



Reducing Faulty Jobs by Job Submission Verifier in Grid Engine



Misha Ahmadian[†], Eric Rees[†], Yu Zhuang[†], Yong Chen[†]

[†]Department of Computer Science, [‡]High Performance Computing Center
Texas Tech University

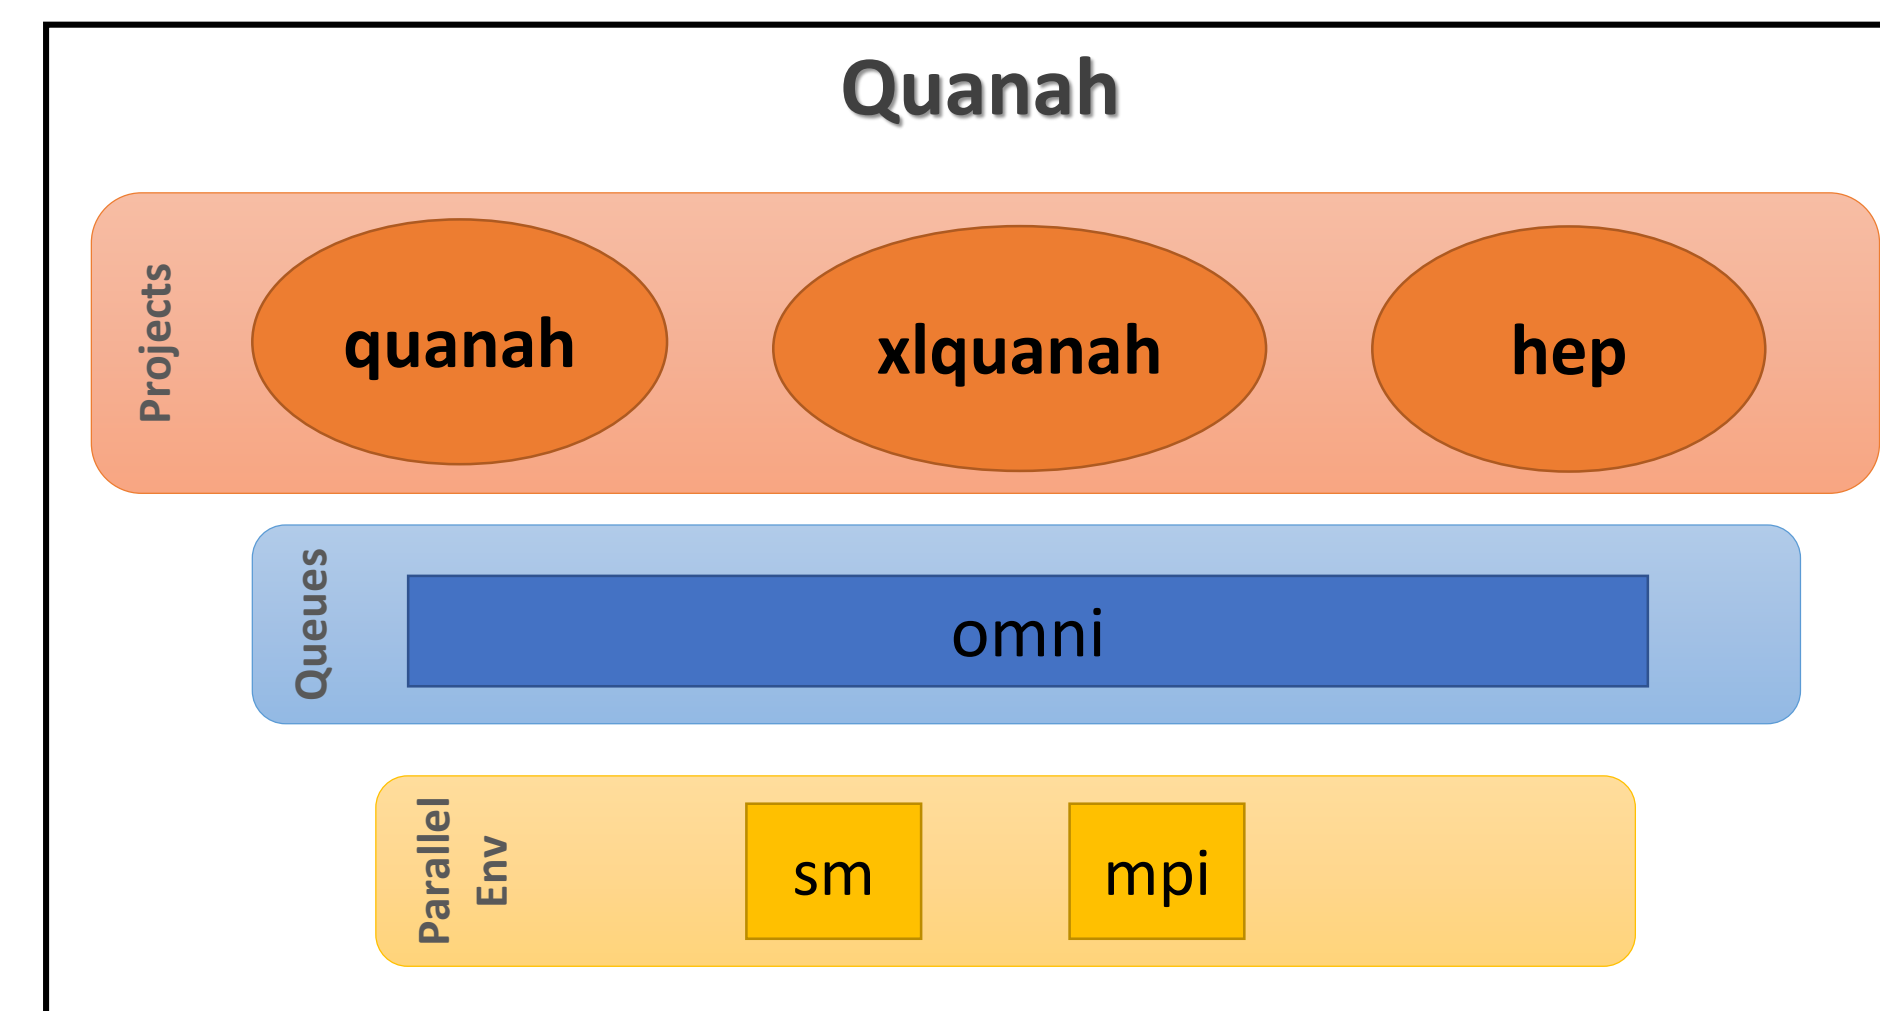
Abstract

- Grid Engine is a Distributed Resource Manager (DRM), that manages the resources of distributed systems (such as Grid, HPC, or Cloud systems).
- Grid Engine applies scheduling policies to allocate resources for jobs while maintaining optimal utilization of all resources.
- the complexity of Grid Engine's job submission commands and complicated resource management policies increases the number of faulty job submissions in data centers.
- We designed and implemented a Job Submission Verifier (JSV) to communicate with Univa Grid Engine (UGE)
- JSV verifies whether a submitted job should be *accepted* as is, or *modified* then accepted, or *rejected* due to improper requests for resources.
- Our evaluation shows a substantial positive impact on reducing the number of faulty jobs submitted to UGE on HPC clusters of Texas Tech University.

JSV Design for Quanah Cluster:

Resources and Policies on QUANAH Cluster:

- 467 nodes, 16,812 cores (36 core per node)
- 87.56 TB total RAM (192 GB per node)
- Intel OmniPath internal network (100 Gbps)



PE	Policies		
mpi	<ul style="list-style-type: none"> Must request for a multiple of 36 cores All MPI applications must define the 'mpi' PE Can request between 1 and 36 cores 		
sm	<ul style="list-style-type: none"> Slots are guaranteed to be in one node. 		
	quanah	xlquanah	hep
Max # of cores	16,182	144	720
default runtime	48 hours	72 hours	48 hours
Max runtime	48 hours	120 hours	∞
Allowed (PE)s	'sm', 'mpi'	'sm'	'sm', 'mpi'

JSV Design and Reference Implementation:

- Assign default runtime and PE values to jobs.
- Ensure that the (s_rt) is not greater than (h_rt).
- Confirm that (mpi) jobs request for CPU cores in multiples of a total number of cores per node.
- Confirm that (sm) jobs do not request for more than total number of CPU cores per node.
- Ensure that the requested amount of memory does not exceed the overall size of the memory across the requested nodes.
- Notify users of their job submission rejection with descriptive error messages.

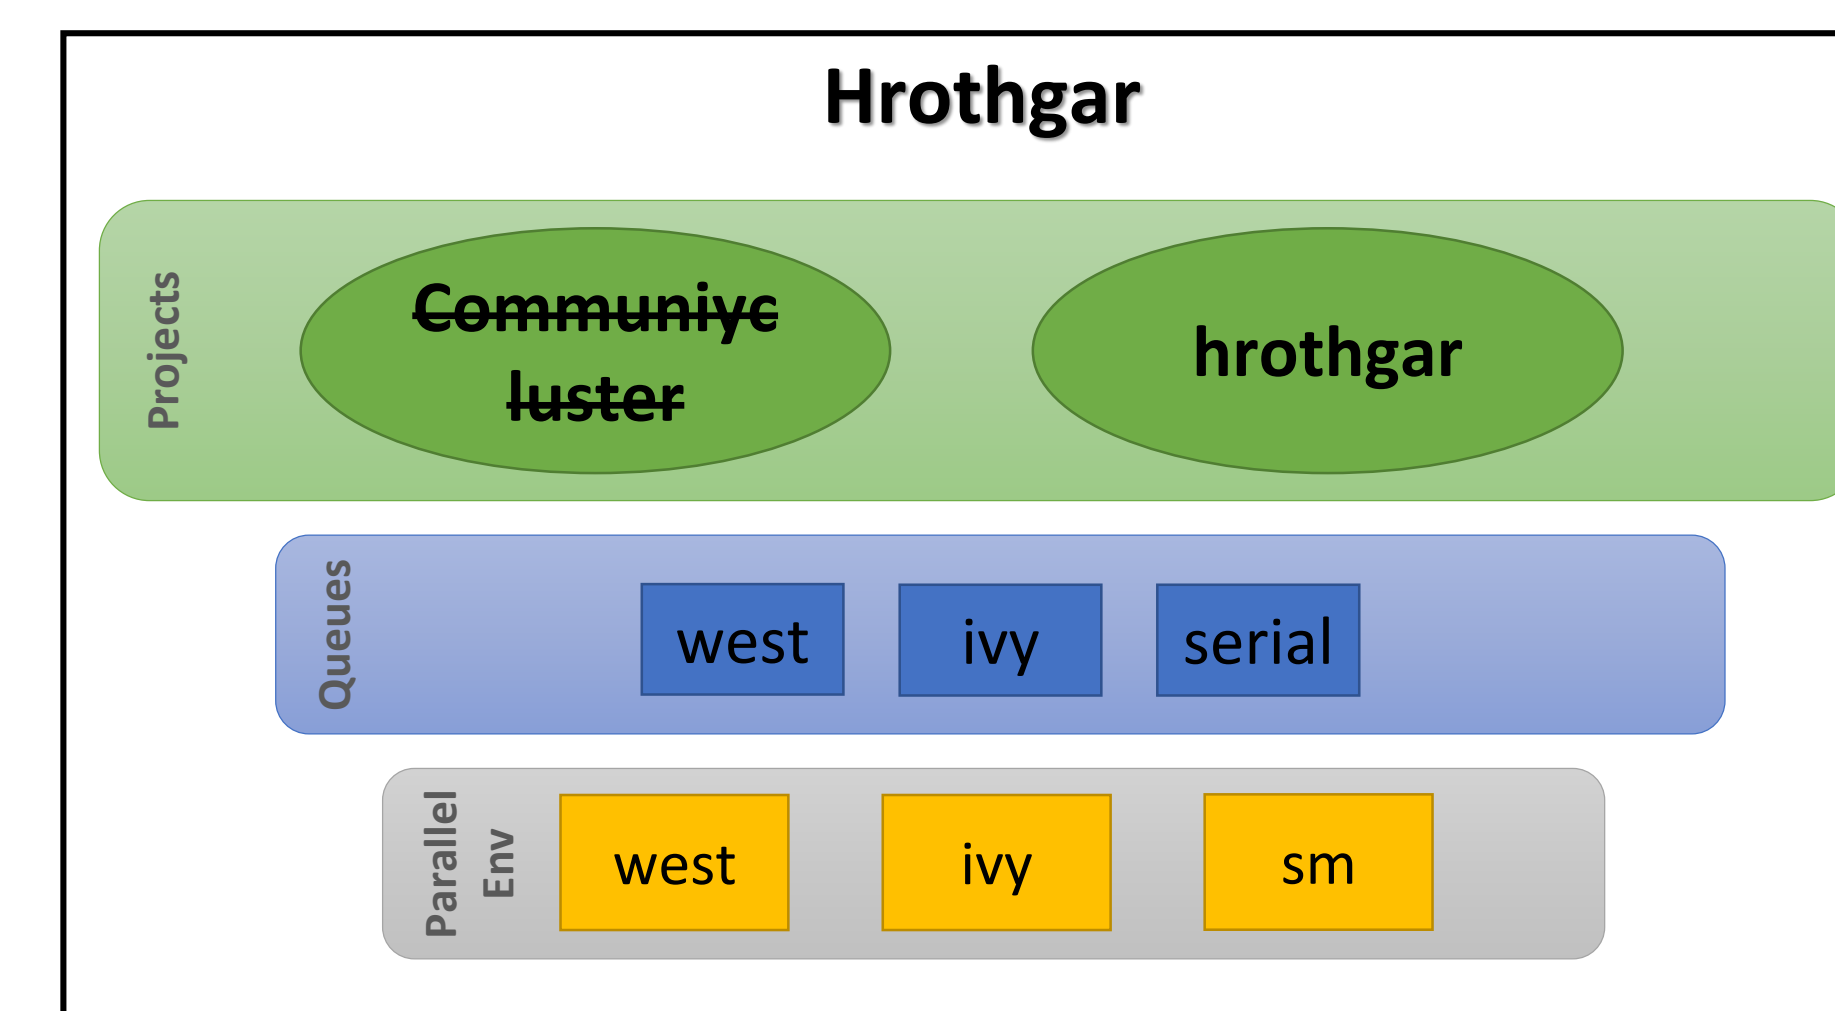
JSV analysis (09/10/2018 – 02/10/2019)

