


Platform

LSF Universus



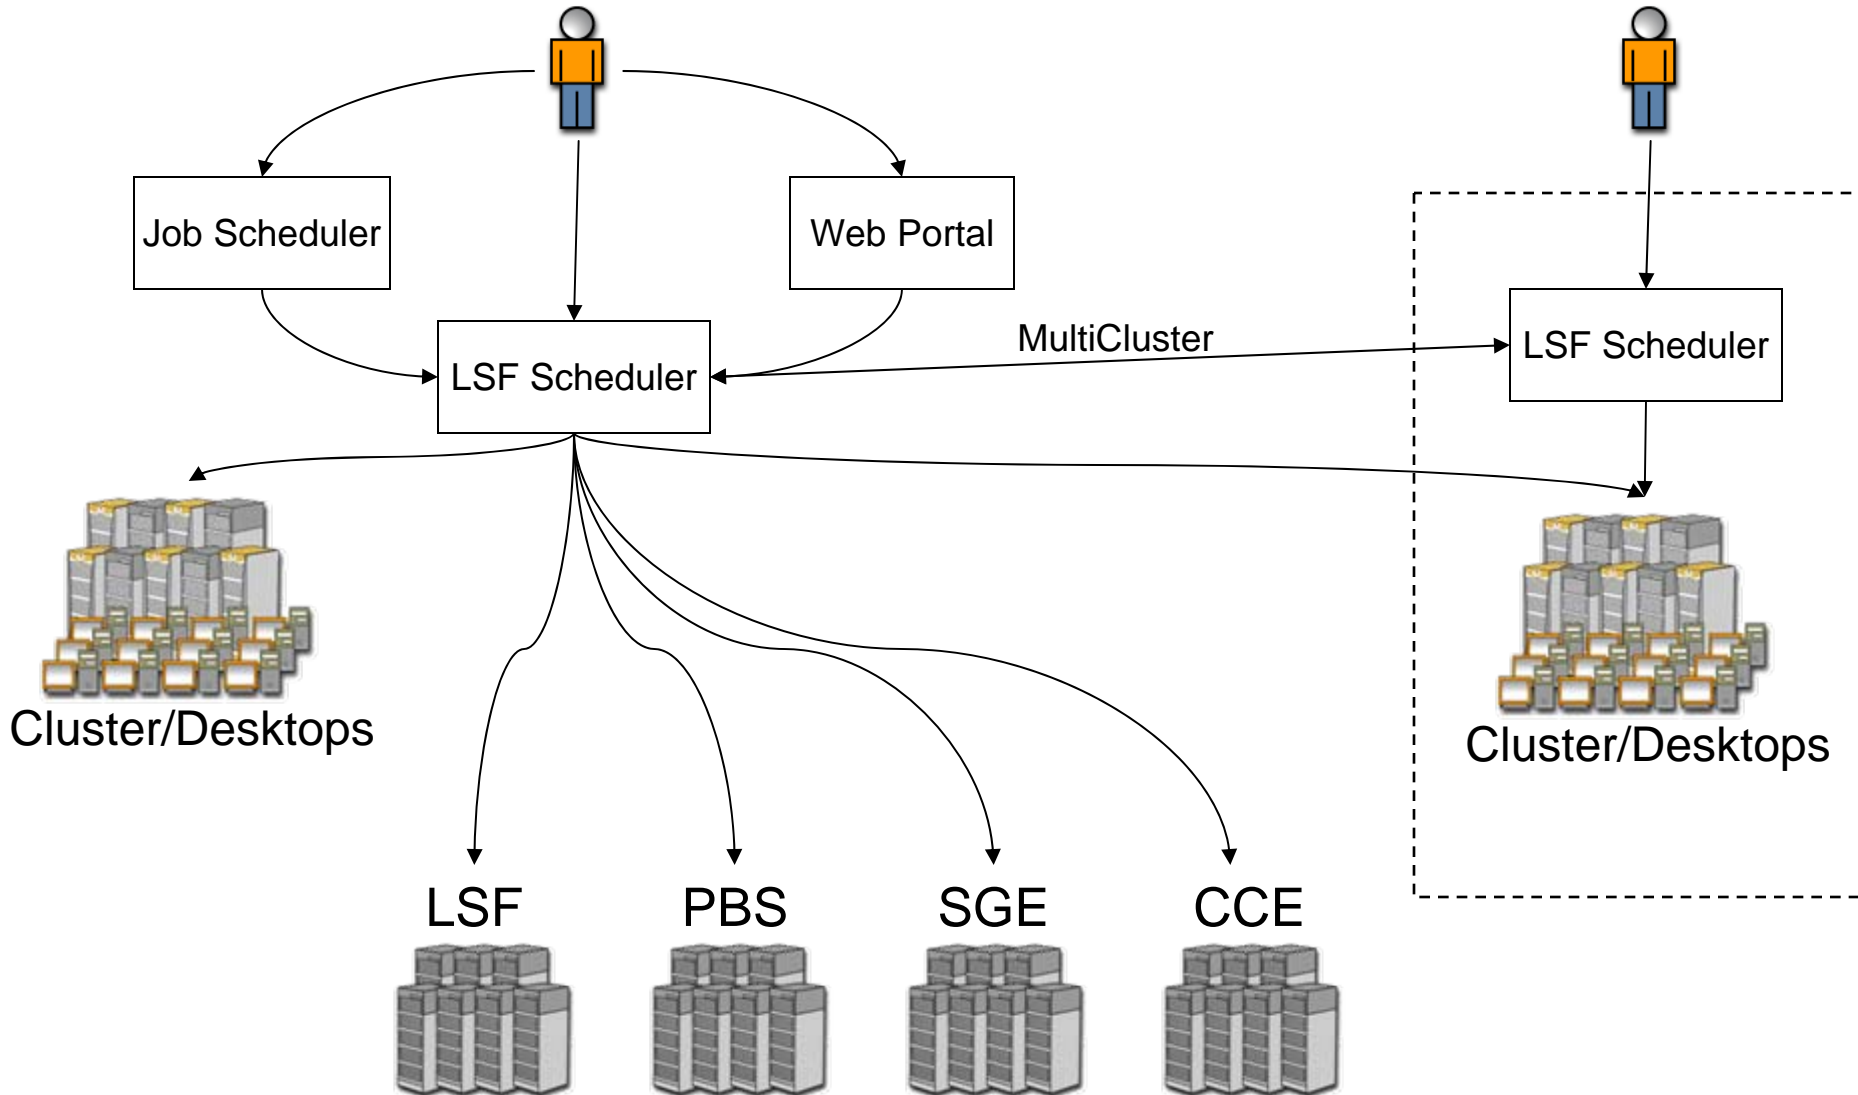
By Robert Stober
Systems Engineer
Platform Computing, Inc.

- ❑ Overview
- ❑ How it Works
- ❑ Security
- ❑ Summary
- ❑ Q & A



Overview

- ❑ LSF Universus is an extension of LSF that provides a secure, transparent, one-way interface from an LSF cluster to any foreign cluster.
- ❑ A foreign cluster is a local or remote cluster managed by a non-LSF workload management system.
- ❑ Universus allows organizations to tie multiple clusters running various workload management systems together into a single logical cluster.
- ❑ Provides users with a single, secure interface to all the computing resources



Benefits

- ❑ Users can use all computing resources without having to log into them and using only LSF commands to submit, monitor and control jobs
- ❑ Centralized scheduler, but sites retain local control
- ❑ Local users can continue using local resources during and after implementation
- ❑ Low cost, fully supported solution
- ❑ As secure as you need it to be

Current Implementations

- ❑ Sandia National Laboratory
 - Used to link OpenPBS, PBS Pro, and LSF Clusters in New Mexico and California

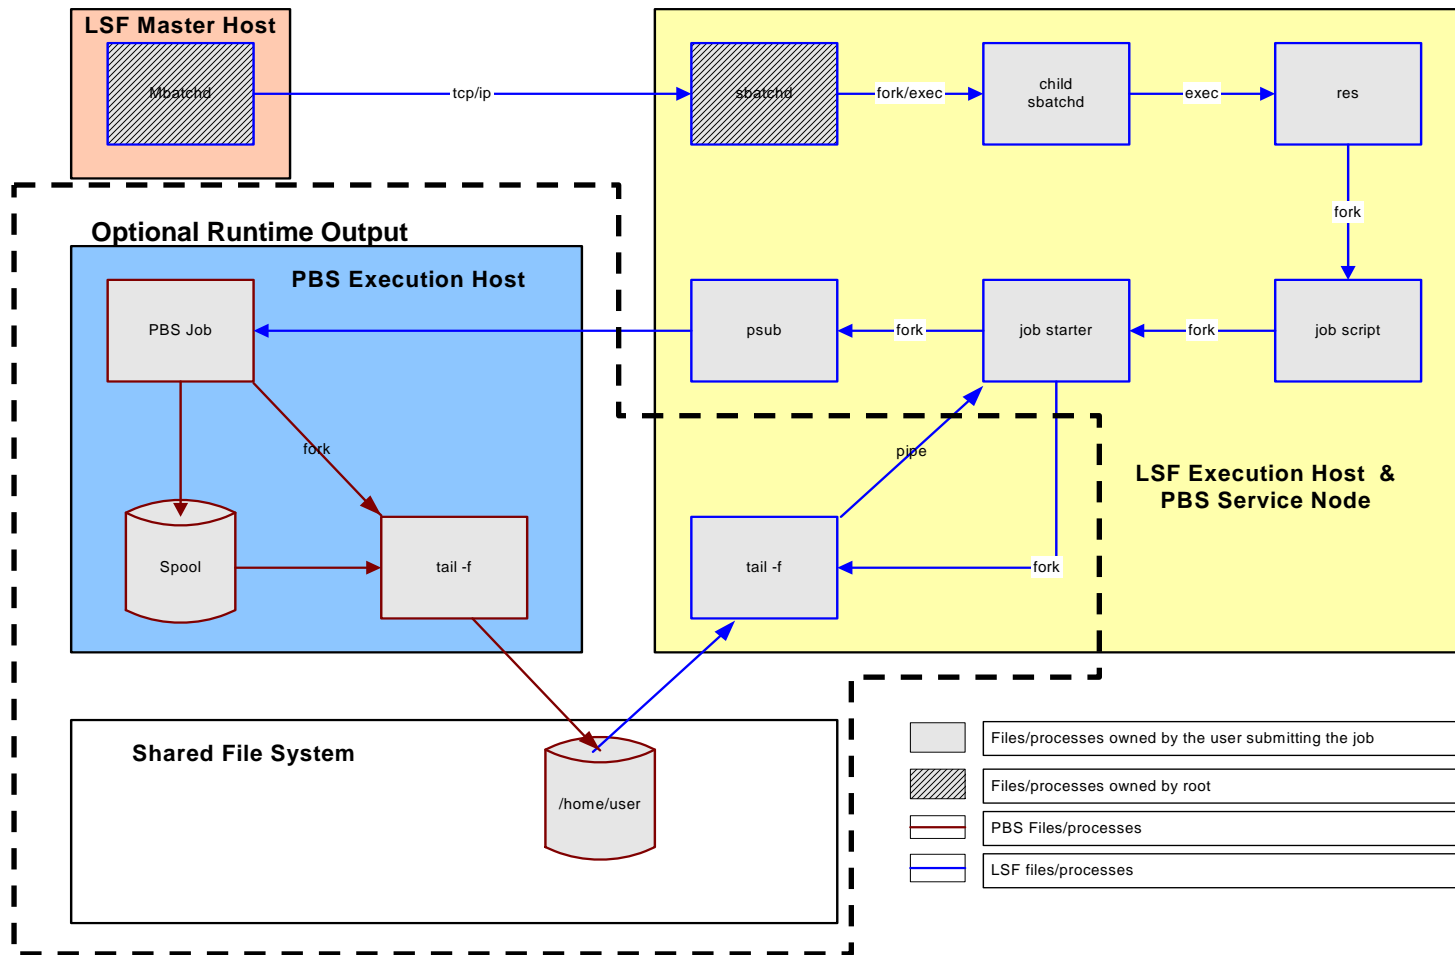
- ❑ Singapore National Grid (NG)
 - Used “to provide seamless access to NGPP resources”
 - Resources include LSF, PBS Pro, and N1GE clusters
 - Completed in Q3 2005

- ❑ Distributed European Infrastructure for Supercomputing Applications (DEISA)
 - Universus has been deployed on top of the DEISA native batch systems to enable users to easily access resources allocations on different remote platforms.

