The astronomer  , the software engineer  , and the cloud 

Frossie Economou, Josh Hoblitt,
J. Matt Peterson, Jonathan Sick
Science Quality and Reliability Engineering team
Large Synoptic Survey Telescope
Greg Daues,
Univ. of Illinois at Urbana-Champaign

or

How I Learned to Stop Worrying and Love Silicon Valley

Advantages

- Developer self-service == Happiness & Profit!

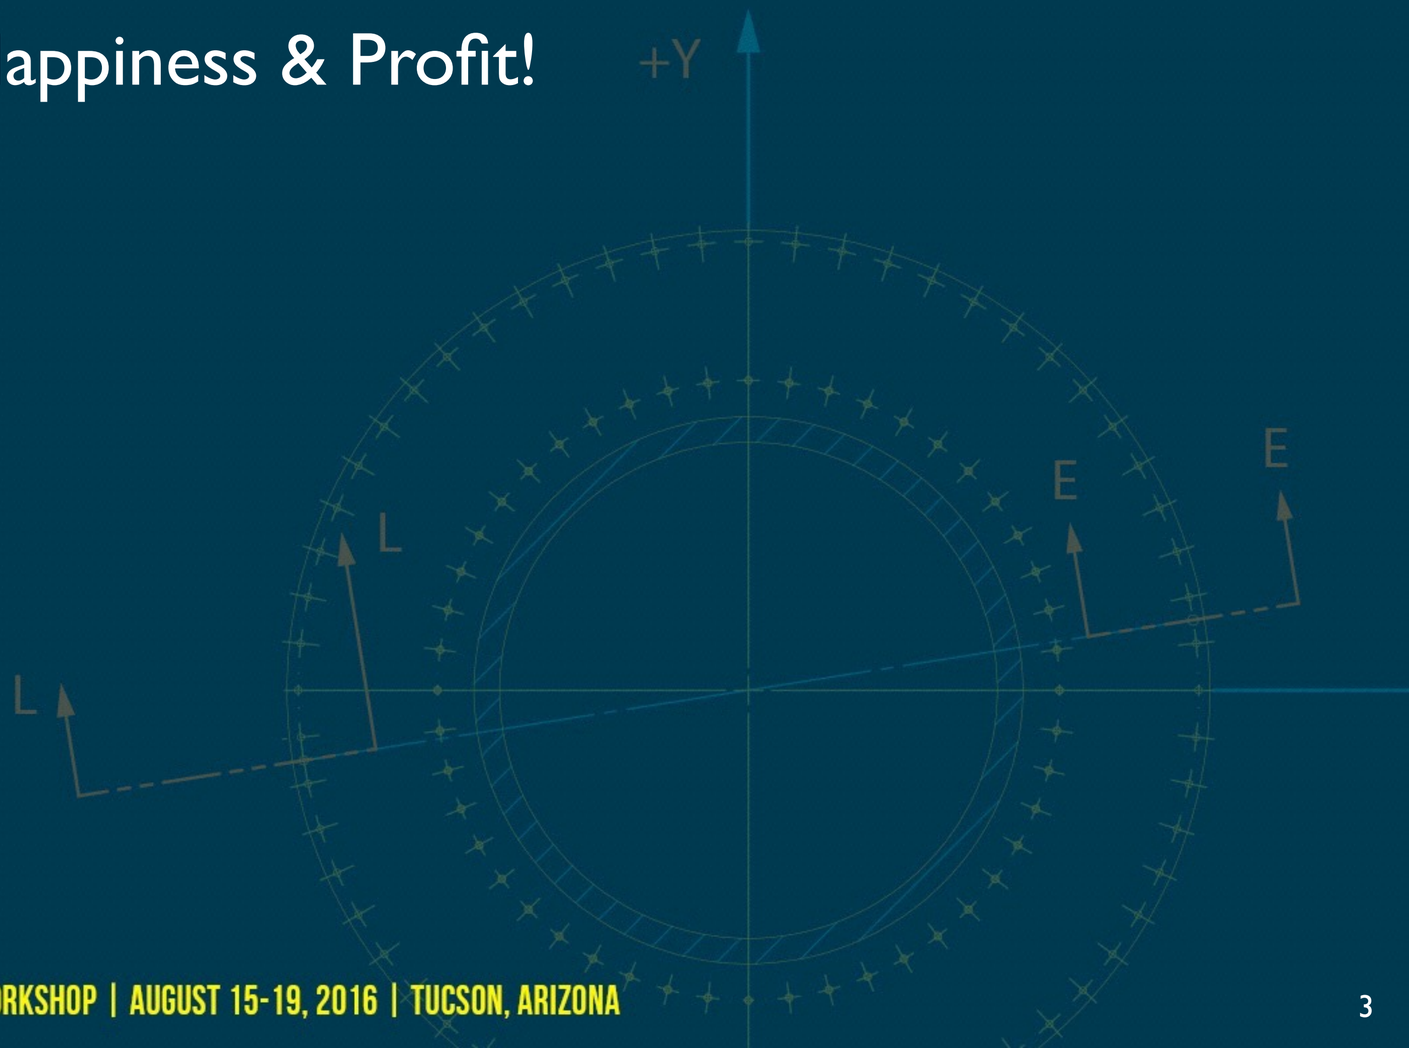

256 MB

This is what living in the future looks like

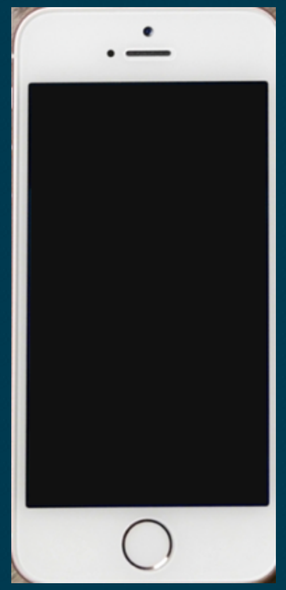

2 GB

Scarcity Mindset

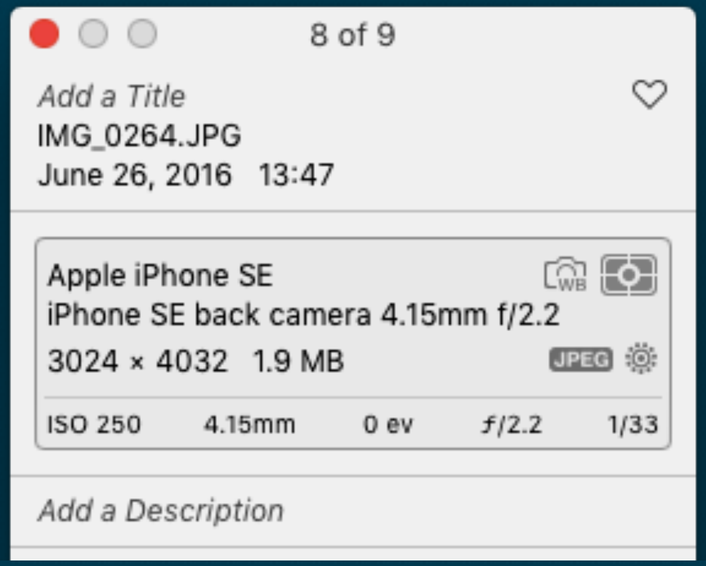

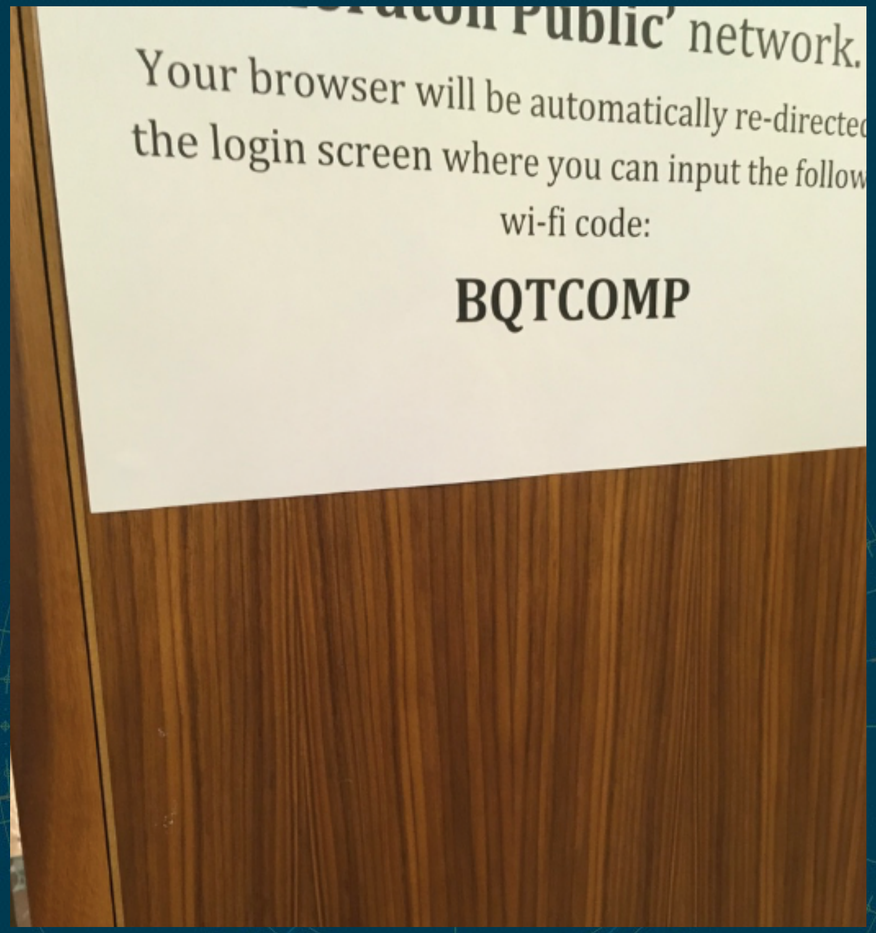

- The trade-off between human time and resources (compute, storage) has changed radically

