



Jet Propulsion Laboratory
California Institute of Technology

Unifying HPC

Bryan Bales, HPC Systems Architect, IT Engineering and Project Management, OCIO
Jet Propulsion Laboratory, California Institute of Technology

March 6, 2017



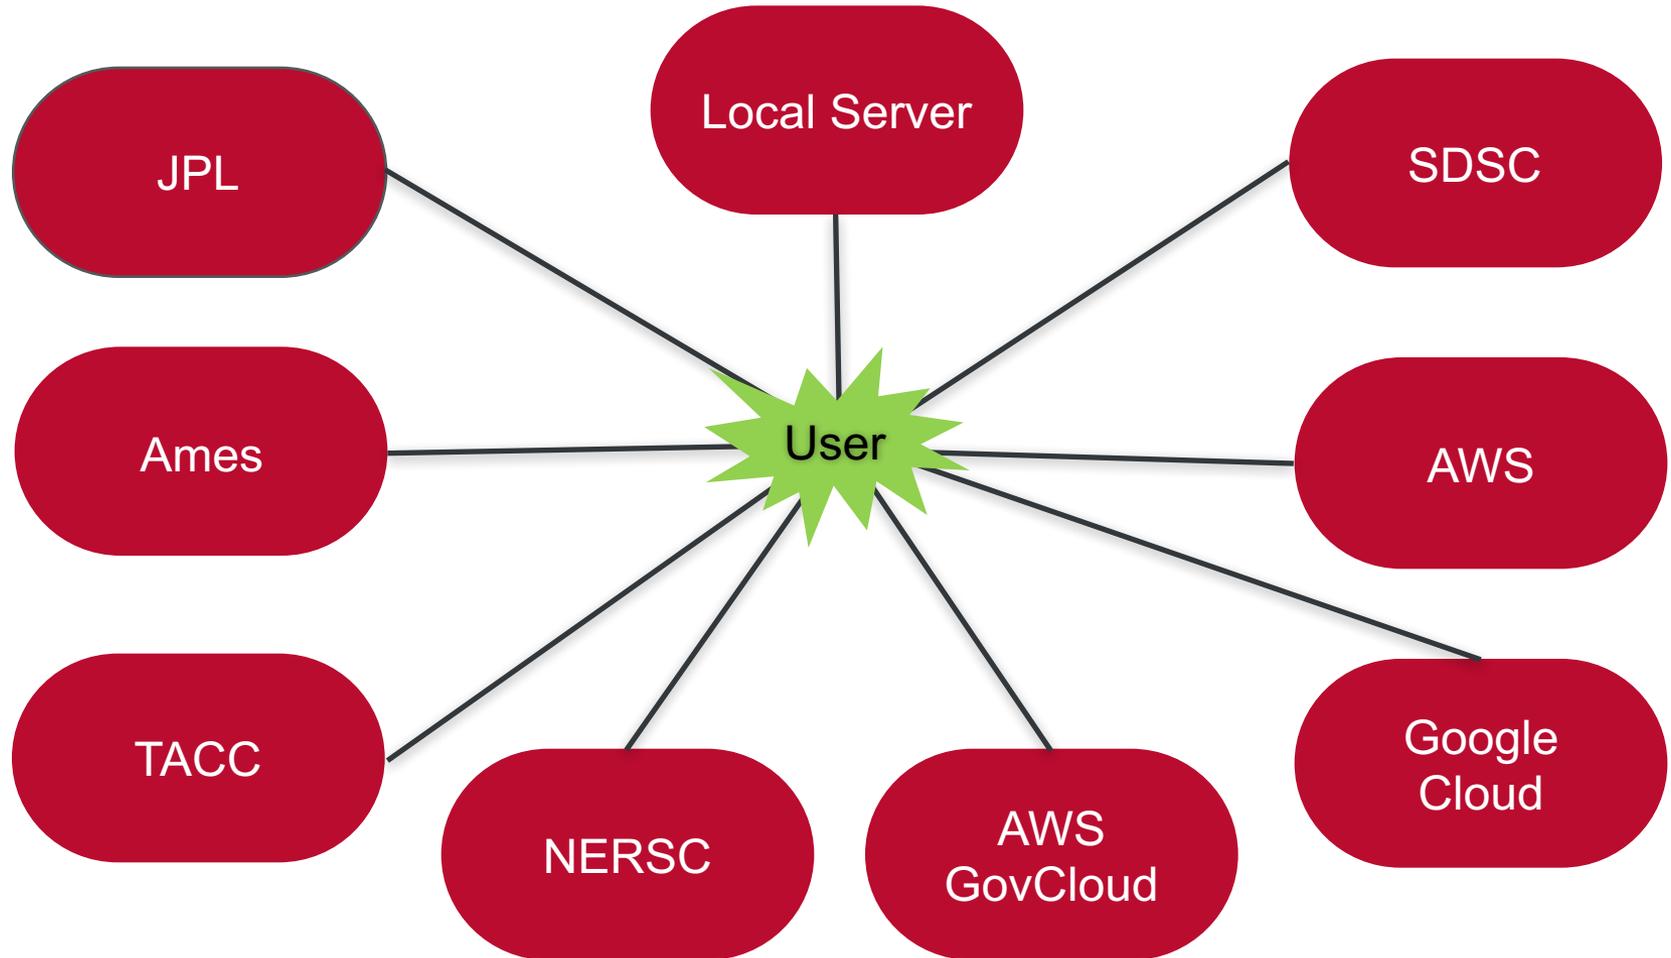
Unifying HPC

Problem Definition

- Job 1 – The Missions
 - All science and research computing is secondary to mission critical needs
 - Luckily, missions typically have few critical periods spread over many years
- Job 2 – The Science
 - Scientific users must either arrange for external resources or live with computing around mission needs