Submission Status	
ACCEPTED	70.6%
CORRECTED	29.0%
Set default Run Time and Memory Size	66.9%
Set default Memory Size only	30.7%
Set default Run Time, Memory, and PE	1.3%
Set default Run Time, Memory, and Reservation	0.9%
Enable Reservation Only	0.2%
REJECTED	0.3%
Non-Existing Project Name	86.5%
Non-Existing Parallel Environment (PE)	6.1%
Incorrect Number of Requested Slots *	4.7%
Incorrect Run Time Format	1.7%
Memory Size out of Bound *	0.9%
Total	100%

JSV Design for Hrothgar Cluster:

Resources and Policies on HROTHGAR Cluster:

- Hrothgar (West):** 563 nodes, 6,756 cores (12 cores/node), 13.19 TB total RAM (24 GB/node), and DDR Infiniband.
- Hrothgar (Ivy):** 96 nodes, 1,920 cores (20 cores/node), 6.14TB Total RAM (64 GB/node), and QDR Infiniband.
- Hrothgar (Serial):** Identical hardware as Hrothgar West. No Infiniband fabric.



PE	Policies		
west	Must request for a multiple of 12 cores		
ivy	Must request for a multiple of 20 cores		
sm	Can request between 1 to 12 cores.		
	west	ivy	serial
Runtime limit	48 hours	48 hours	120 hours
Allowed PE	'west'	'ivy'	'sm'
# cores per node	12	20	12
Max Memory size	24GB	64GB	48GB

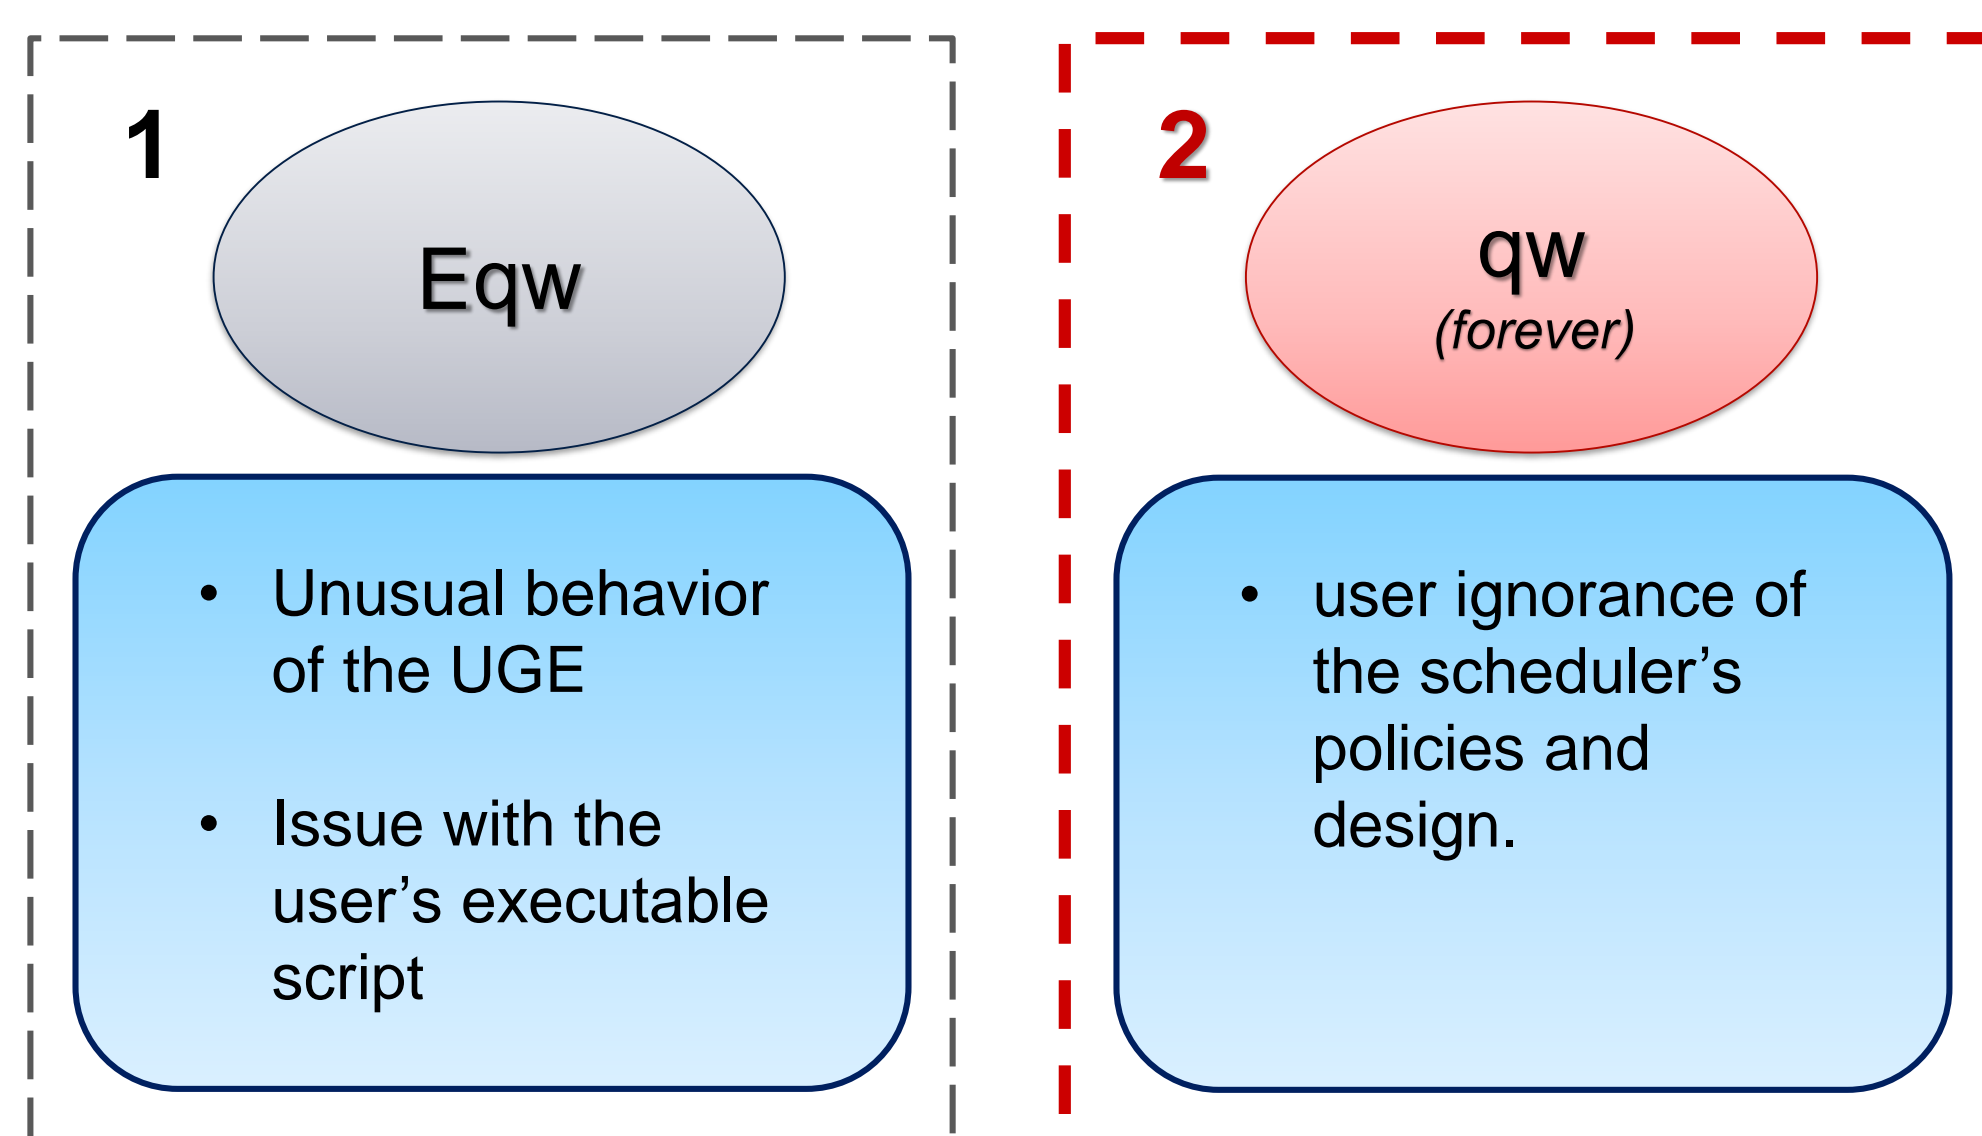
JSV Design and Reference Implementation:

- Only verifies the job submissions without correcting them.
- JSV works only when 'hrothgar' project is
- If job submissions requests for either 'west', 'ivy', or 'serial' queue, then the PE should be defined as 'west', 'ivy', or 'sm'.
- Ensure that the number of requested CPU cores is multiple of 12 for 'west' PE and 20 for 'ivy' PE.
- Confirme that (sm) jobs do not request for more than total number of 12 cores per node.
- Ensure that the requested amount of memory does not exceed the 24GB on 'west', 64GB on 'ivy' and 48GB on 'sm' queue.