Computers are not pets

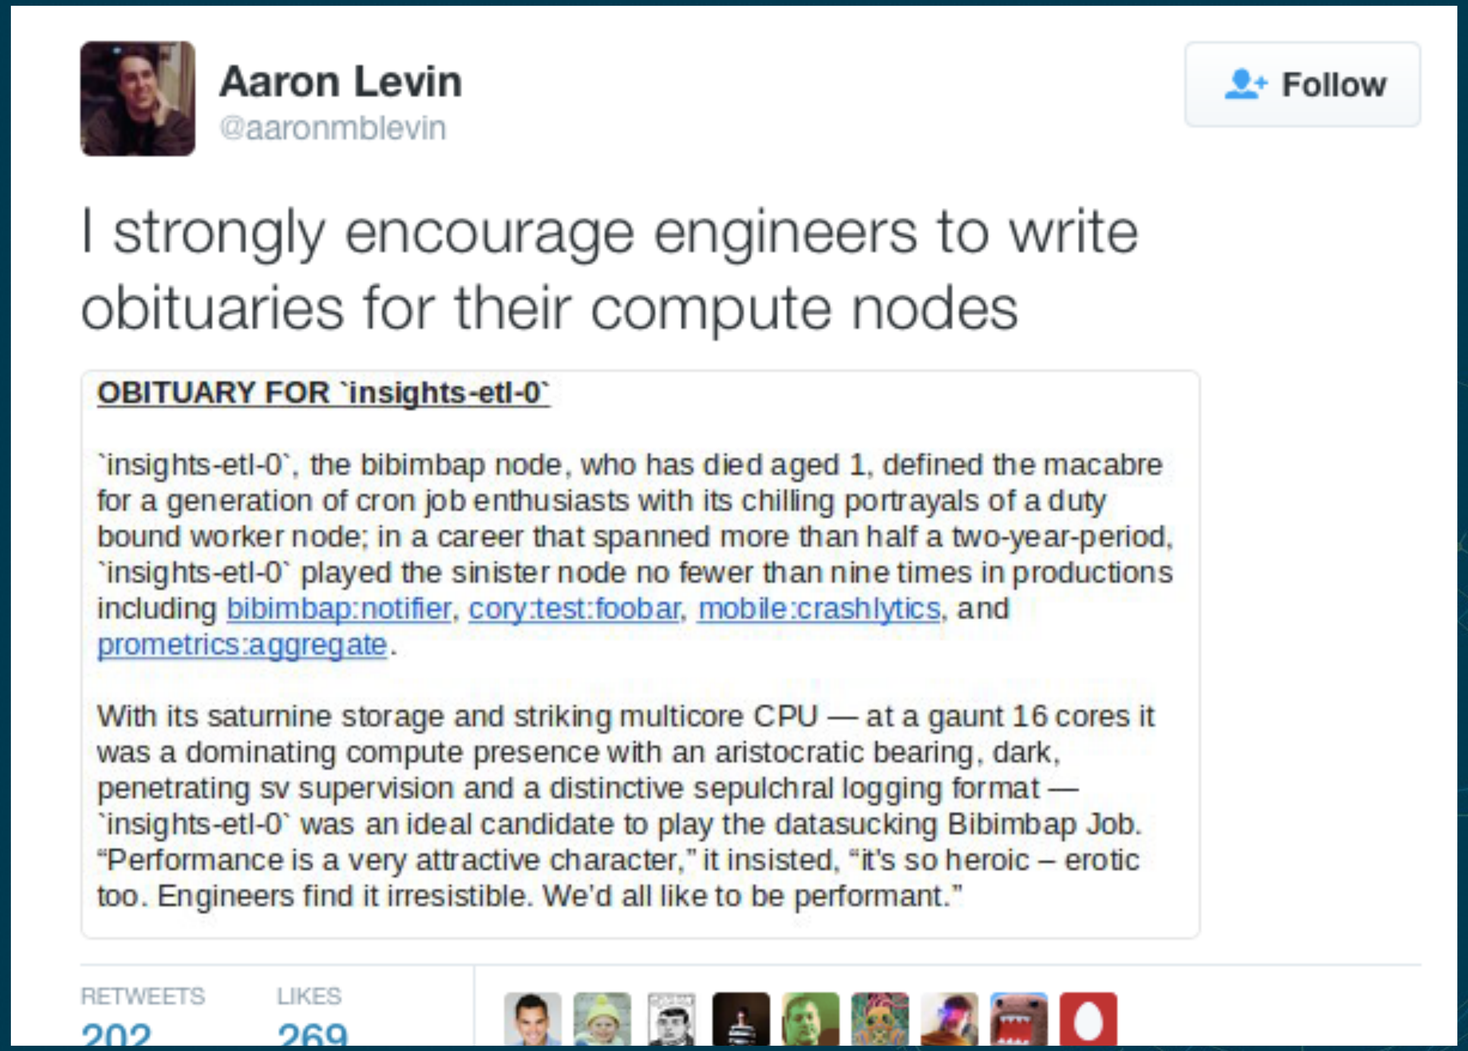

Aaron Levin
@aaronmblevin

Follow

I strongly encourage engineers to write obituaries for their compute nodes

OBITUARY FOR `insights-etl-0`

`insights-etl-0`, the bibimbap node, who has died aged 1, defined the macabre for a generation of cron job enthusiasts with its chilling portrayals of a duty bound worker node; in a career that spanned more than half a two-year-period, `insights-etl-0` played the sinister node no fewer than nine times in productions including bibimbap:notifier, cory:test:foobar, mobile:crashlytics, and prometrics:aggregate.

With its saturnine storage and striking multicore CPU — at a gaunt 16 cores it was a dominating compute presence with an aristocratic bearing, dark, penetrating sv supervision and a distinctive sepulchral logging format — `insights-etl-0` was an ideal candidate to play the datasucking Bibimbap Job. "Performance is a very attractive character," it insisted, "it's so heroic – erotic too. Engineers find it irresistible. We'd all like to be performant."

RETWEETS 202 LIKES 269

Advantages

- Developer self-service
- Reliability

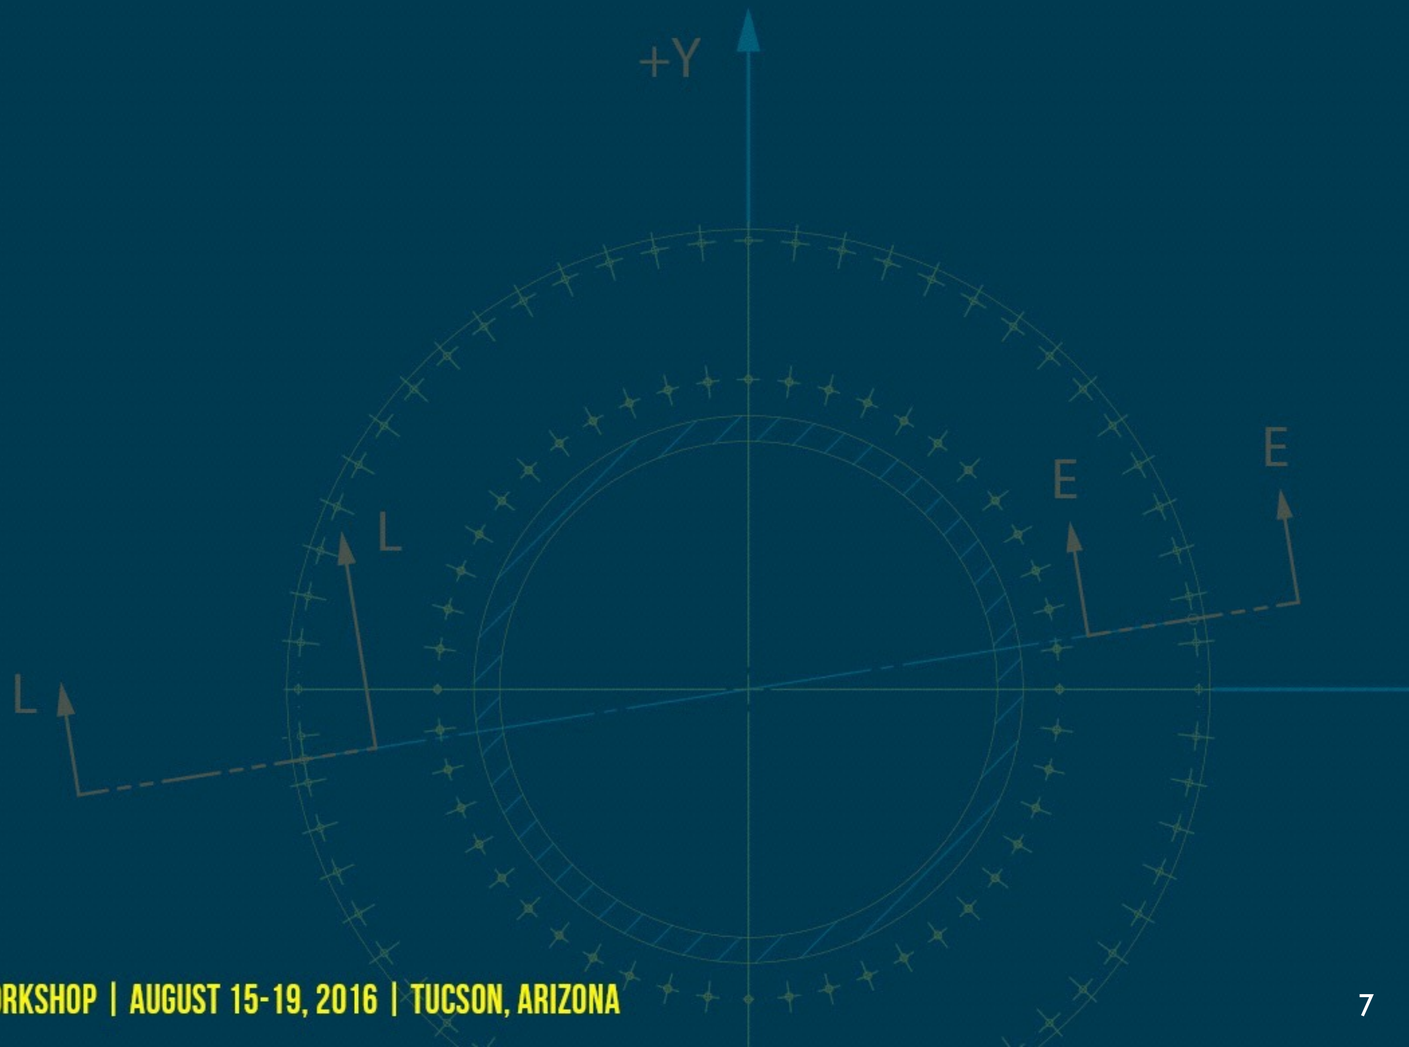

Reliability

- We always cared about uptime
- … maybe too much…
- Regardless, our users' bar for what reasonable availability is has changed