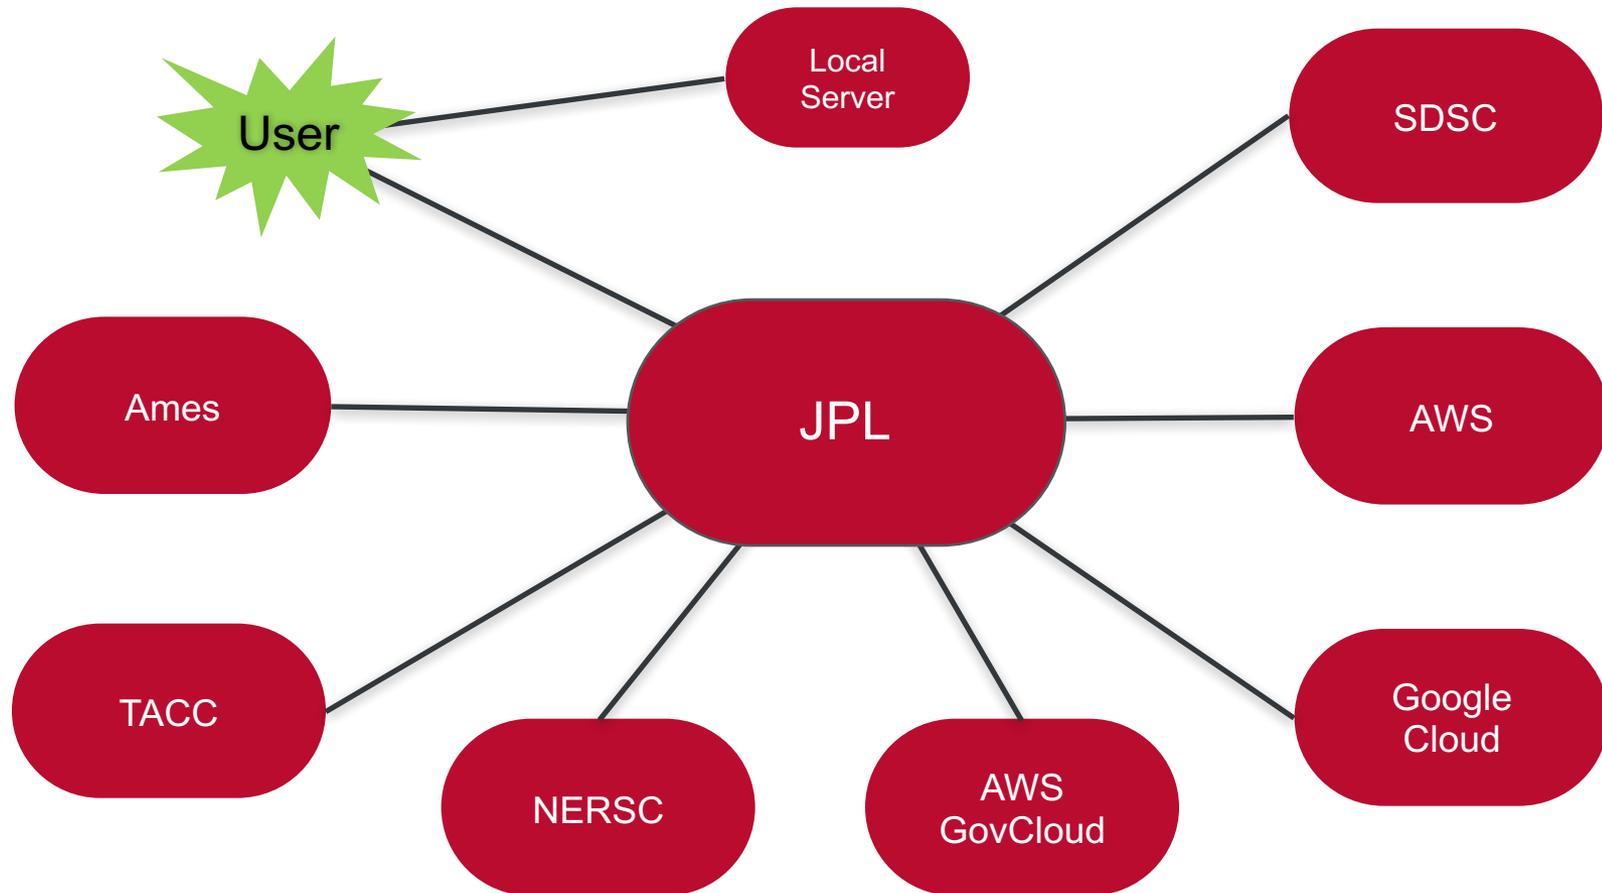
Unifying HPC

Where users compute



Unifying HPC

How we'd like users compute



Unifying HPC

Implementation

- Update JPL Infrastructure
 - Network
 - 100 Gbps – minimum
 - Storage
 - Unified Architecture for shared storage
 - NFS v4.x with ACLs for security
 - High Speed, Low Latency Scratch
 - HPC clusters
 - Expansible management, analytics
 - Compute PODs

Unifying HPC

Implementation

- **Extend JPL Infrastructure**
 - **Cloud**
 - Add capability to dynamically deploy compute PODs in cloud
 - Explore service like Scalr, Rightscale
 - **Job Submission, Execution and Data Movement**
 - Open Science Grid, Globus, HTCondor
 - Extend for heavy MPI use cases

Unifying HPC

User Acceptance

- No mandate for users
 - Users have choice whether to use JPL computing or go external
 - Ease of use is key to acceptance
 - Must be able to quantify advantages with better metrics of system utilization and throughput
- Keep user community involved in HTC & HPC efforts



Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov