LHC on the Pacific Research Platform

Jeff Dost (UCSD)

What is PRP?

- Pacific Research Platform:
  - 100 gbit network extending from Southern California to Washington
  - Interconnects Science DMZ between institutions

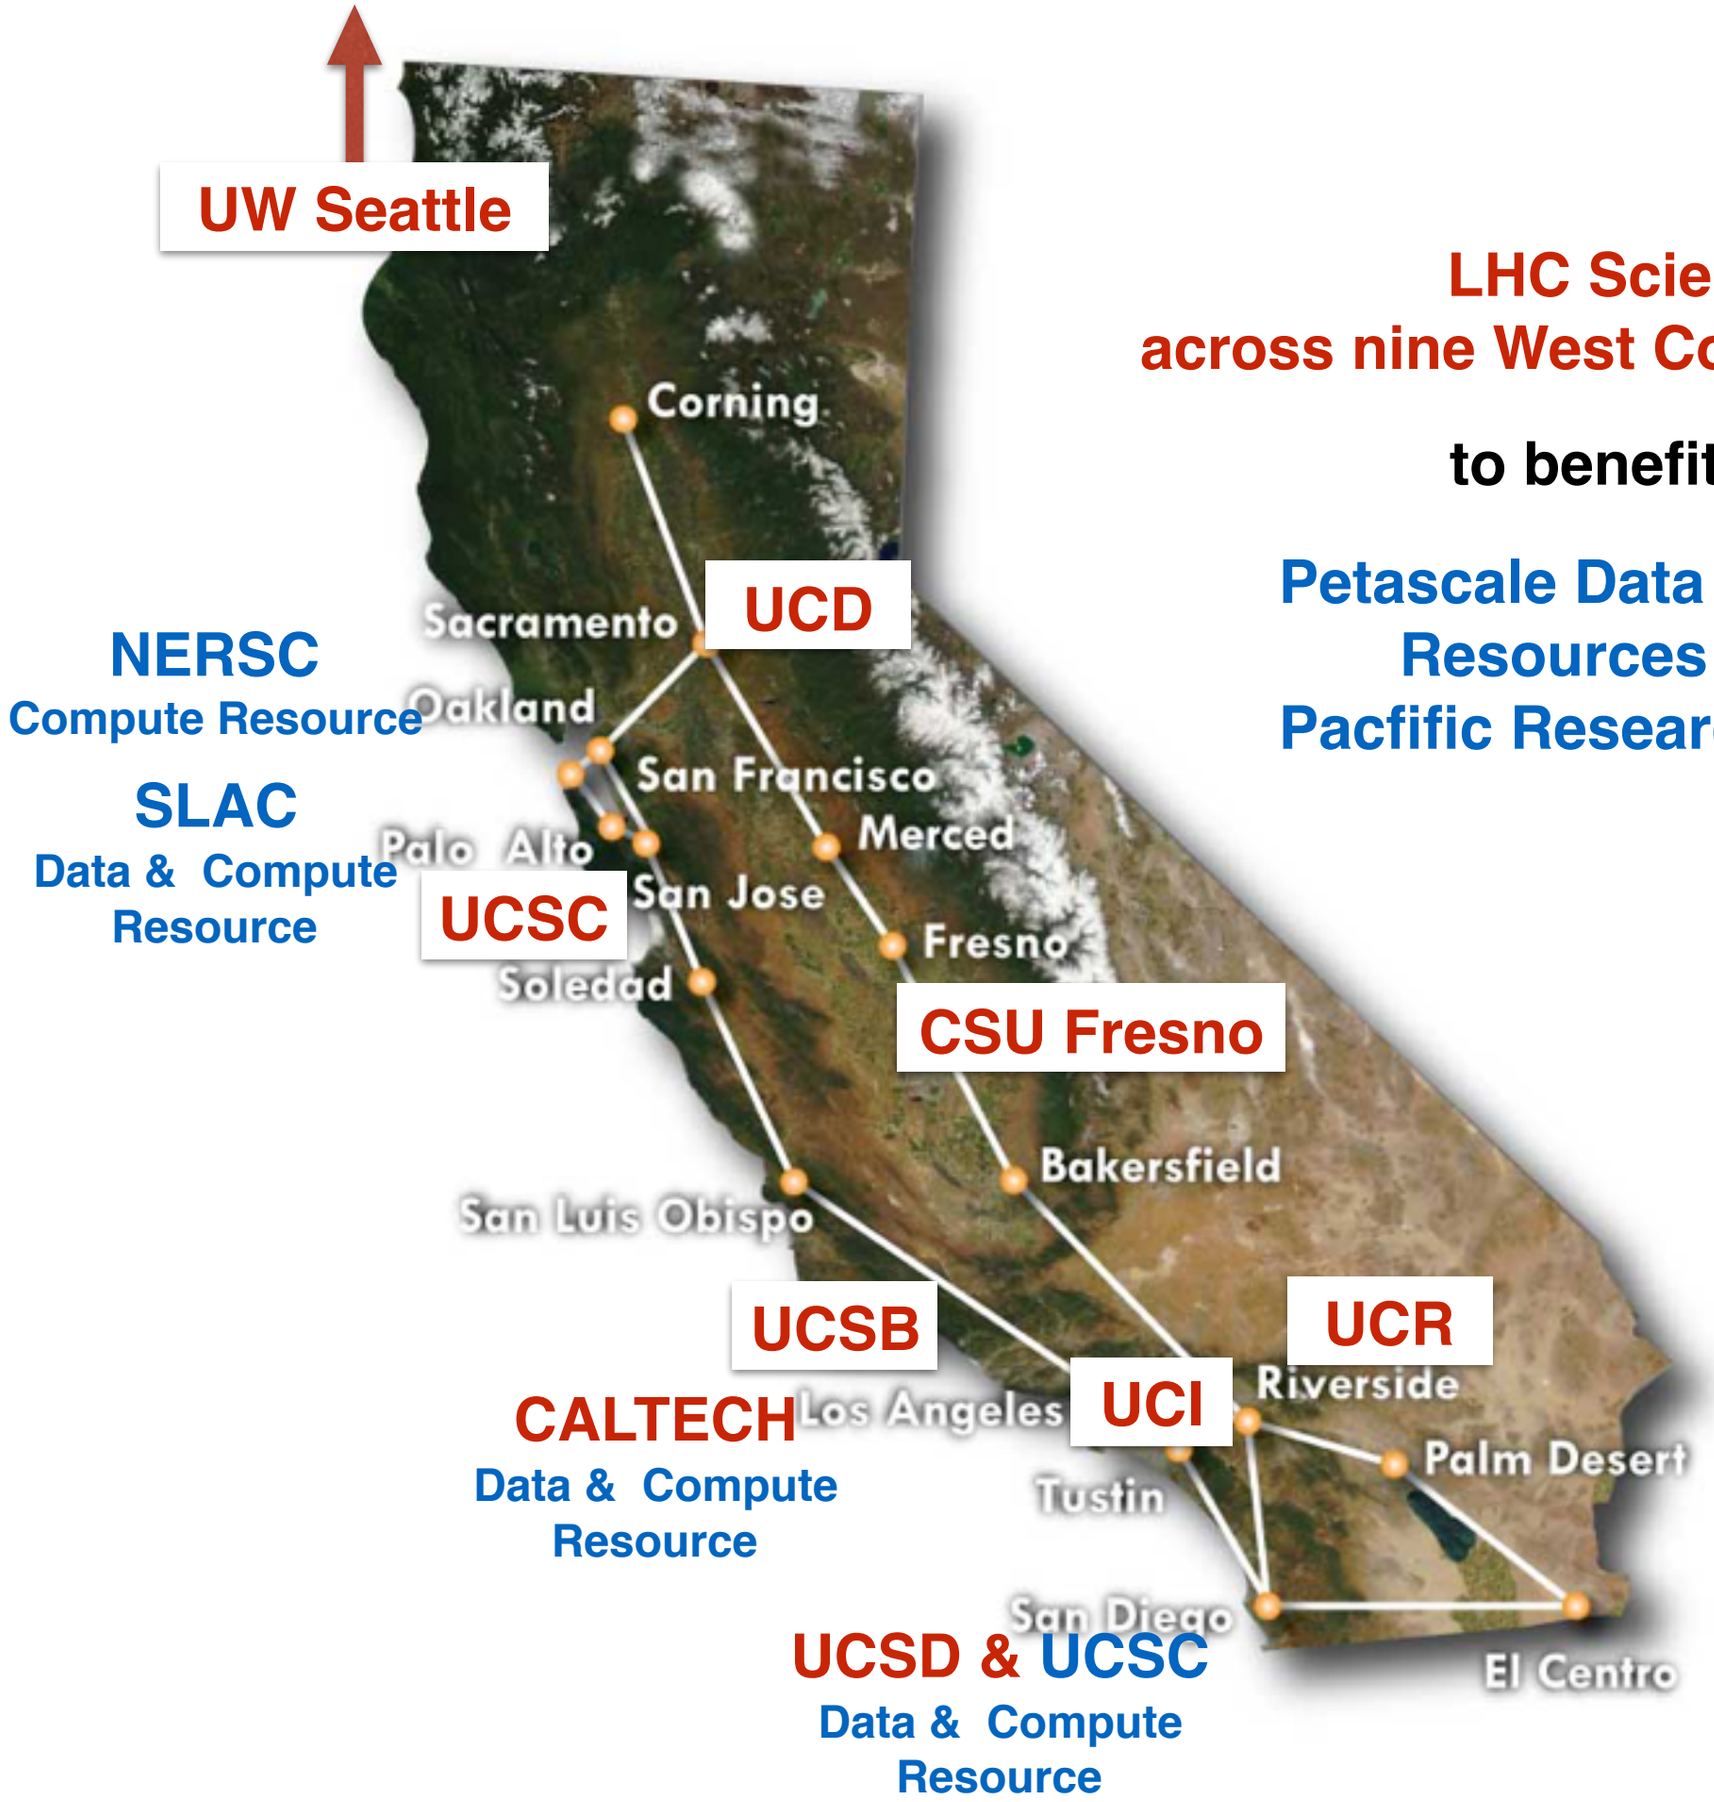

LHC Scientists across nine West Coast Universities

to benefit from

Petascale Data & Compute Resources across Pacfific Research Platform

What is LHC @ UC?

- Pilot project that unifies University of California Tier 3 campus compute infrastructures on top of PRP network
- Provides each UC the ability to:
  - Utilize external compute resources
    - Provided by GlideinWMS HTCondor pool
  - Access data from anywhere
    - Provided by XRootD
- Integrates seamlessly with local compute and storage resources

What is LHC @ UC?

- Participating UCs:
  - ATLAS
  - CMS
- All share resources
  - Including San Diego Super Computing Center

External Compute Resources

- Currently:
  - Each UC
  - 50k core Comet cluster at SDSC
    - See Edgar's talk
- Eventually:
  - Any other non-UC participating PRP site
  - Any OSG site beyond PRP
  - Other NSF XSEDE and DOE super computing centers
  - Cloud resources

Data Access

- Built on top of XRootD
- Jobs don’t need to run where the data is located
- Local UC and external compute resources both cache remote data accesses
- Arbitrary data from local UC can be exported and made available to all compute resources

Hardware shipped to UCs

(aka the “brick”)

Hardware:

- 40 cores
- 12 x 4TB data disks
- 128 GB ram
- 2 x 10 gbit network interface

Software:

- Full HTCondor pool
- XRootD server, redirector, and proxy cache
- cvmfs w/ optional Squid