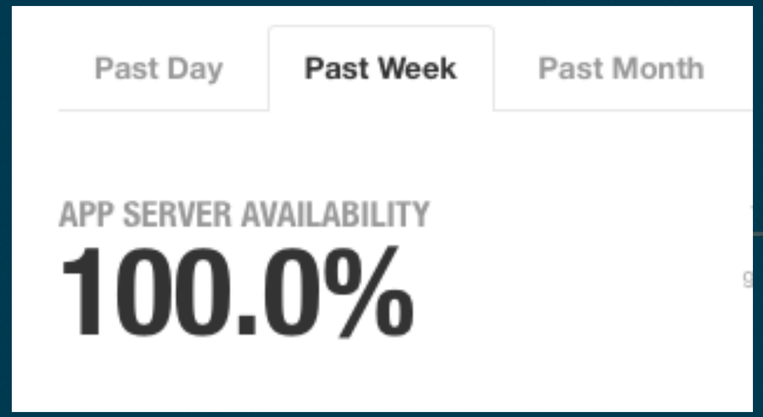

Advantages

- Developer self-service
- Reliability
- Re-use (double-dip, triple-dip, N-dip)

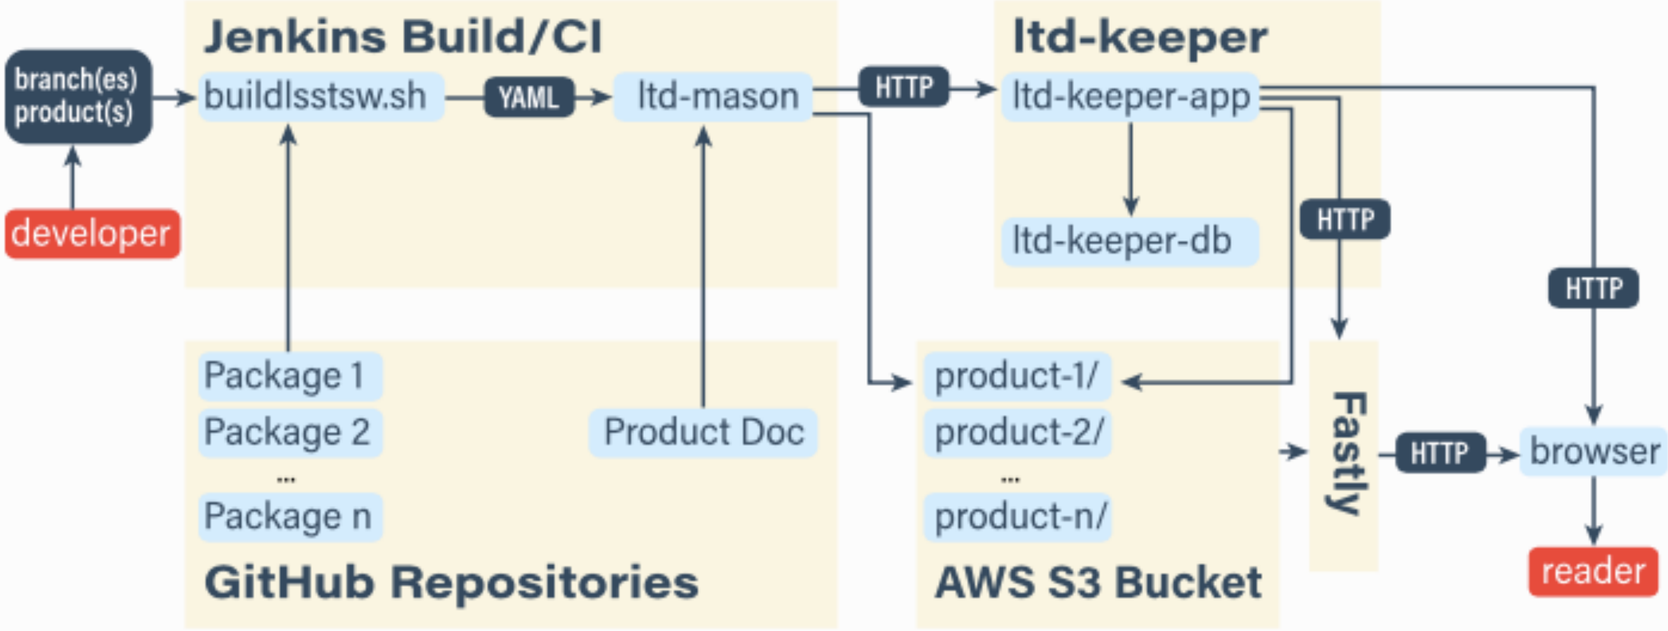

Figure 1 Architecture of LSST the Docs (LTD).

Advantages

- Developer self-service
- Reliability
- Re-use (double, triple, N-dip)
- Architecture (modularity)

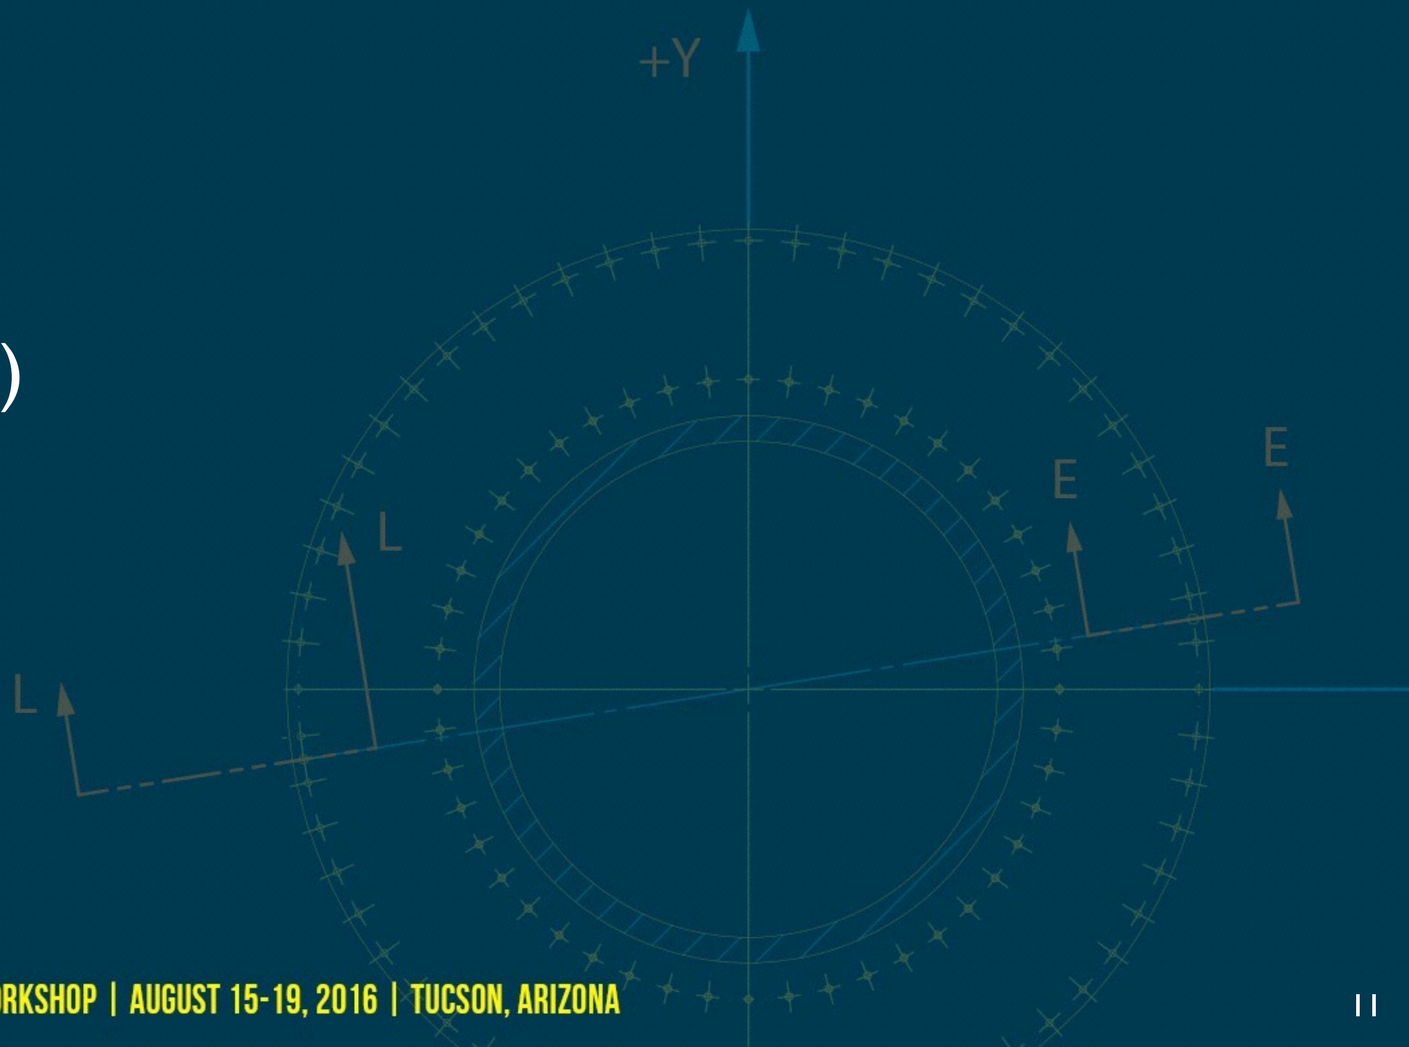


Jenkins


GitHub


slack


JIRA

Advantages

- Developer self-service
- Reliability
- Re-use (double, triple, N-dip)
- Architecture (modularity)
- Scalability

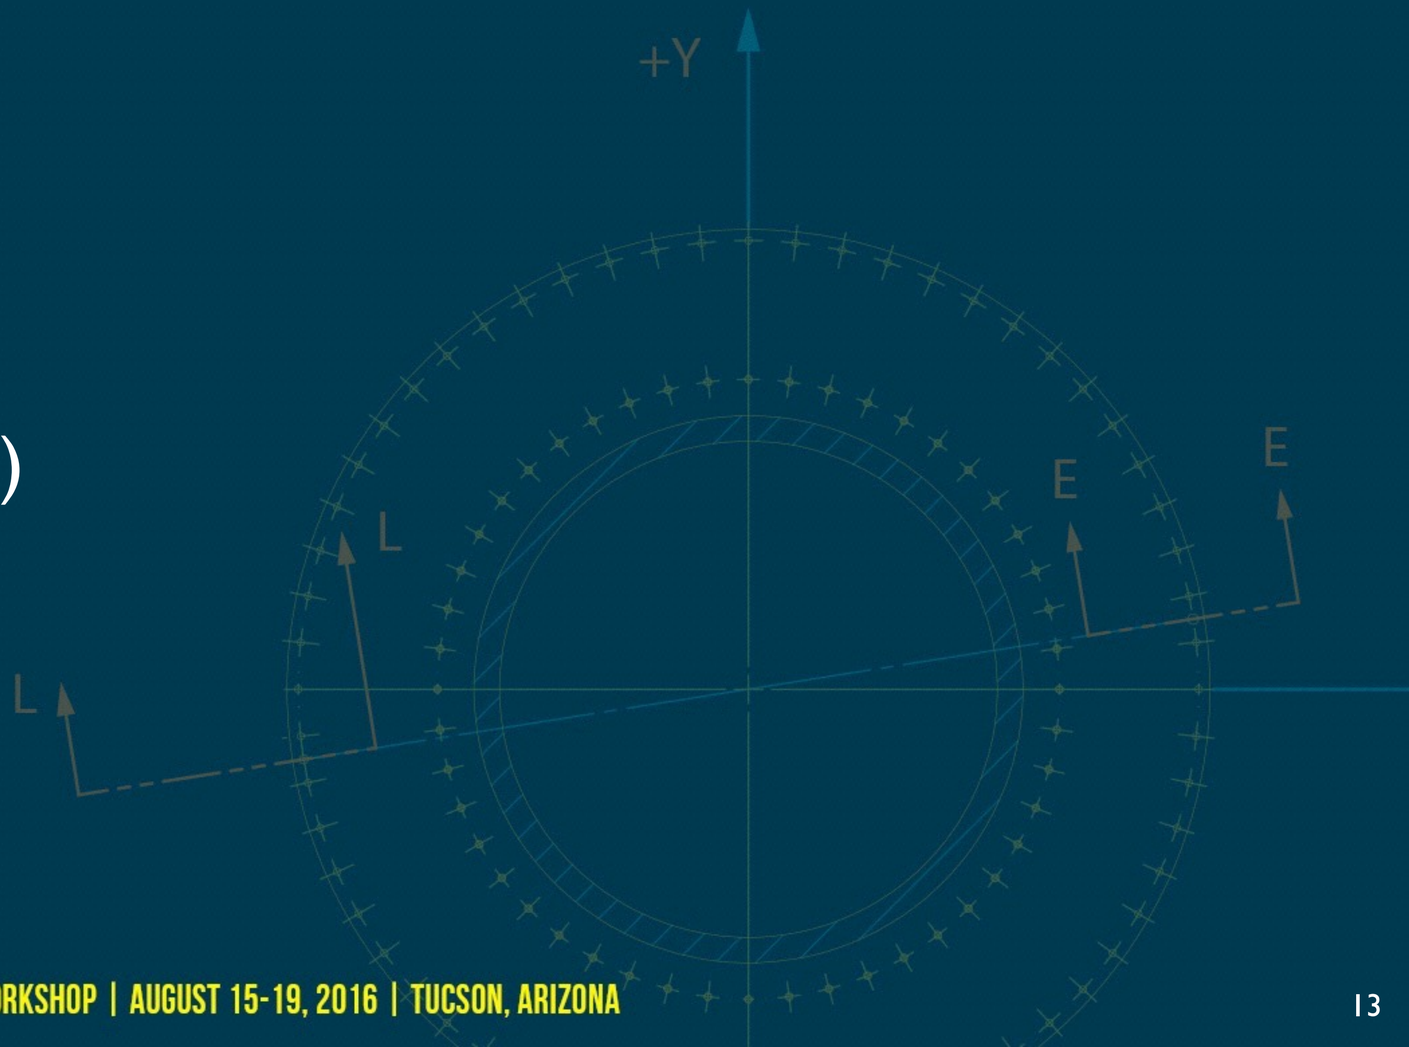

Concerns

- Cost 💸
- Cost 💸
- Cost 💸

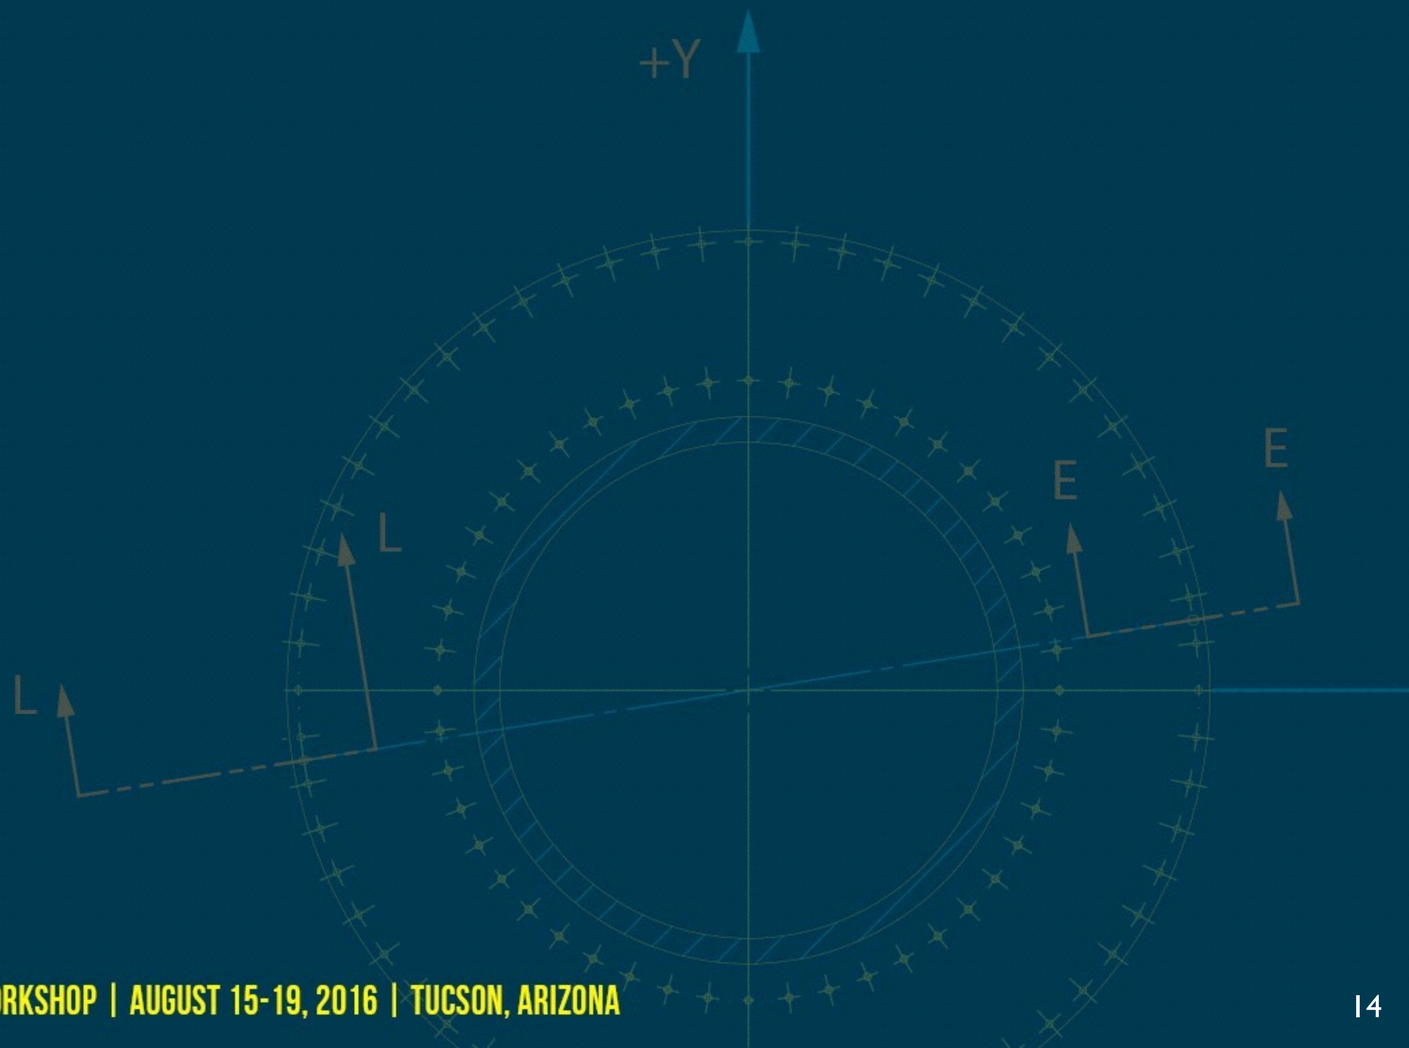

The problem is “data out”… or is it?