Requirements & Assumptions

- ❑ Needs to support any foreign workload management systems that provides a command line interface and/or supports embedded directives in job script
- ❑ Needs to support Kerberos authentication
- ❑ No shared file system between LSF master and execution host
- ❑ All file transfers and job content must be encrypted
- ❑ LSF daemons are installed on the head node of the remote cluster

How it Works



Uses LSF Extension Facilities

- `esub`
 - Records job submission options
 - Ensures that job submission record is copied to the execution host

- `Job Starter`
 - Reads in the job submission record and builds the command line to submit the job to the foreign workload management system
 - Monitors job status and reports the information back to LSF
 - Propagates kill, suspend, and resume signals to the remote job
 - Reads job stdout and stderr into LSF output stream in real time

- `lsrcp`
 - Uses scp to securely transfer files to and from execution hosts

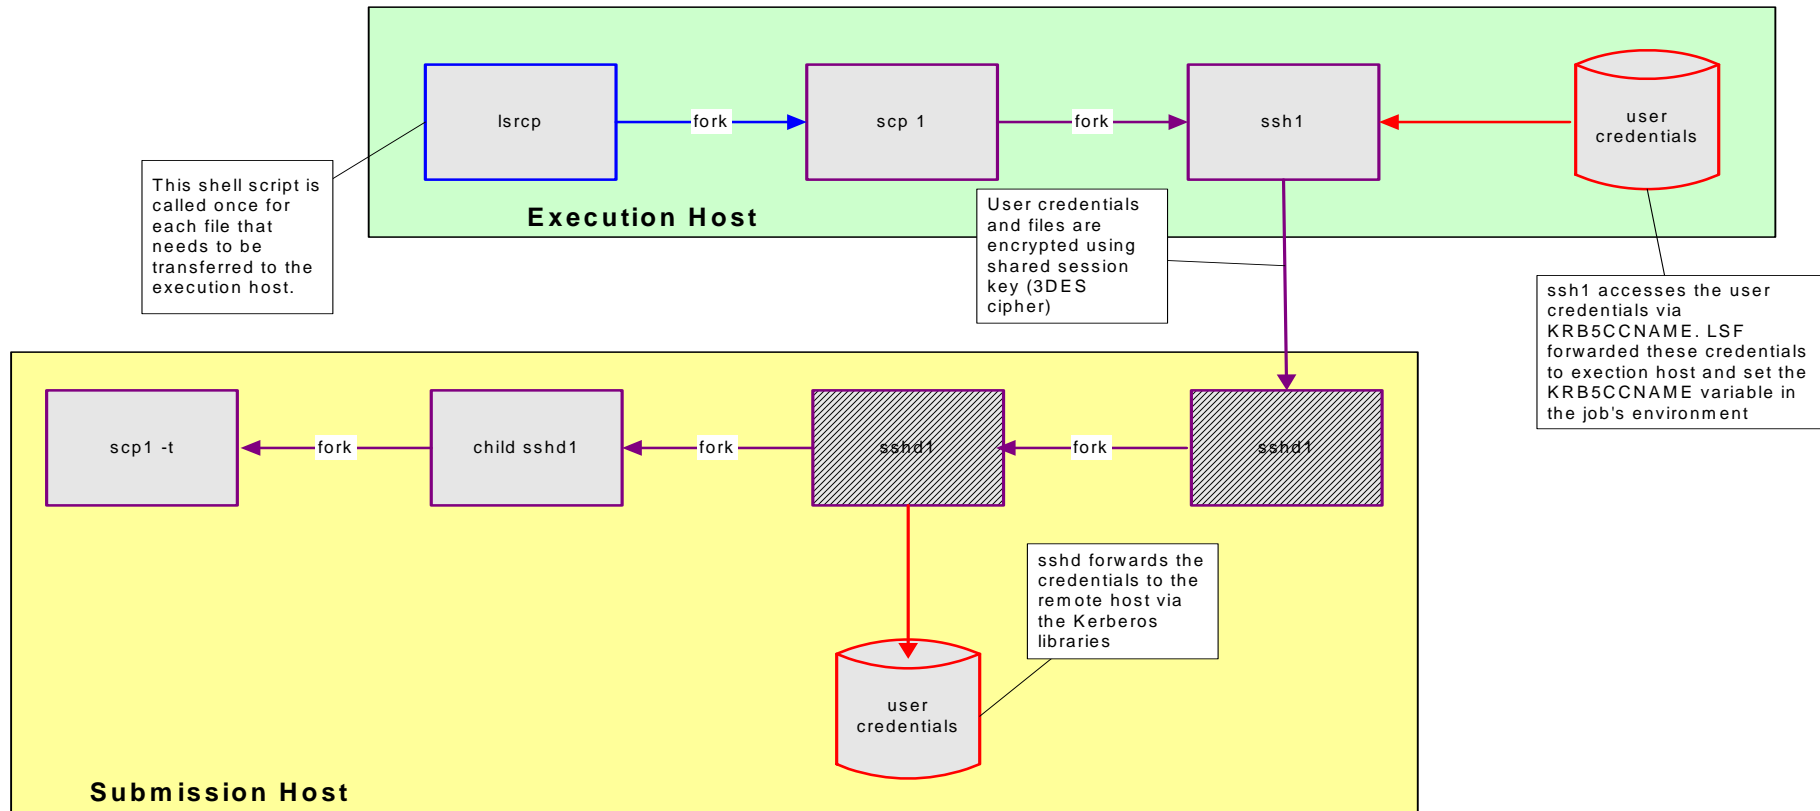
- `eauth`
 - Uses Kerberos to authenticate users, hosts, and services