The brick is effectively a site in a box

Traditional T3

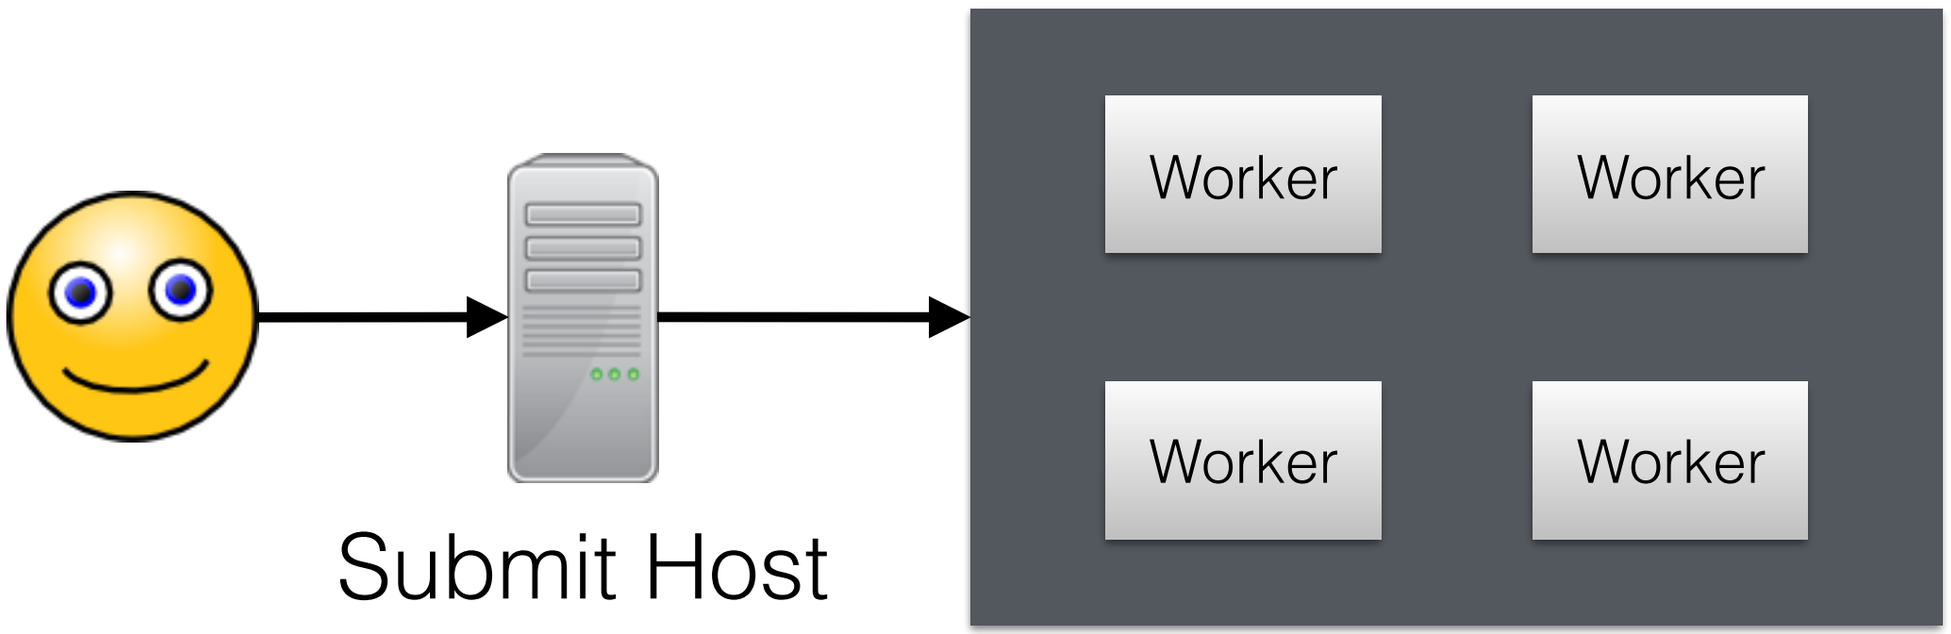

Submit Host

Local Batch

Uber Tier 3

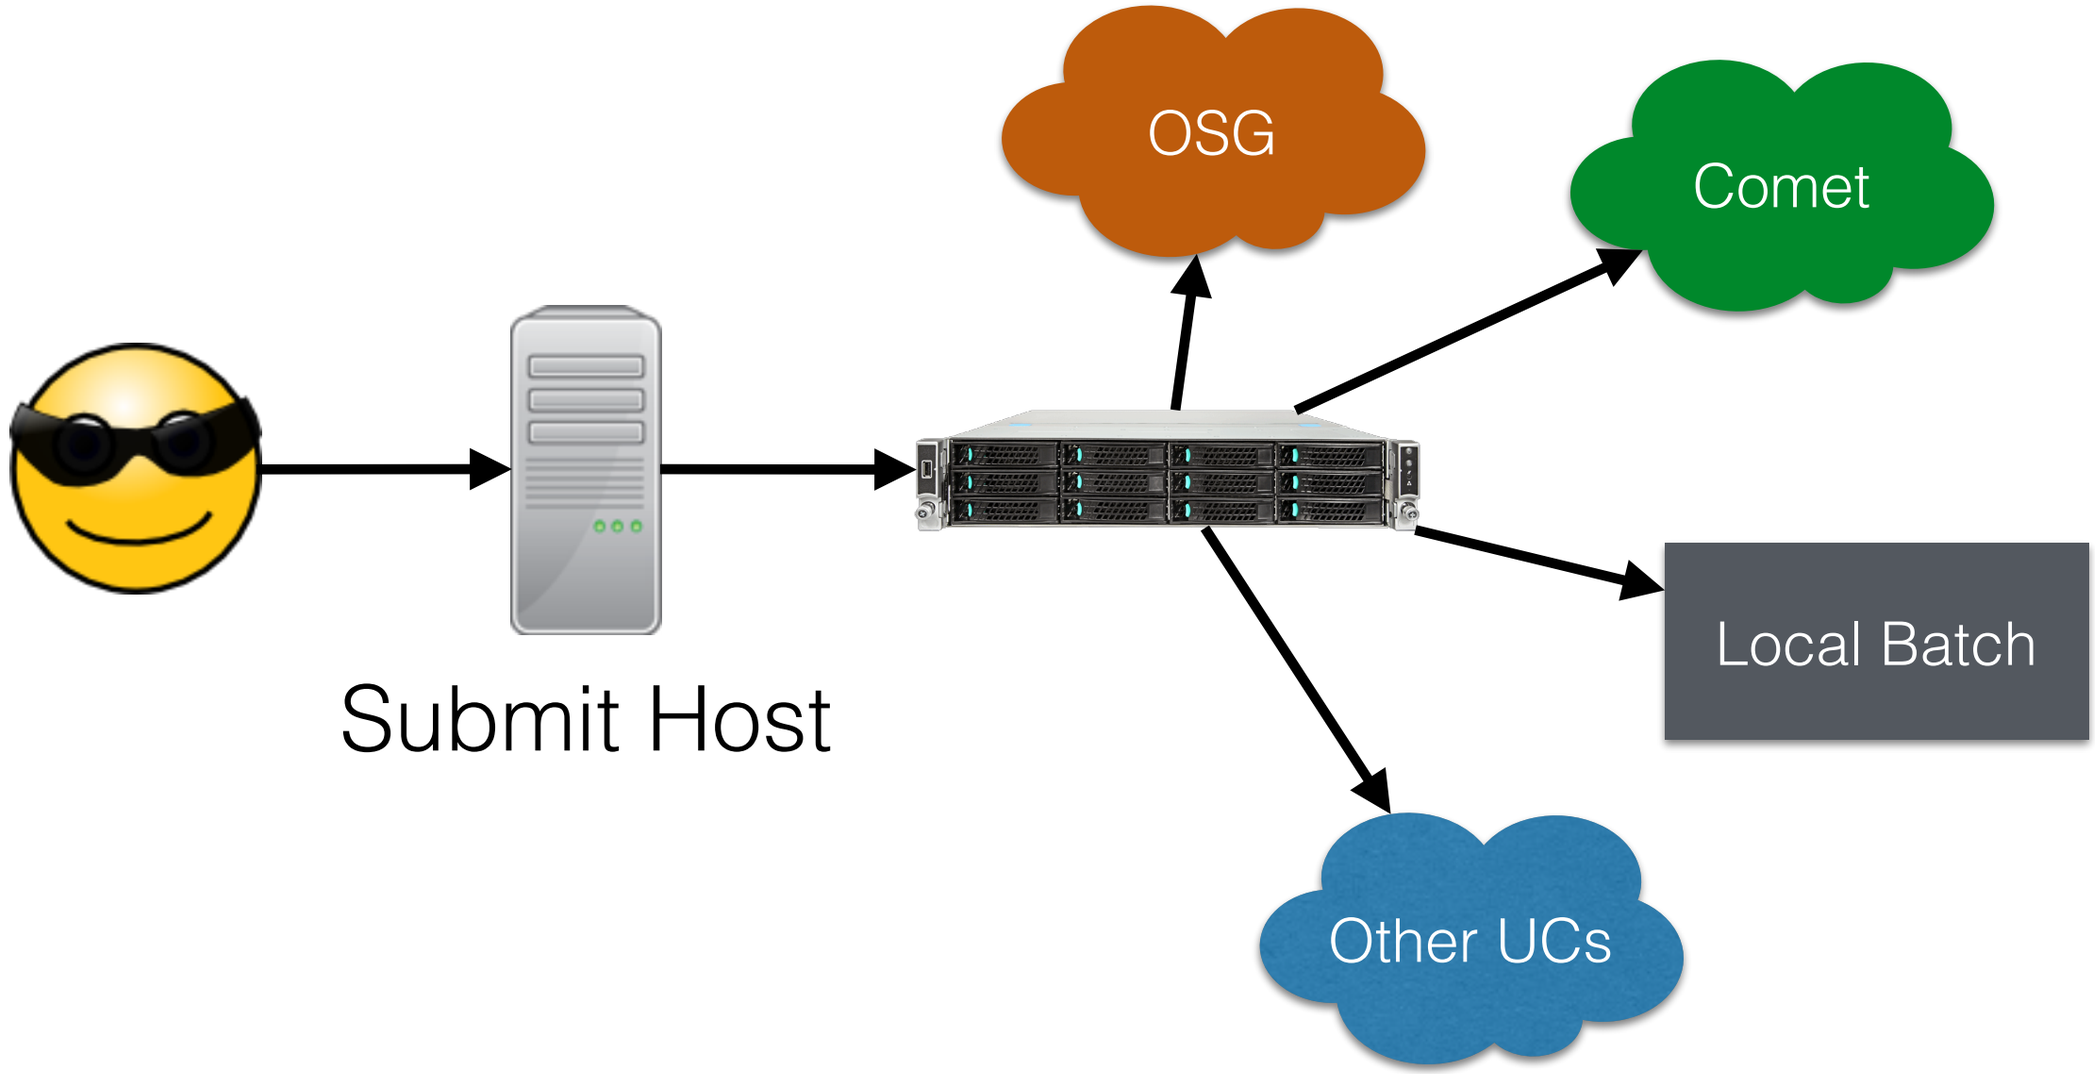

Milestones

- Jul 2015 — set up UCLHC XRD Federation
- Aug 2015 — deployed* first brick at UCI
- Sep 2015 — glidein pool configured to run at 5 out of 6 UCs (Santa Cruz not ready to support incoming grid jobs)
- Dec 2015 — deployed** UCR brick
- Mar 2016 — deployed** UCSC and UCD bricks

* **Fully operational, being actively used**
* **brick host is up and connected but still working out configuration details**

Technical Challenges

- T3 heterogeneity:
  - Different Network topologies — Using 1 network card connected to DMZ only, or 2, one also for private internal network?
  - Workers behind a NAT?
  - Different firewall requirements
  - Different batch systems (fortunately most are HTCondor)