It's not just about compute

- Route 53: $1.54/mo
- Data Transfer: $33.31/mo
- File storage: $9.68/mo

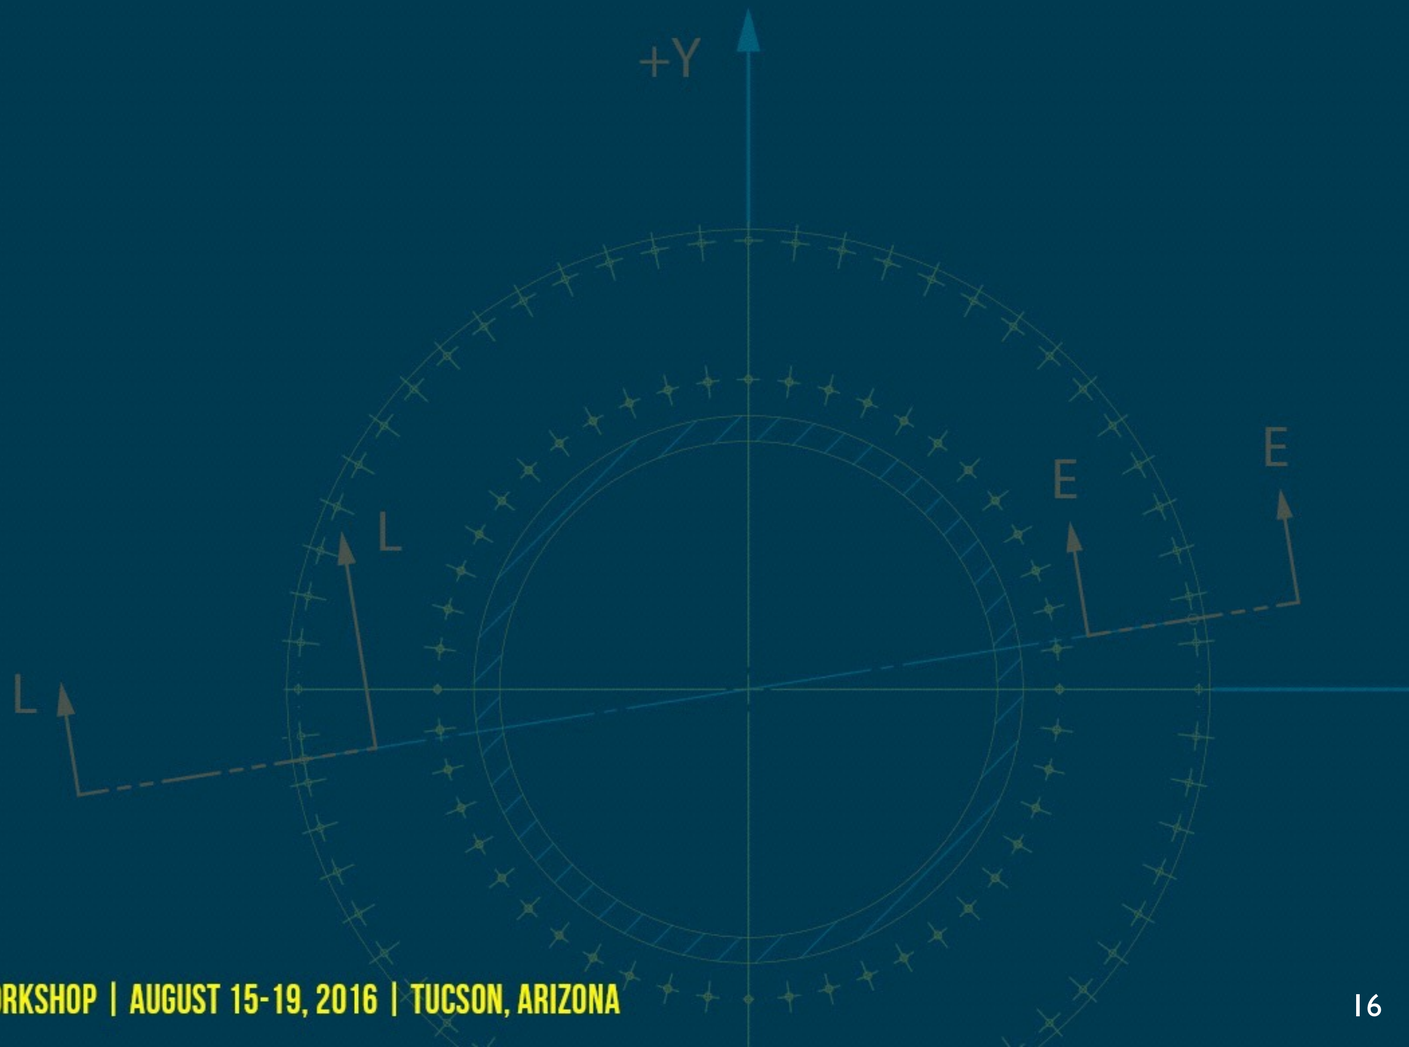

IT is never “free as in beer”

- Sysadmins (set-up, upgrades, RMAs)
- Overhead on sysadmins
- Floor space
- Power (racks & A/C)
- Security
- Inventory

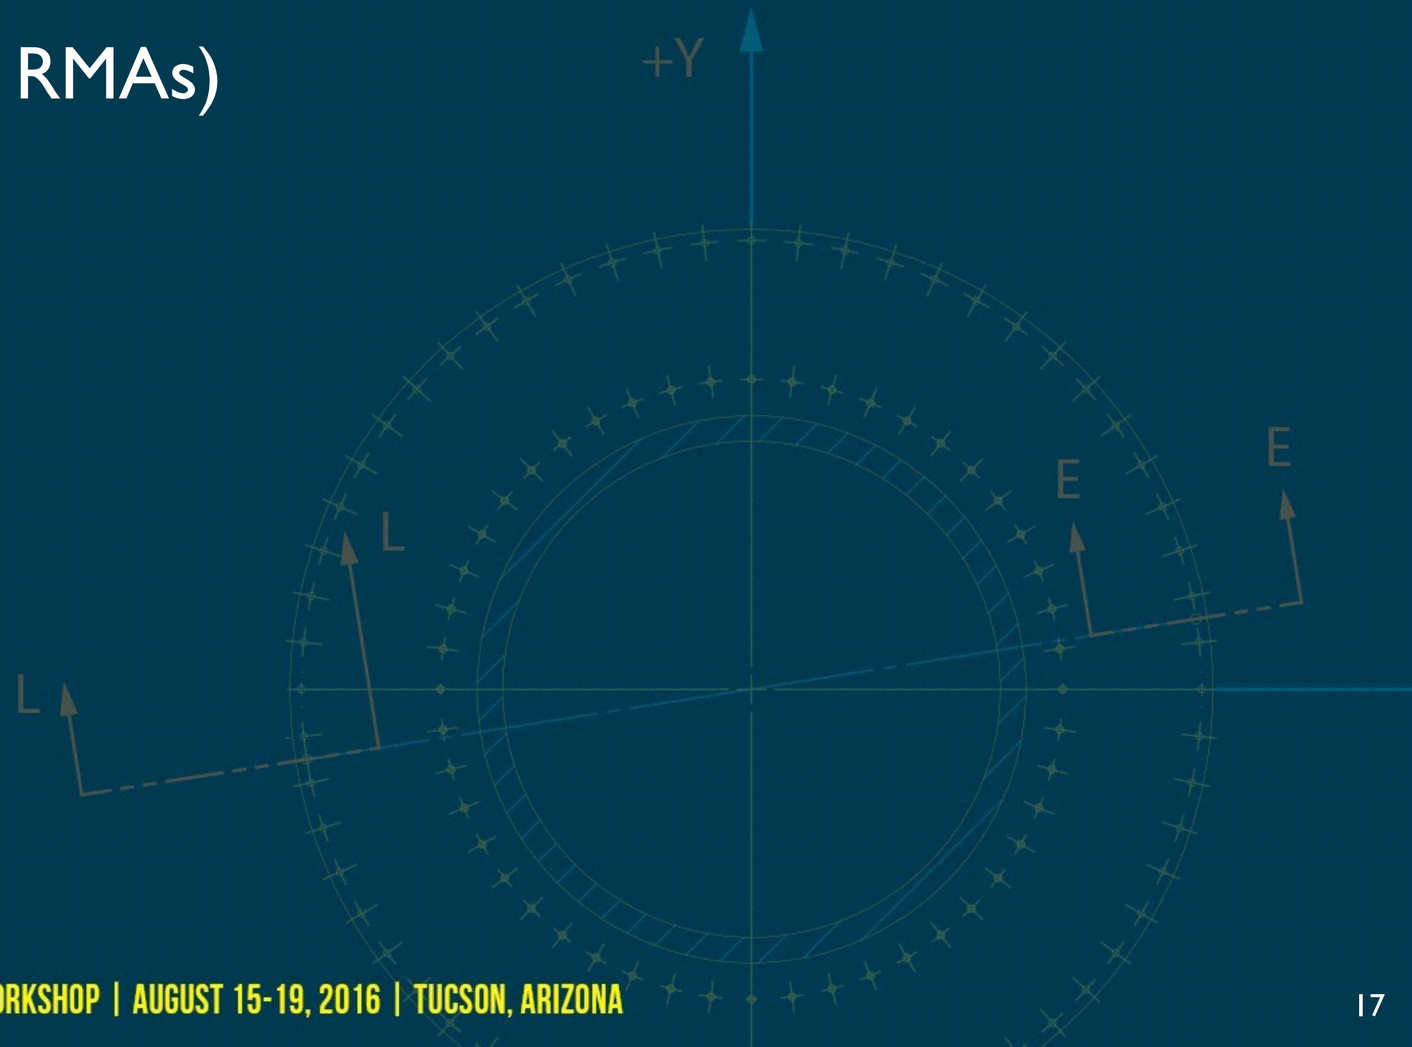

Private clouds are clouds too

- OpenStack has similar architectural advantages
- Non-trivial to set-up and operate
  - University/Lab system
  - Buy hosted (eg. HP Helion)
  - Buy turnkey (eg. Mirantis)
  - Your Funding Agency here? <- your science user almost certainly

