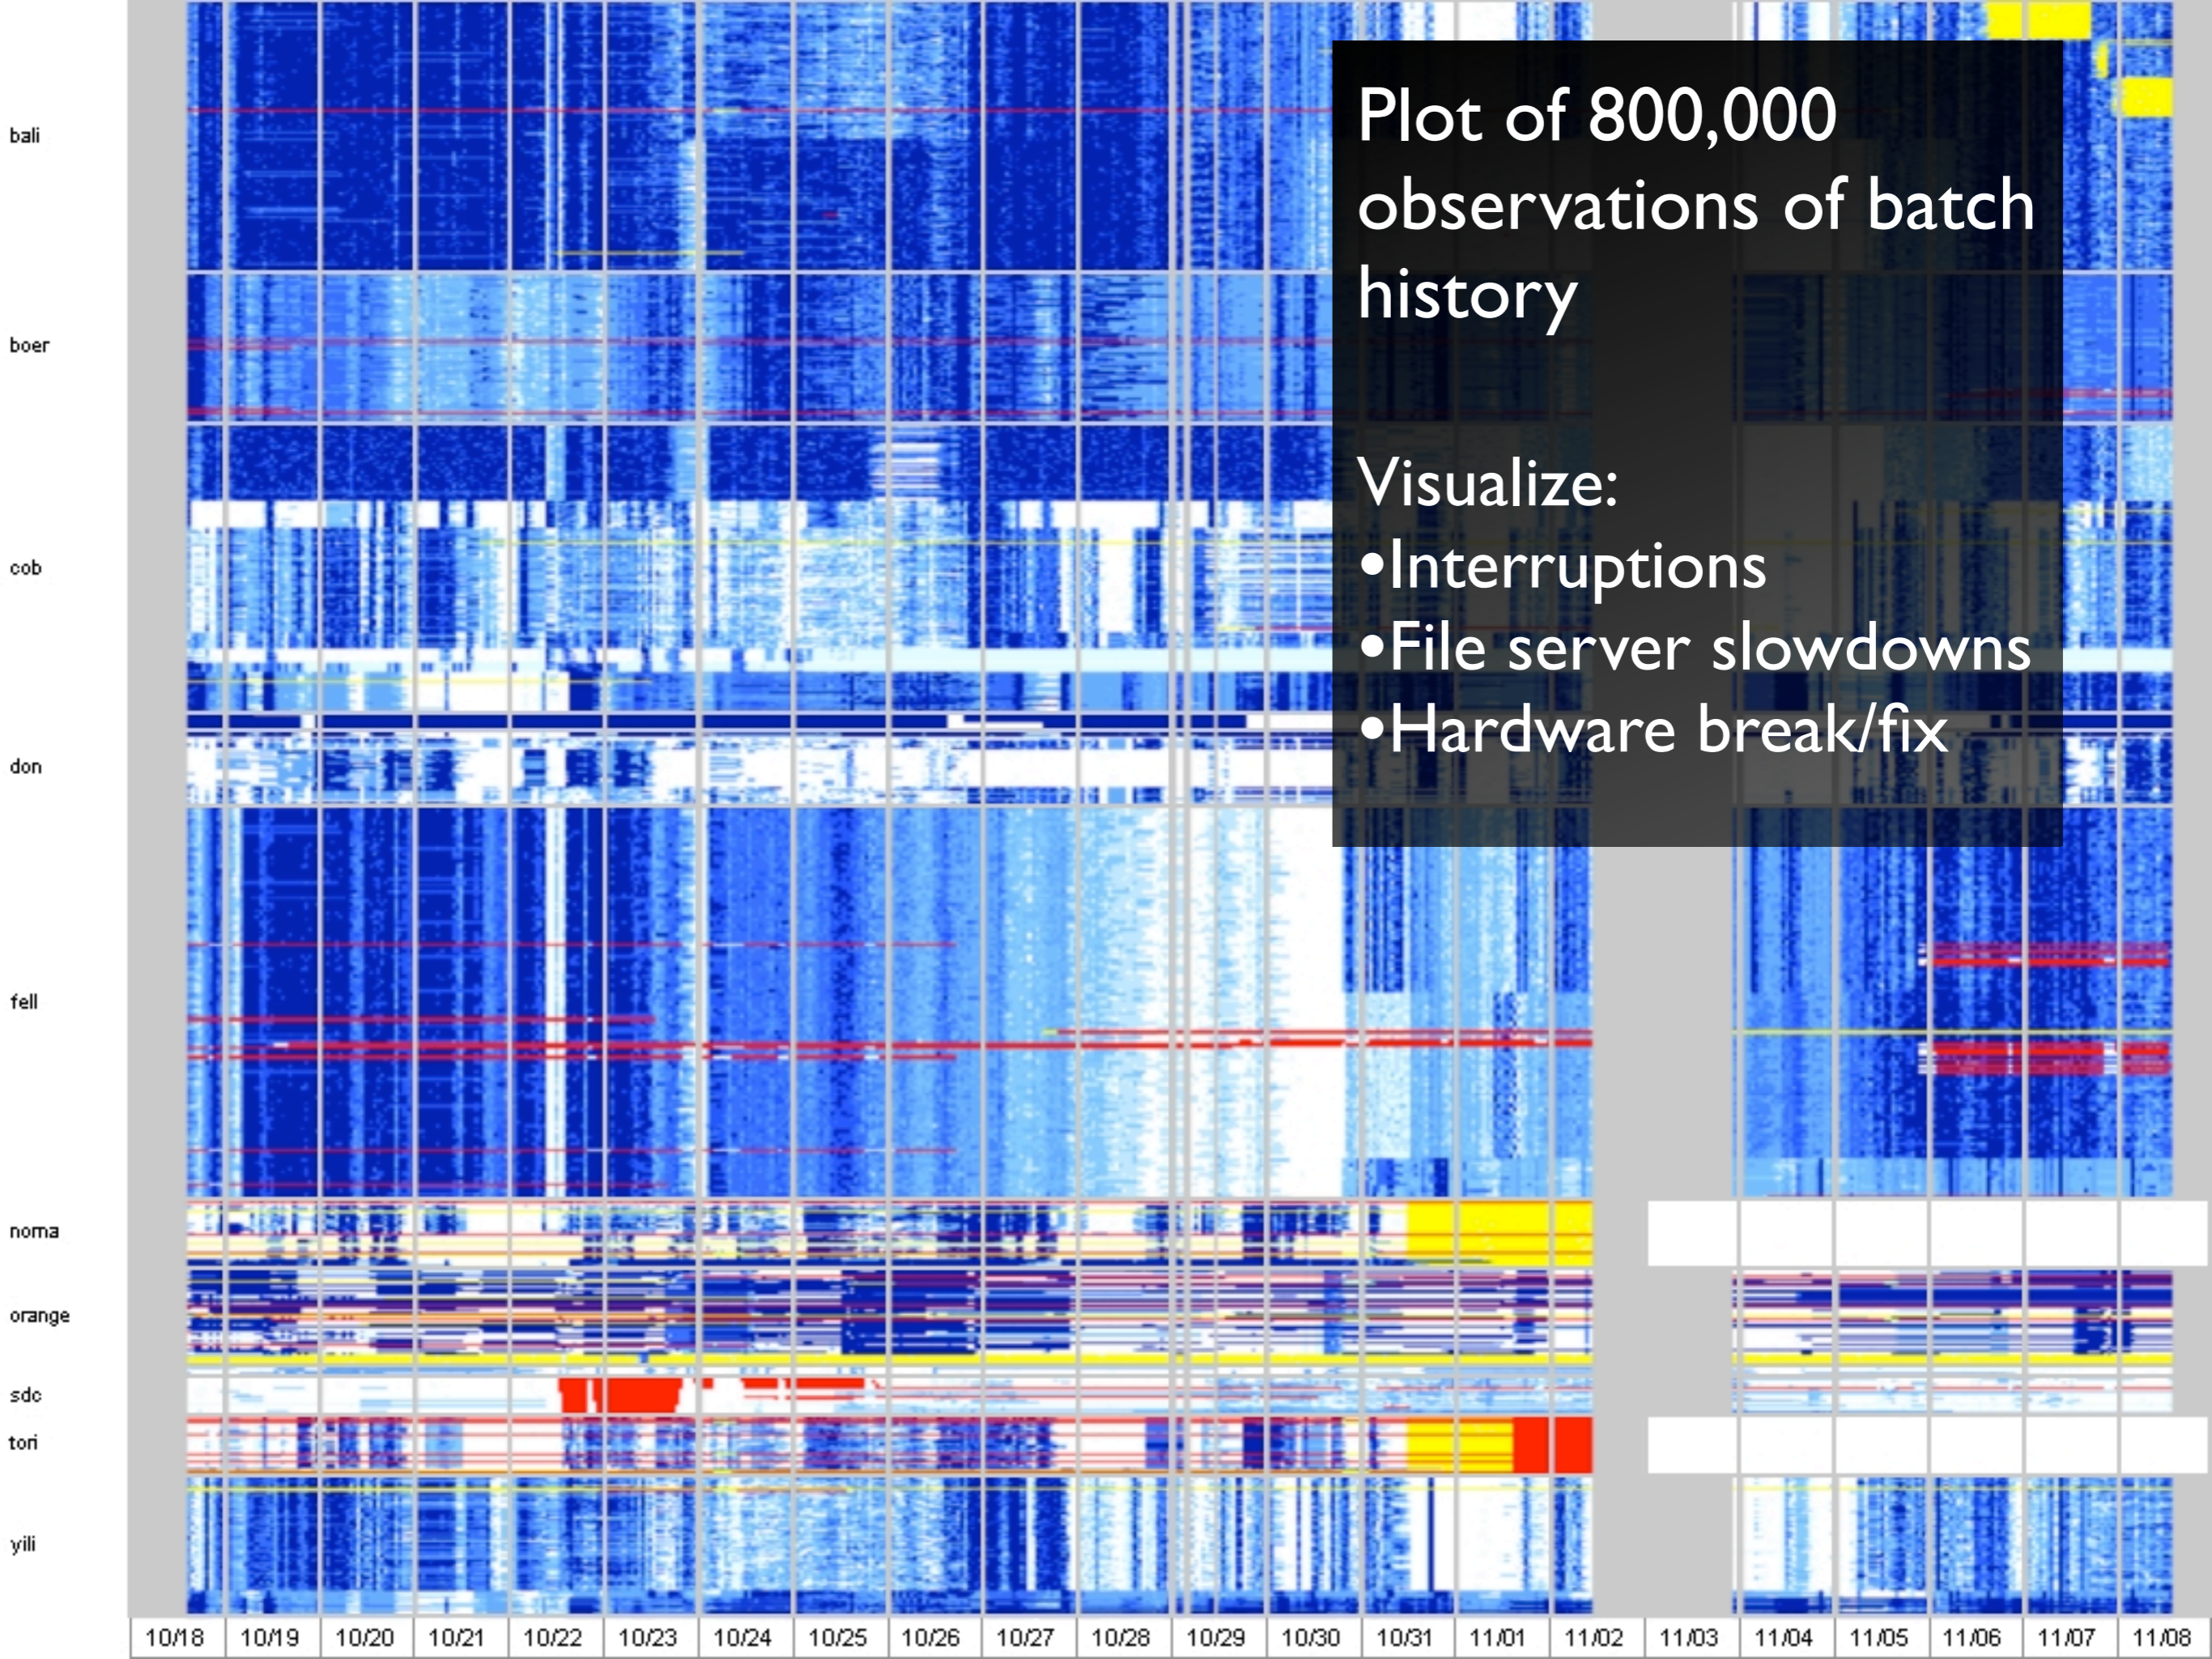
Identifies large users

Correlates queue  
time / run time with  
user

# Bar chart of pending and running jobs

Shows details of each job





# Plot of 800,000 observations of batch history

Visualize:

- Interruptions
- File server slowdowns
- Hardware break/fix


10/18 10/19 10/20 10/21 10/22 10/23 10/24 10/25 10/26 10/27 10/28 10/29 10/30 10/31 11/01 11/02 11/03 11/04 11/05 11/06 11/07 11/08

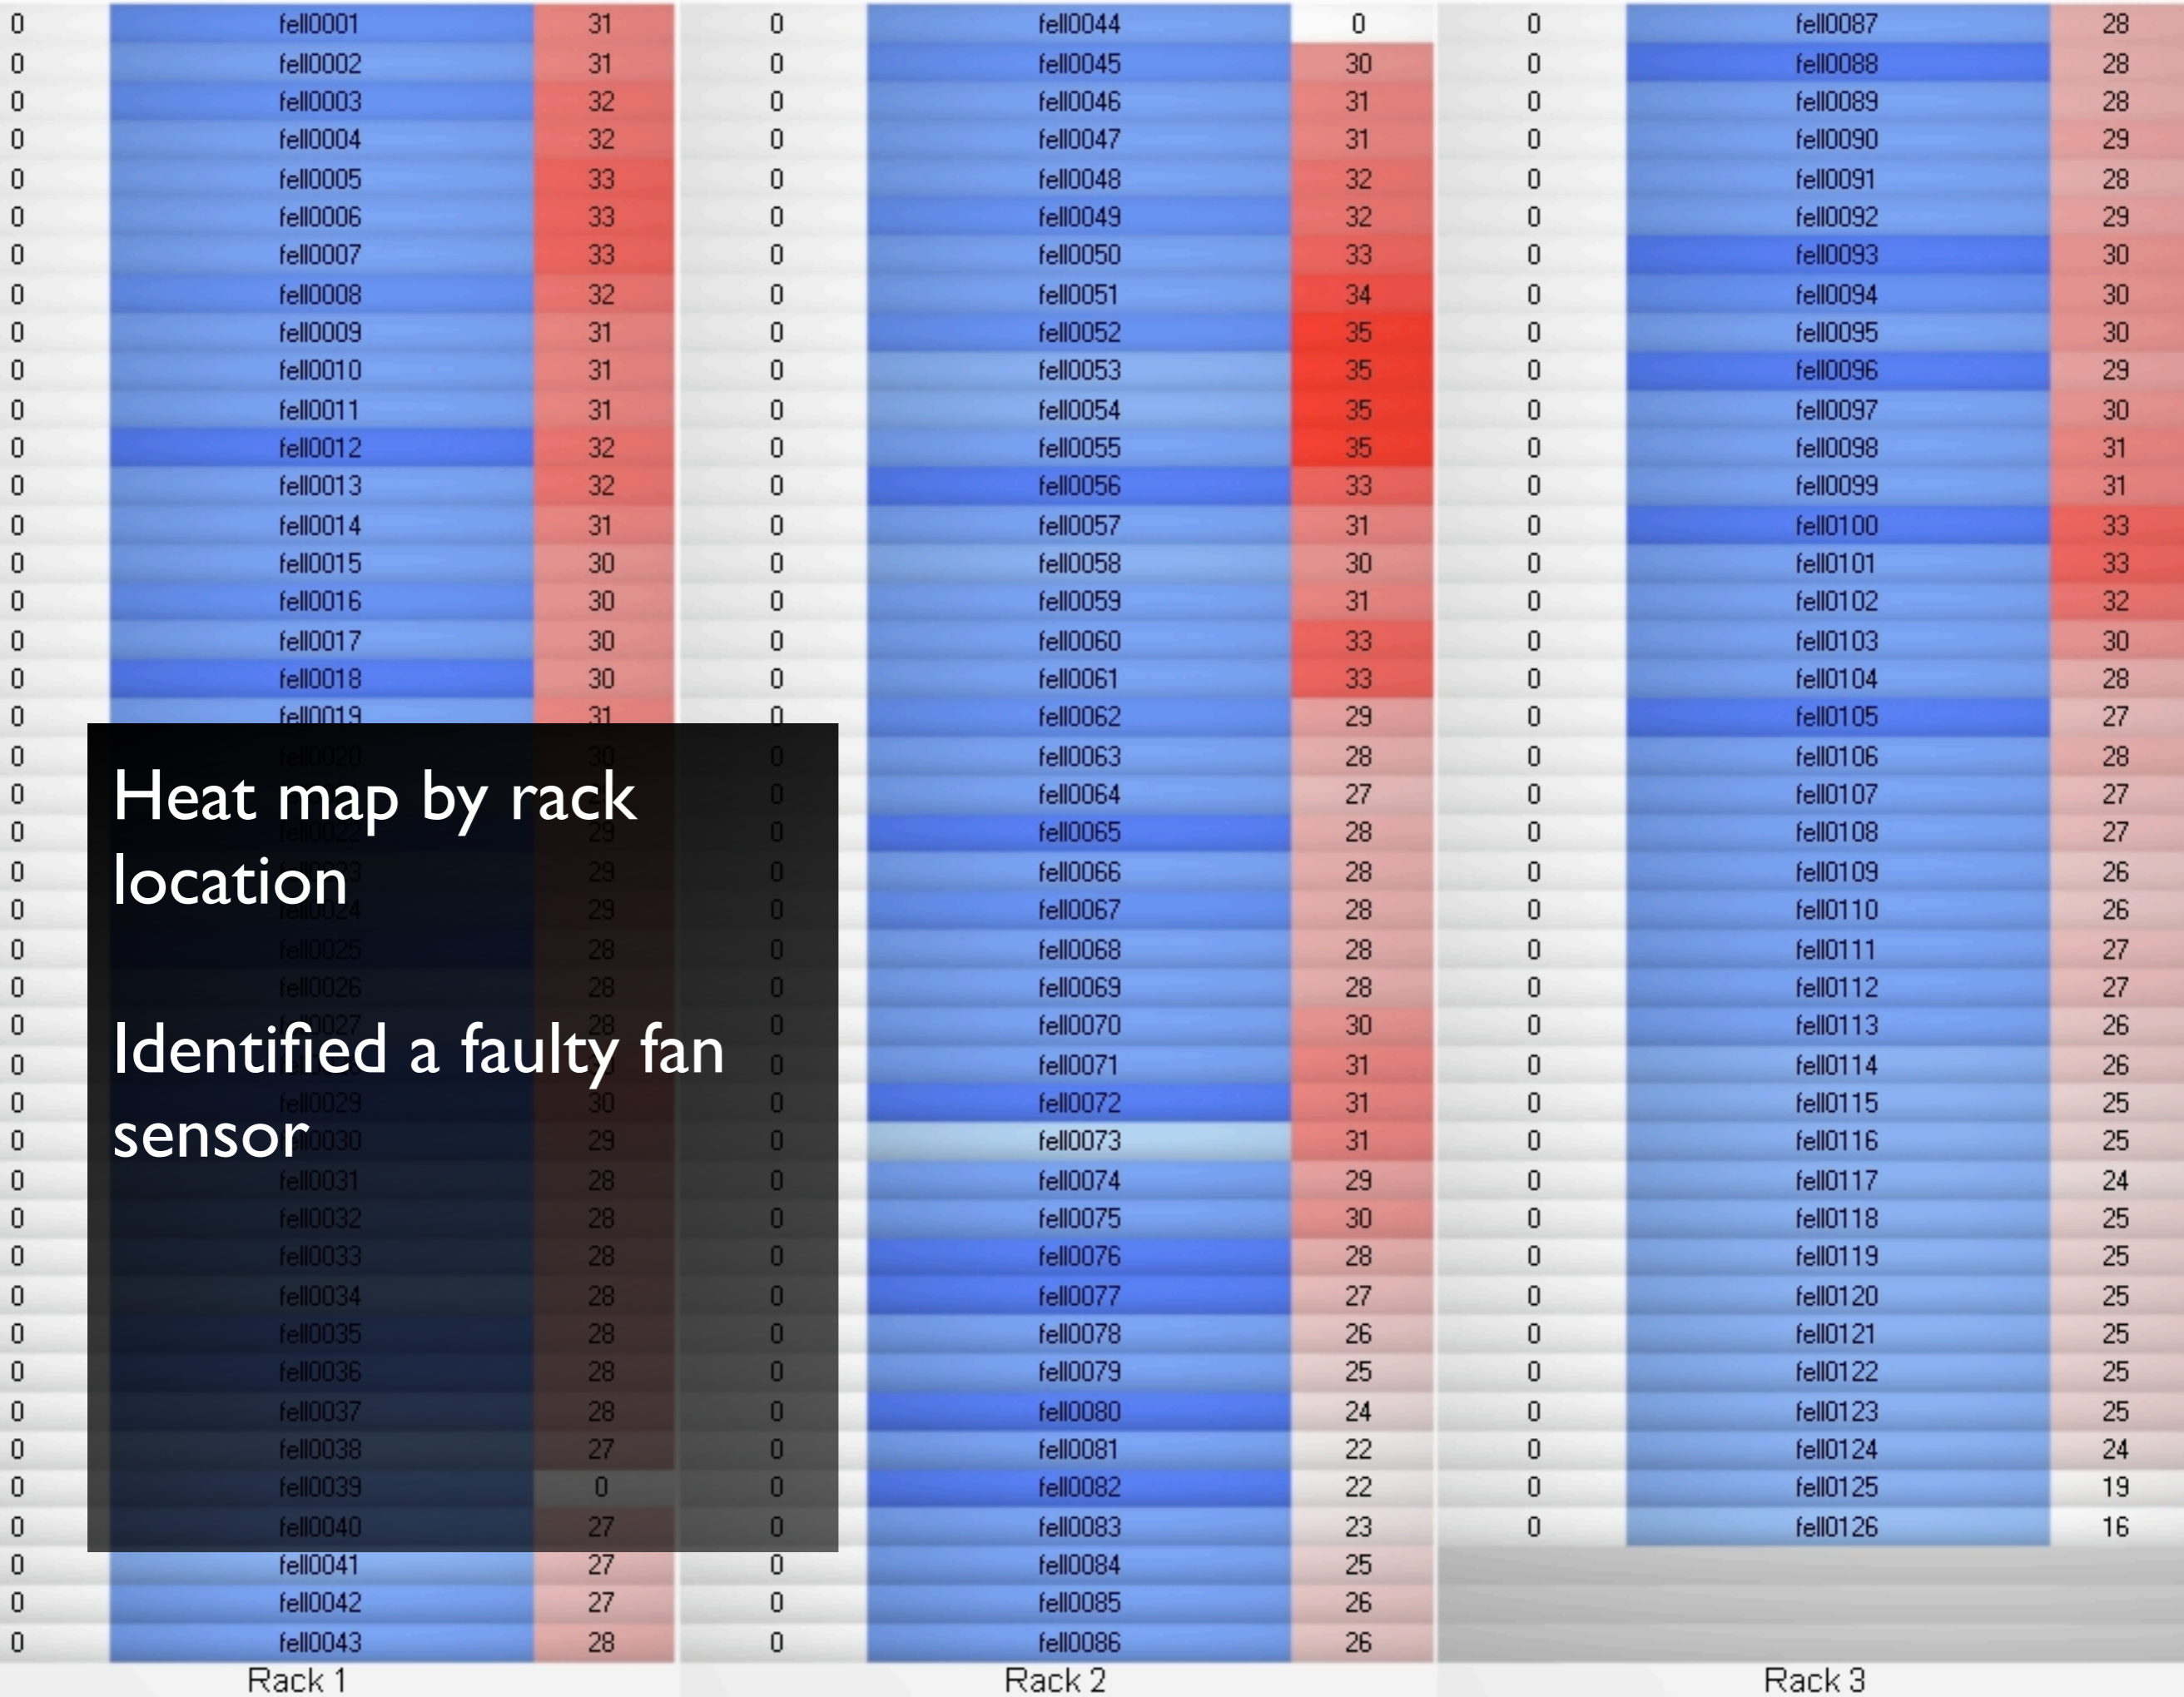
# Batch Slot Occupancy

% of job slots on each host occupied

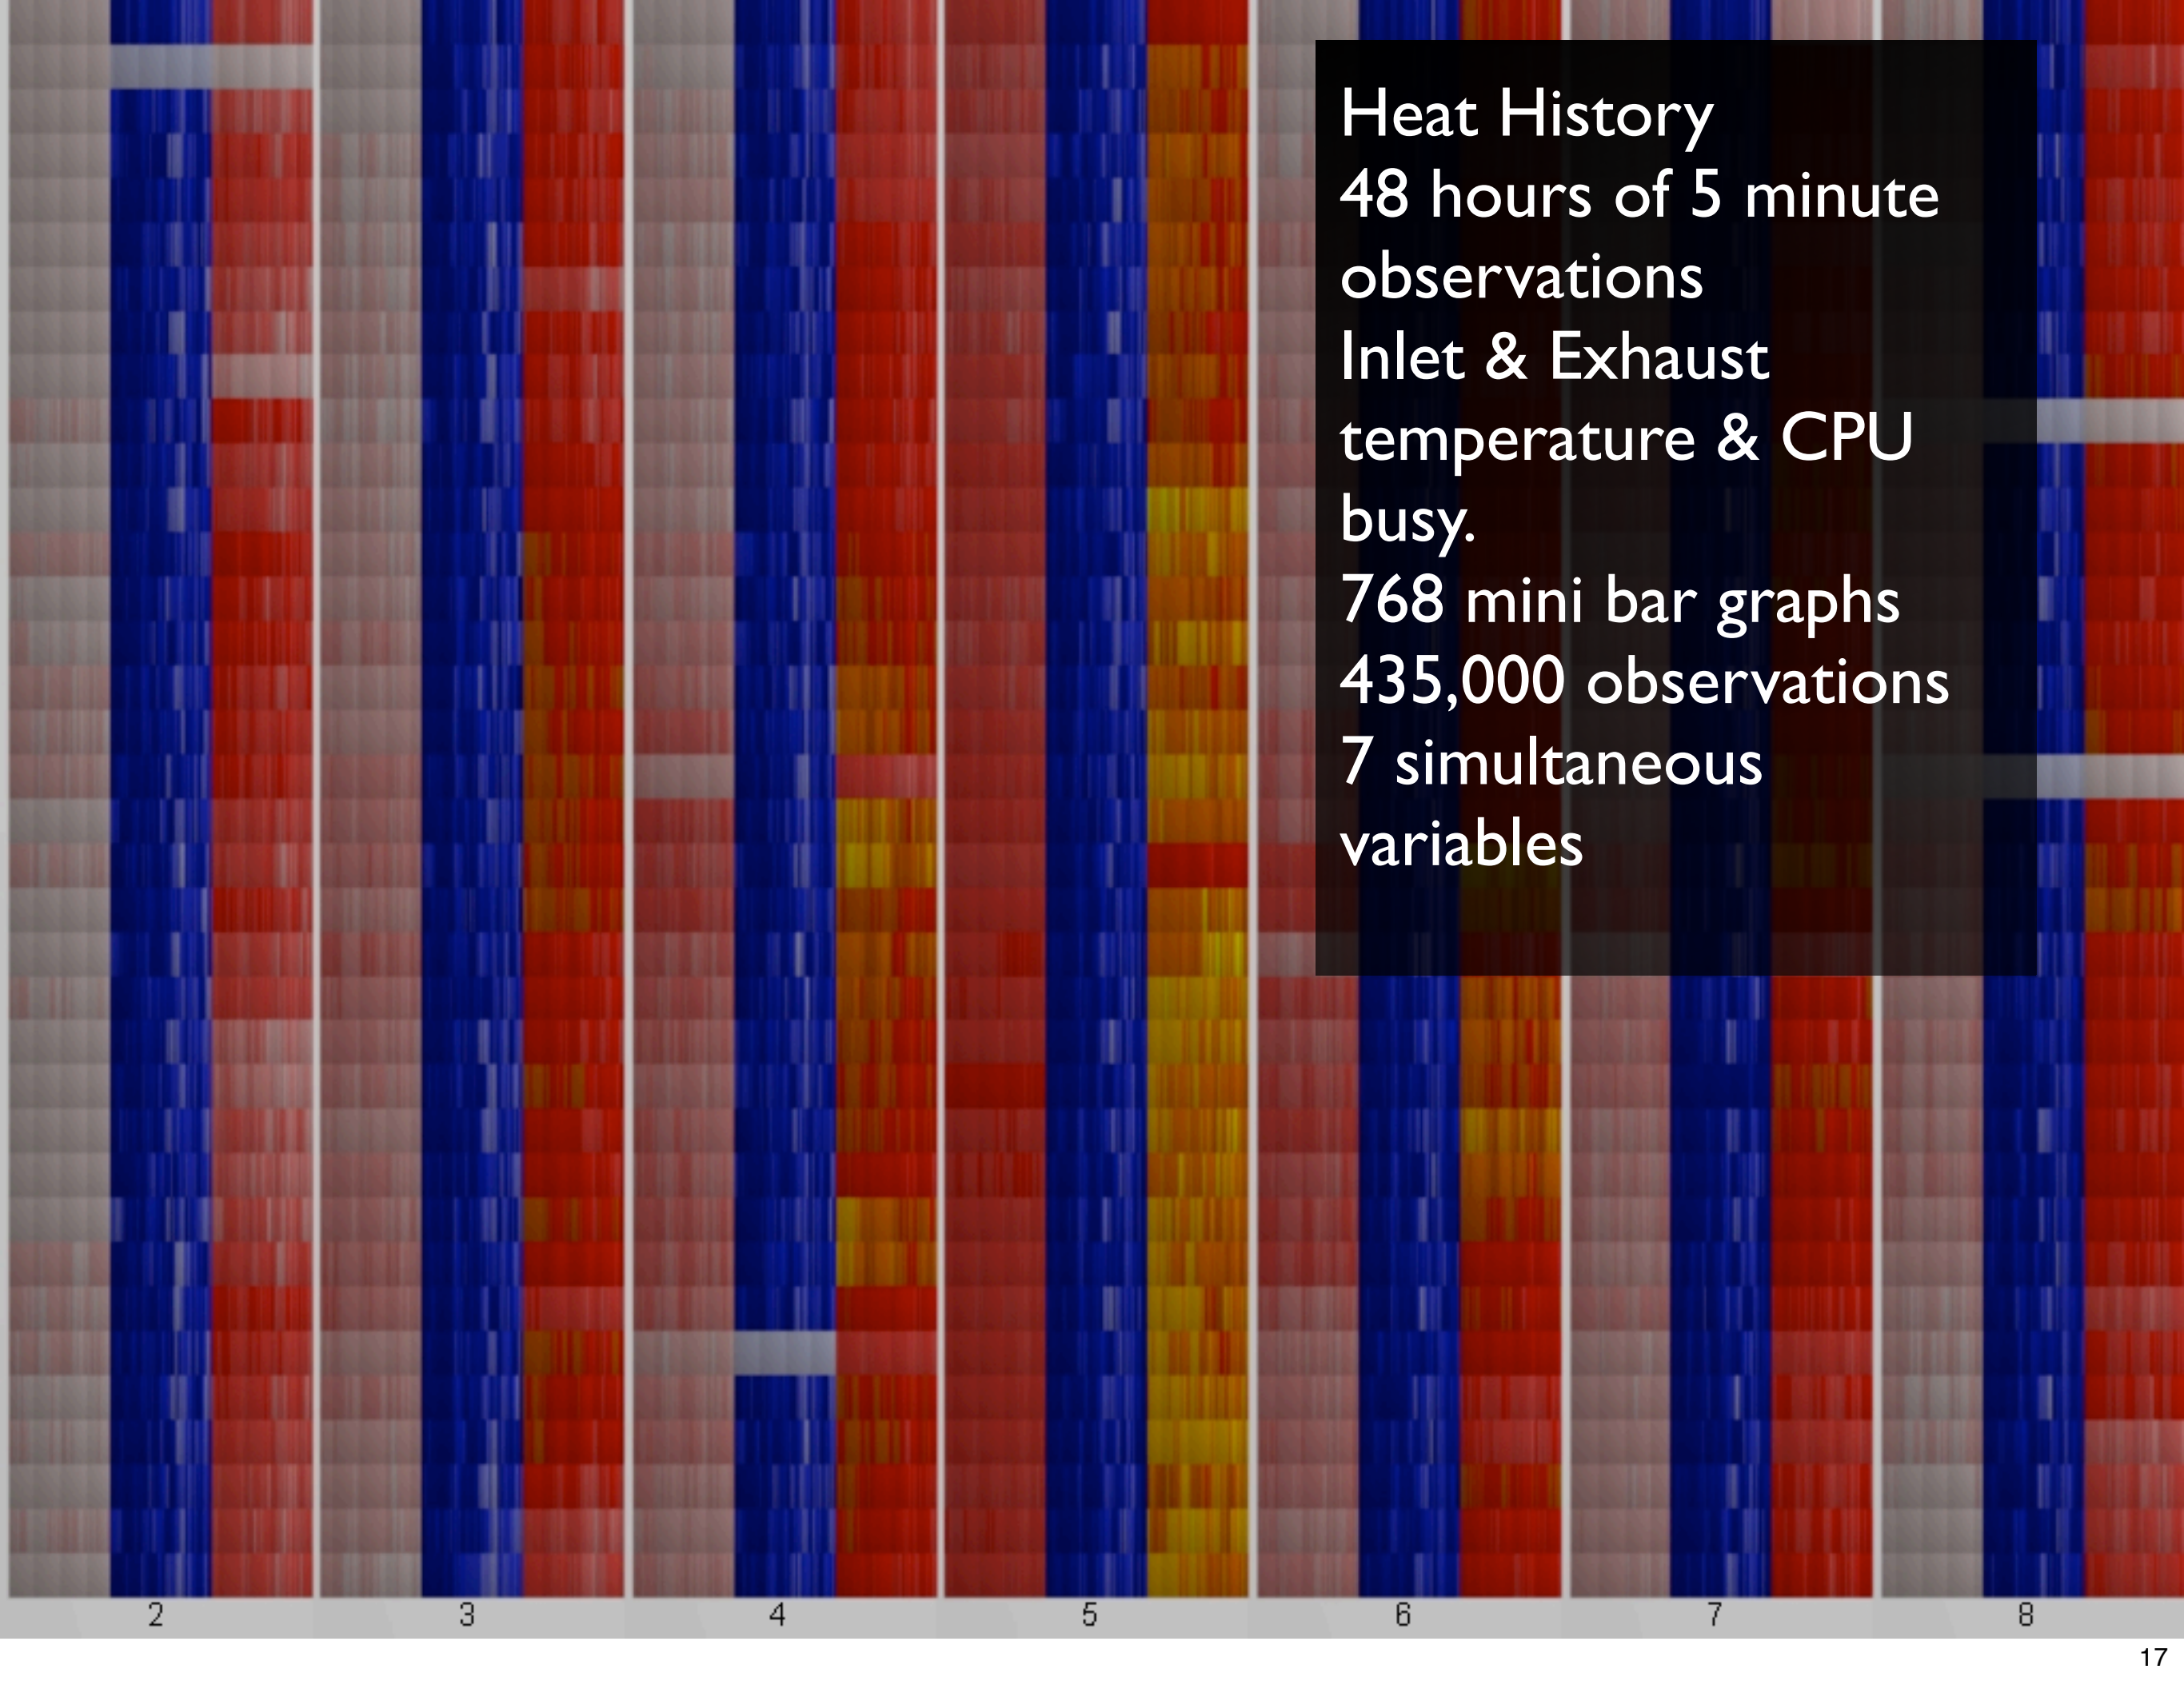


# I/O Activity map









Heat History  
48 hours of 5 minute  
observations  
Inlet & Exhaust  
temperature & CPU  
busy.  
768 mini bar graphs  
435,000 observations  
7 simultaneous  
variables

2

3

4

5

6

7

8



A map of the storage pools at SLAC

Color is coded to indicate the most dangerously full file systems



# Conclusions

Dense displays allow more meaningful information to reach the user

There are over 18 million bits on a screen - use them

GUIs can scale

Questions?