  - Different storage backends
  - VO config differences (ATLAS vs CMS)

Technical Challenges

- T3 heterogeneity:
  - **Different Network topologies — Using 1 network card connected to DMZ only, or 2, one also for private internal network?**
  - **Workers behind a NAT?**
  - Different firewall requirements
  - Different batch systems (fortunately most are HTCondor)
  - Different storage backends
  - VO config differences (ATLAS vs CMS)

Work in Progress

Dealing with Diversity

- We use a combination of Puppet + Foreman and Hiera to manage brick configs
- Hiera allows us to set up tiered configurations:
  - General — Top level that all bricks share
  - VO specific — Shared config for all ATLAS or CMS bricks respectively
  - Host specific — Anything unique applying to a particular UC
- Heavy use of puppet parameters in combination with the Heira tiers significantly reduces the overhead of maintaining the diverse nature of the UC Tier 3's
- Working out the pending network configs should be just a matter of careful HTCondor and XRootD configuration

Todo

- Complete the network configuration for UCR, UCSC, and UCD to enable usage
- Coordinate with UCSB to deploy their brick
- Install cache node at Comet
- And then…

One Puppet to rule them all

UCSB

UCDAVIS
UNIVERSITY OF CALIFORNIA

UCI

UCR

UC San Diego

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

Conclusion

- LHC @ UC project utilizes the PRP network to enhance the T3s at each site by providing:
  - A unified way to submit locally and compute globally
  - The ability to decouple data placement from where the jobs run
- The central management of the services by dedicated admins at UCSD allows the local UC users to worry less about infrastructure maintenance and focus more on getting science done