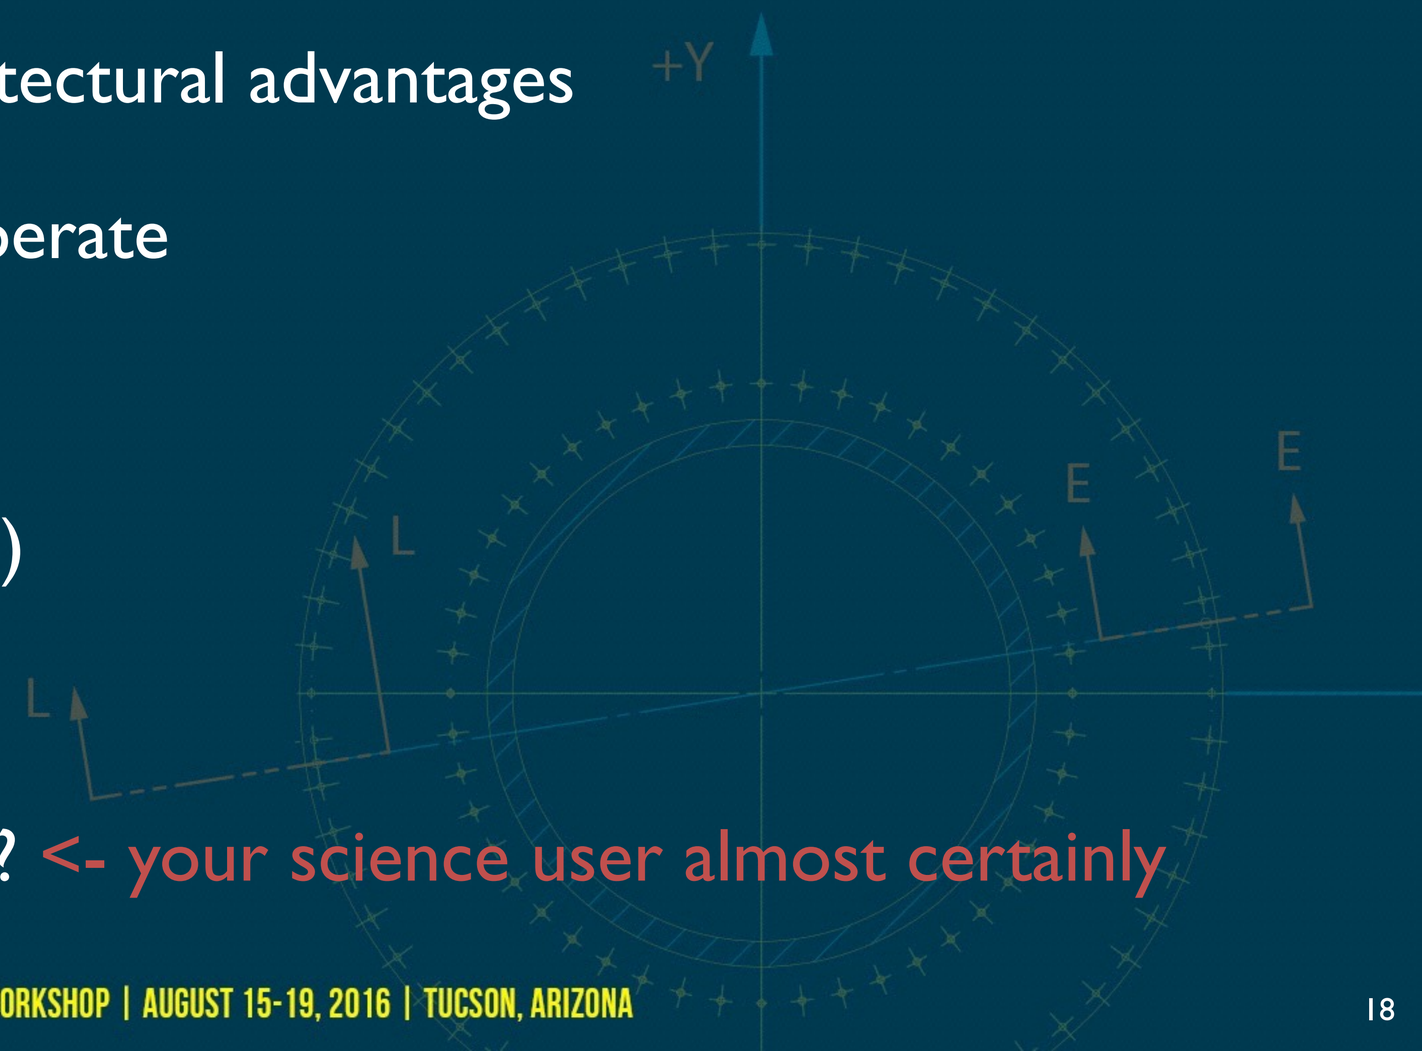

Concerns

- Cost 💸
- Cost 💸
- Cost 💸
- Purchasing rules 💸
- Grant rules 💸

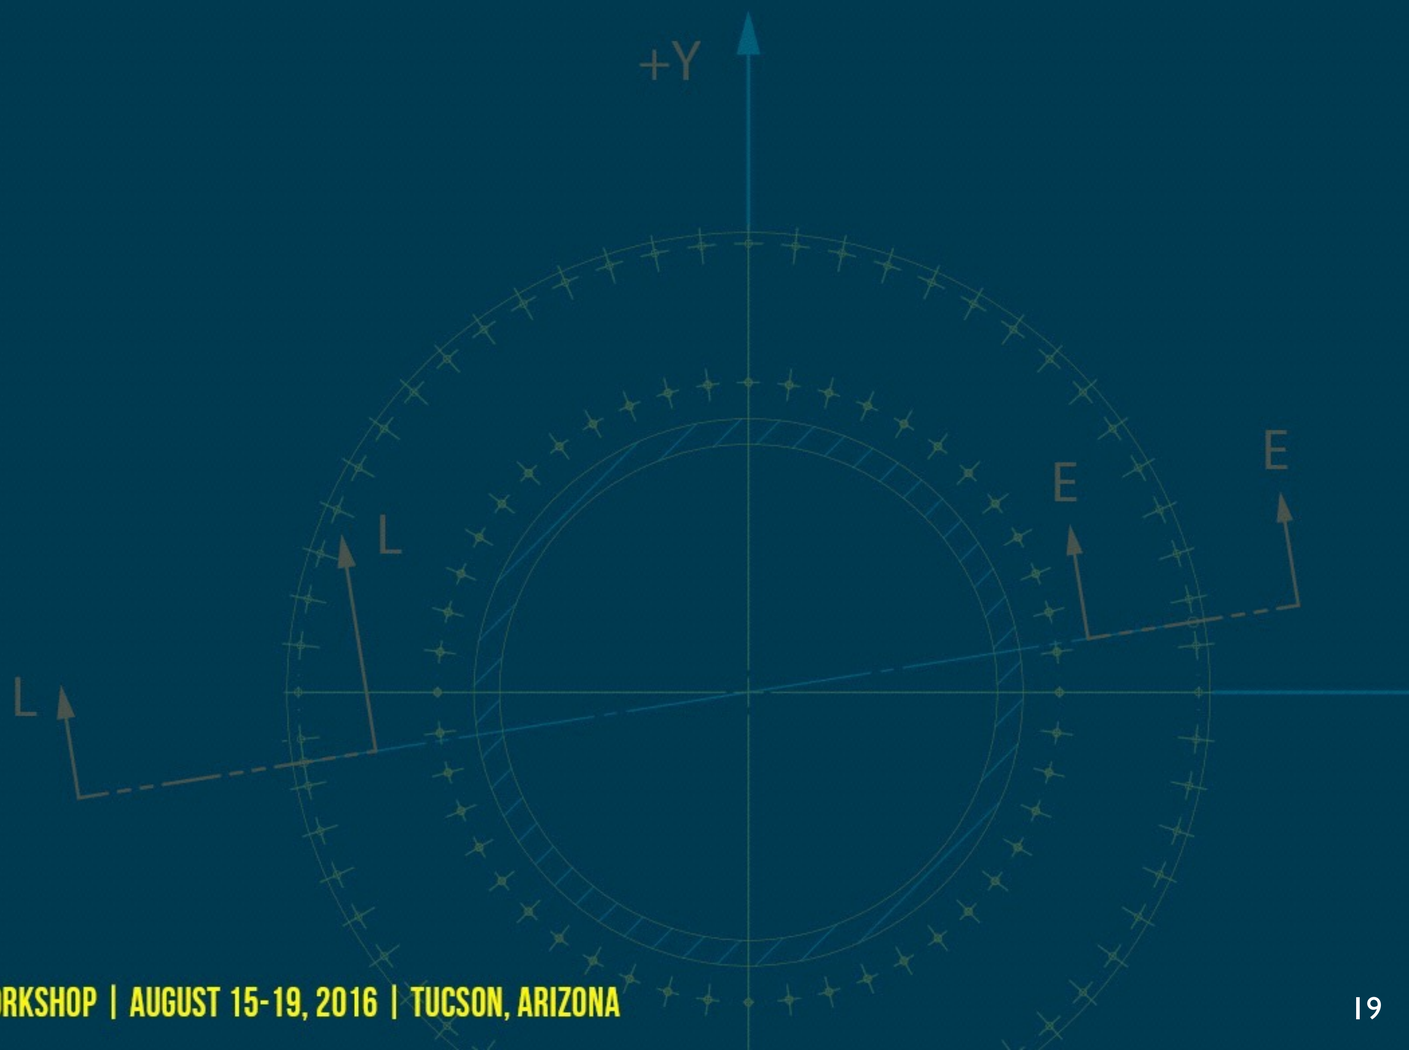

Concerns

- Cost 💸
- Cost 💸
- Cost 💸
- Purchasing rules 💸
- Grant rules 💸
- Lock-in 🔒

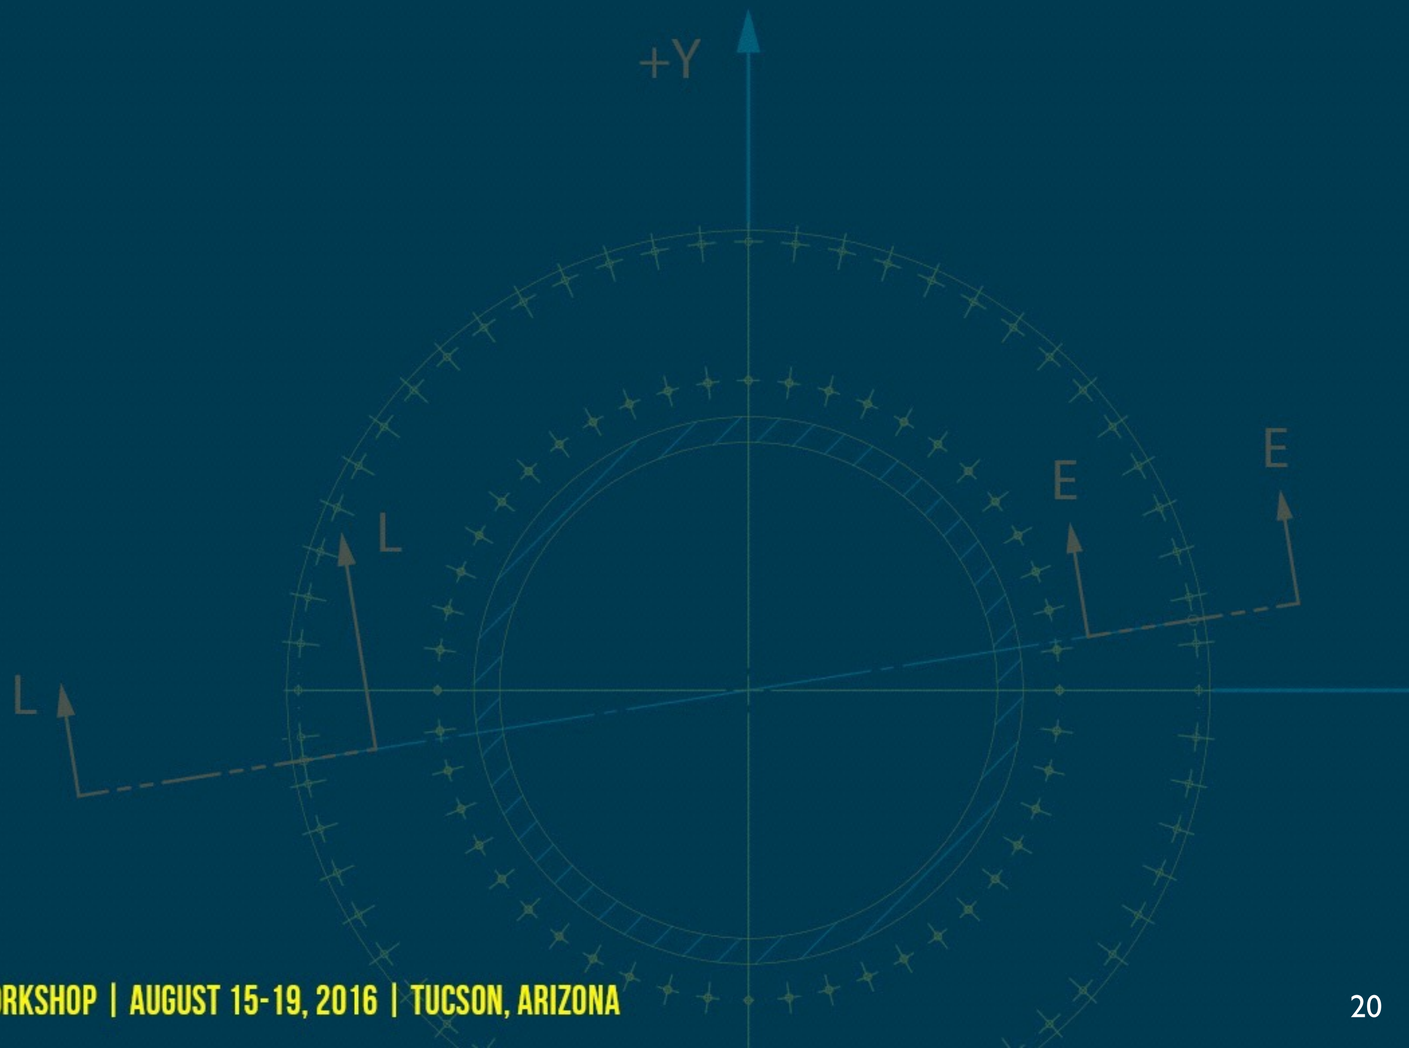

Lock-in? What lock-in?

- Migration paths easier than they have ever been
- Net advantage in keeping up with the “cool kids”
- Once you have seen, you cannot unsee