LSF to PBS Option Map

bsub	Description	qsub
B	Sends mail upon job dispatch	"-m a"
r	makes a job re-runnable	"-r y"
c	[hour:]minute	"-lcput=[hour:]minute"
e	err_file	"-e \$HOME/.assets/assets.LSB_JOBID/spool/stderr"
ext[sched]	"external scheduler options"	"-W x=FLAGS:ADVRES:RESID"
F	file_limit	"-lfile=\${LSB_SUB_RLIMIT_FSIZE}kb"
G	user_group	"-W group_list=user_group"
J	job_name	"-N job_name"
k	"checkpoint_dir[period]"	"-c c=checkpoint_period"
L	login_shell	"-S login_shell"
M	mem_limit	"-lpmem=\${LSB_SUB_RLIMIT_RSS}kb"
n	min_processors[,max]	"-lnodes=(processors/span)"
o	out_file	"-o \$HOME/.assets/assets.LSB_JOBID/spool/stdout"
P	project_name	"-A project_name"
R	Resource requirement	ppn=span or ppn=1 if span[hosts=1]
u	mail_user	"-M user"
v	swap_limit	"-lvmem=\${LSB_SUB_RLIMIT_SWAP}kb"

- ❑ LSF Universus provides for secure file transfers by replacing the standard LSF lsrcp program with a wrapper script that calls scp
- ❑ Kerberos is supported. Just use the Kerberized version of LSF and ssh
- ❑ Setting the LSF_RSH=ssh causes LSF to use ssh instead of rsh
- ❑ While the LSF Kerberos integration provides LSF daemon authentication, daemon communication is not encrypted

Secure File Transfers



	Files/processes owned by the user submitting the job
	Files/processes owned by root
	ssh Files/processes
	Kerberos files/processes
	LSF files/processes

Summary

LSF Universus is:

- ❑ A proven, fully-supported solution built using the standard LSF extension mechanisms
- ❑ A meta-scheduler that sits on top of, and schedules to, a variety and extensible set of workload management systems
- ❑ Cost-effective because it only requires LSF to be installed on the head-node of each resource
- ❑ Easy to implement since sites can retain control of their resources, and since the resources can be used during roll-out

Summary

LSF Universus is:

- ❑ Open and extensible. The executables that comprise the solution are scripts that can be understood and modified to suit your needs
- ❑ As secure as you need it to be. Standard, ssh, and Kerberos authentication are available, or you could extend the solution to support PKI
- ❑ Practical. Universus collects job resource information that be used for metering and accounting.



Platform LSF



Platform™