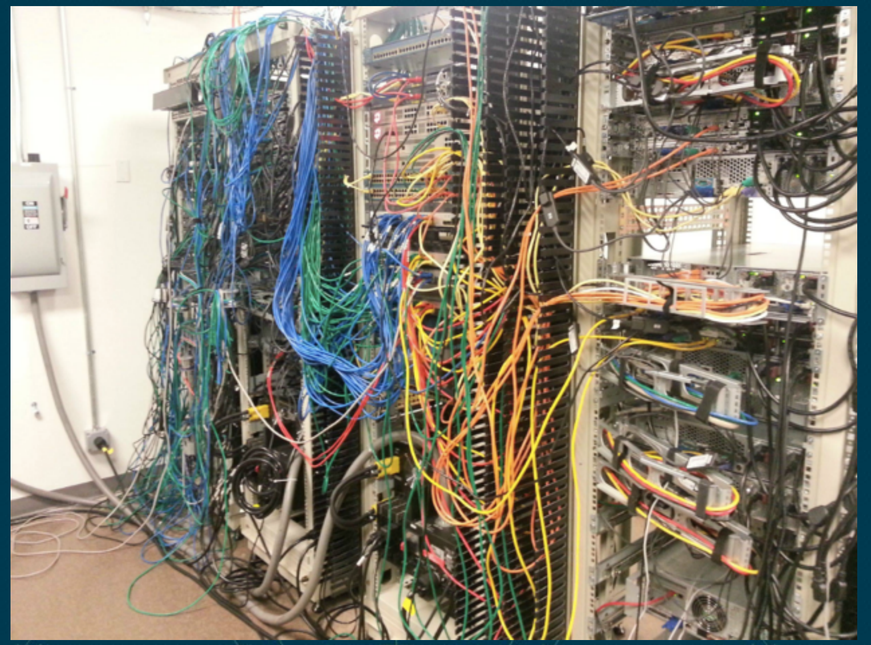

Concerns

- Cost 💸
- Cost 💸
- Cost 💸
- Purchasing rules 💸
- Grant rules 💸
- Lock-in 🔒
- Evilness 👹

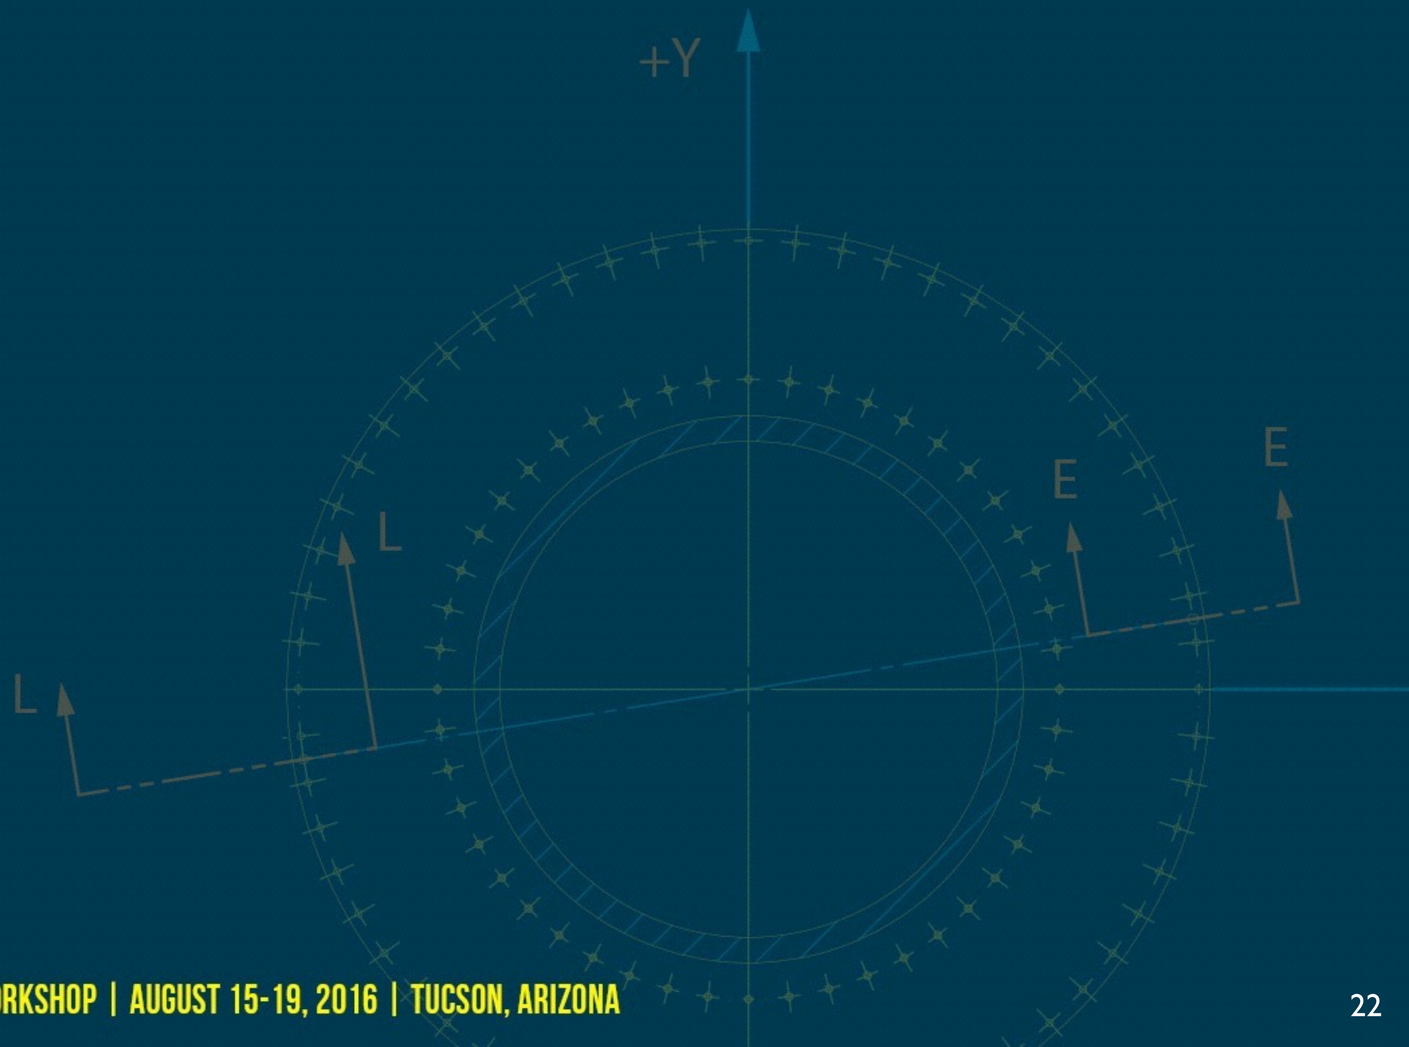

Cloud-hosted SQuaRE - Where?

- Commodity Cloud
  - AWS
  - Google Compute Engine
  - Digital Ocean
- Private Cloud
  - OpenStack
- SaaS
  - Github, TravisCI, Fastly, Hipchat (Slack), Dropbox, 1Password
- Own Metal
  - MacOS build slave

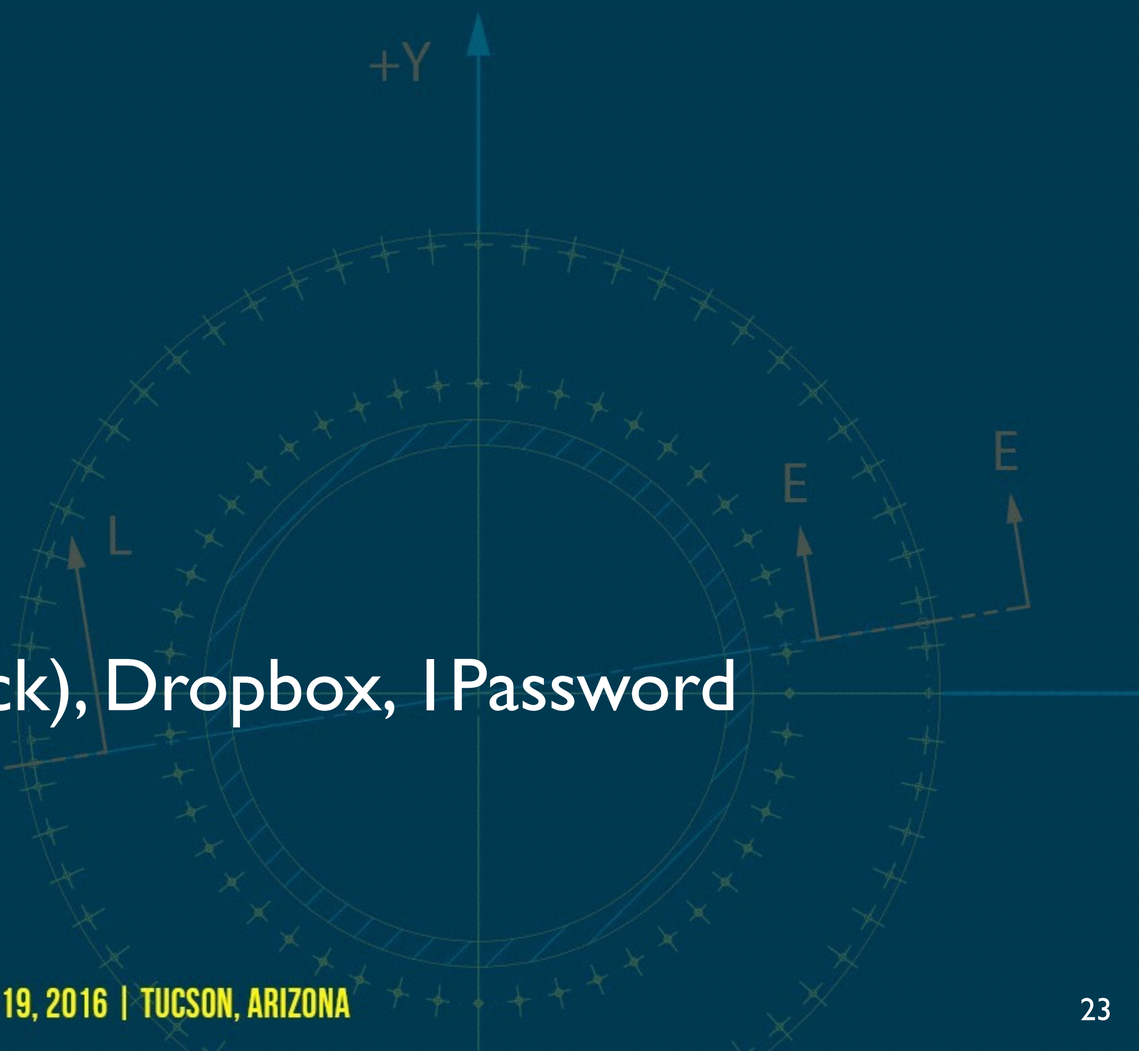

Cloud-hosted SQuaRE - What?

- Repository Management (Github)
- GitLFS service
- Documentation system (LSST-The-Docs)
- Continuous Integration (Jenkins, Travis CI)
- Science QA Harness (in-house)
- Monitoring system (ELK)
- Forum (Discourse)
- Release engineering (Vagrant/puppet)
- Not Cloud Hosted: JIRA

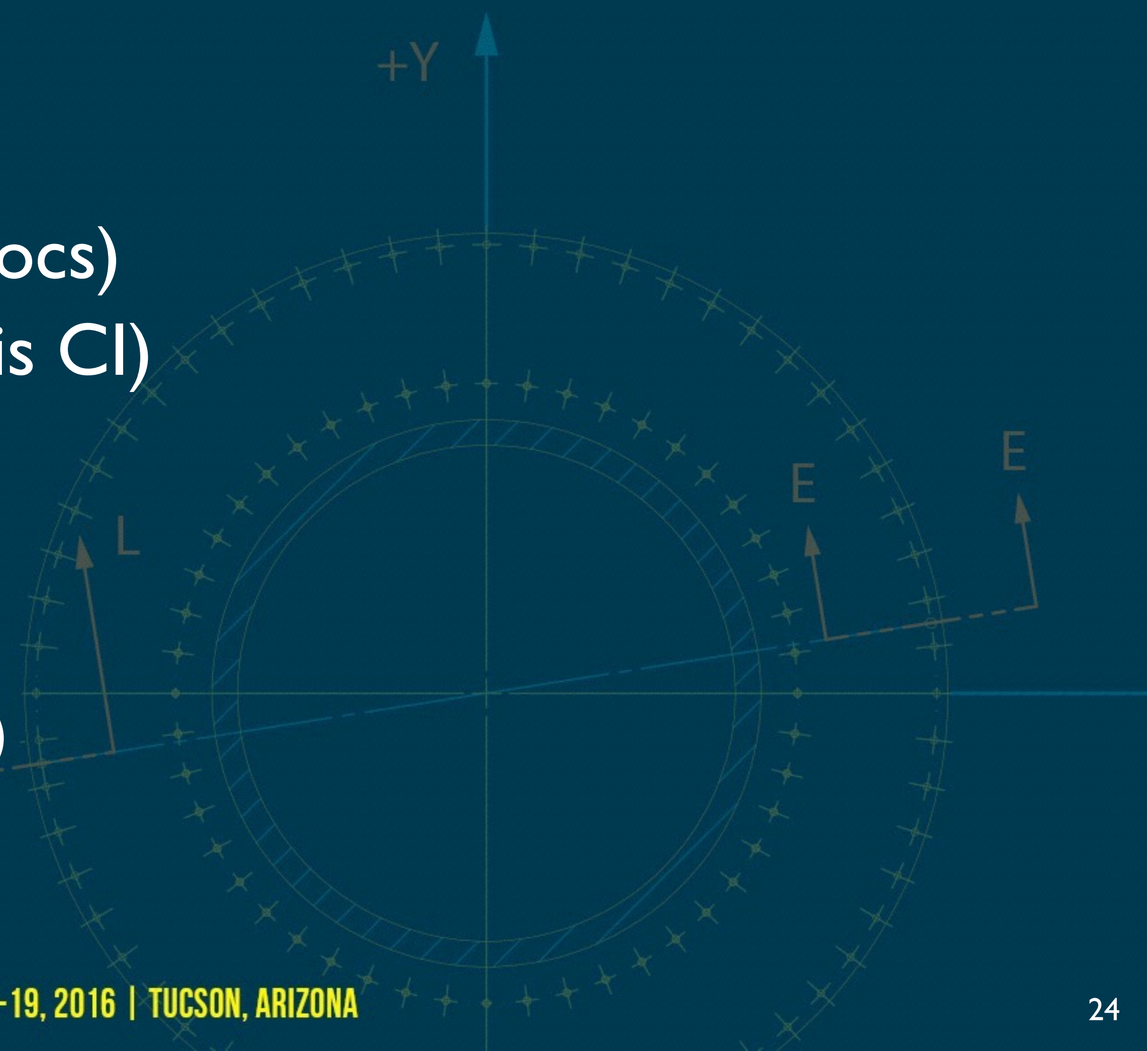

SQuaRE QA and user services tech stack

- AWS: EC2, S3, RDS, Route 53
- Google Compute Engine
- Openstack
- Fastly / Varnish
- S3 / Swift
- Kubernetes (Swarm? Mesos?)
- Puppet & Ansible
- Elasticsearch / (Logstash|fluentd|Riemann)/ Kibana
- Docker / virtualisation
- vagrant / packer / terraform

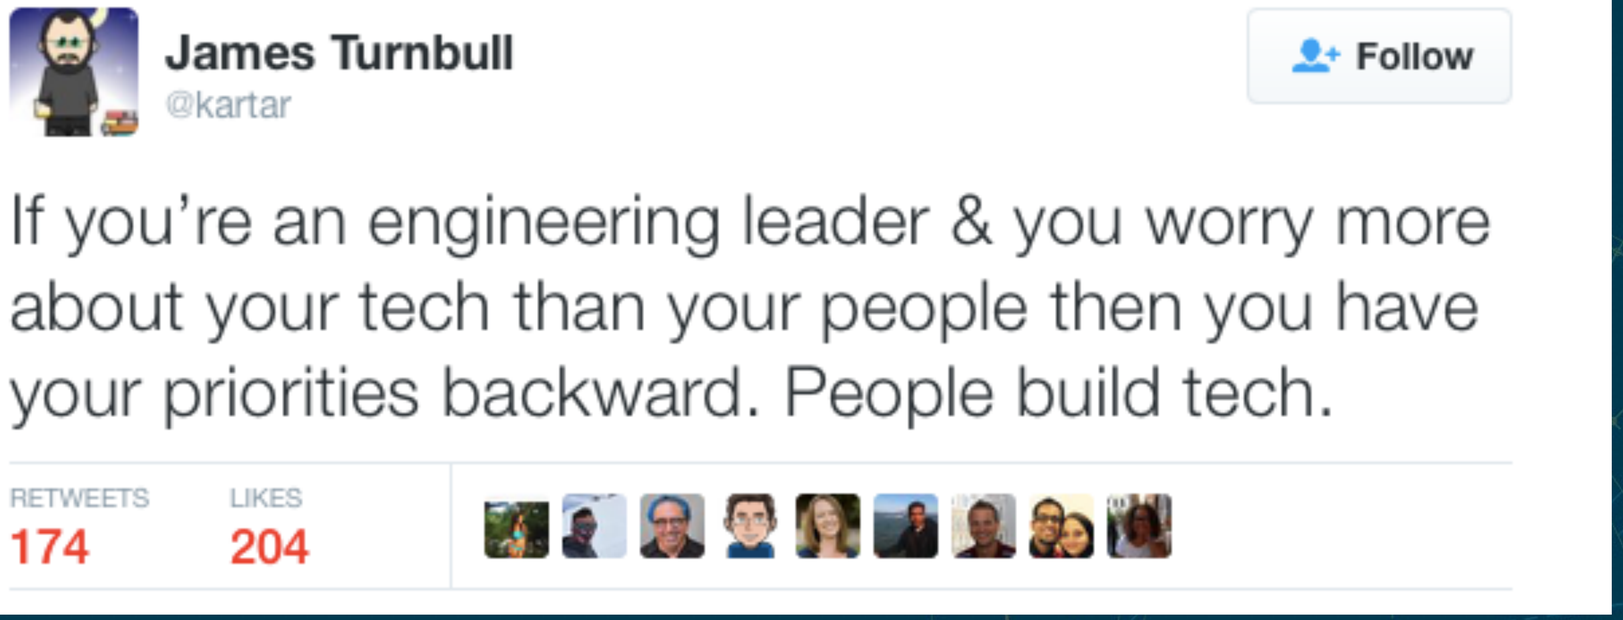
James Turnbull
@kartar
Follow
If you're an engineering leader & you worry more about your tech than your people then you have your priorities backward. People build tech.
RETWEETS
174
LIKES
204