JSV analysis (09/10/2018 – 02/10/2019)

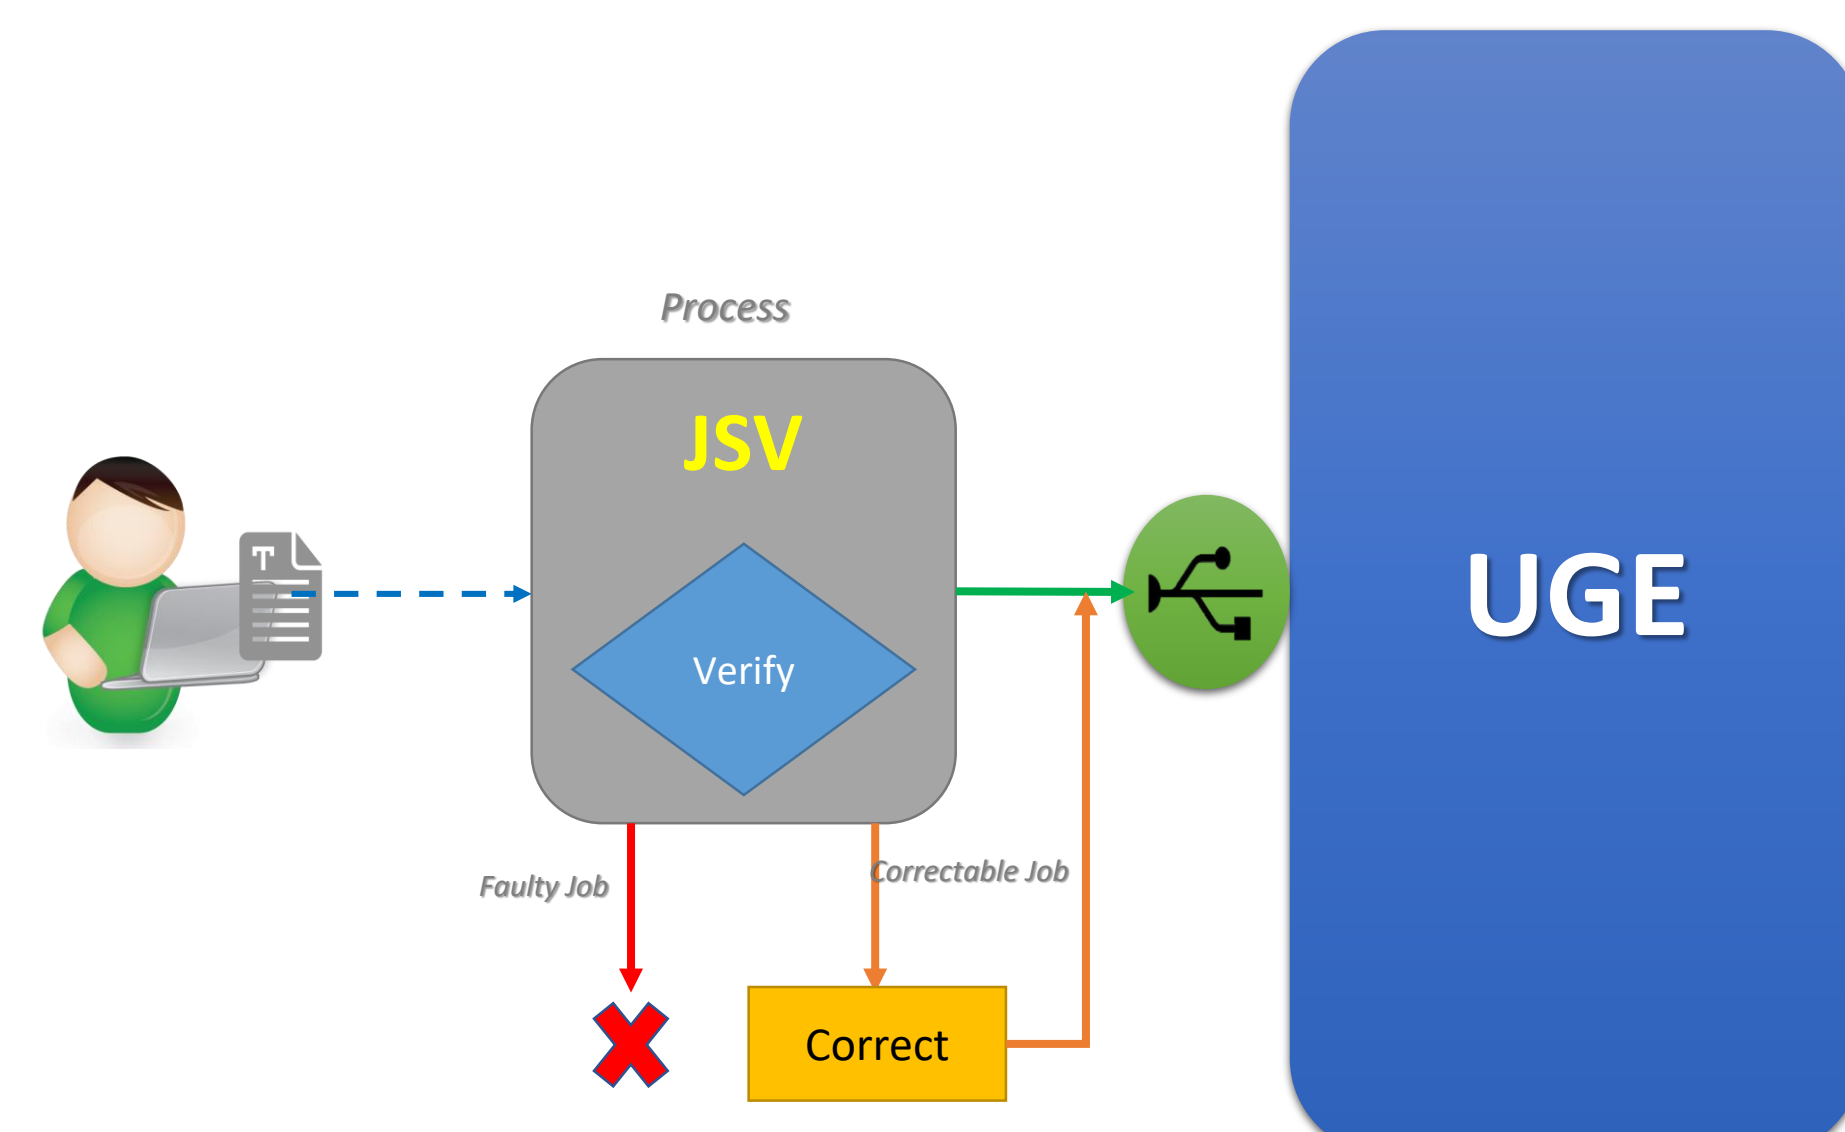
Submission Status	
ACCEPTED	99.5%
REJECTED	0.5%
Incorrect Project Name	24.0%
Non-Existing Queue Name	19.6%
Incorrect Parallel Environment *	19.6%
Undefined Parallel Environment (PE)	10.2%
Incorrect Queue Name *	9.1%
Incorrect Number of Requested Slots *	9.1%
Non-Existing Parallel Environment (PE)	5.1%
Undefined Queue Name	2.2%
Memory Size Out of Bound *	1.1%
Total	100%

Faulty Job Submissions




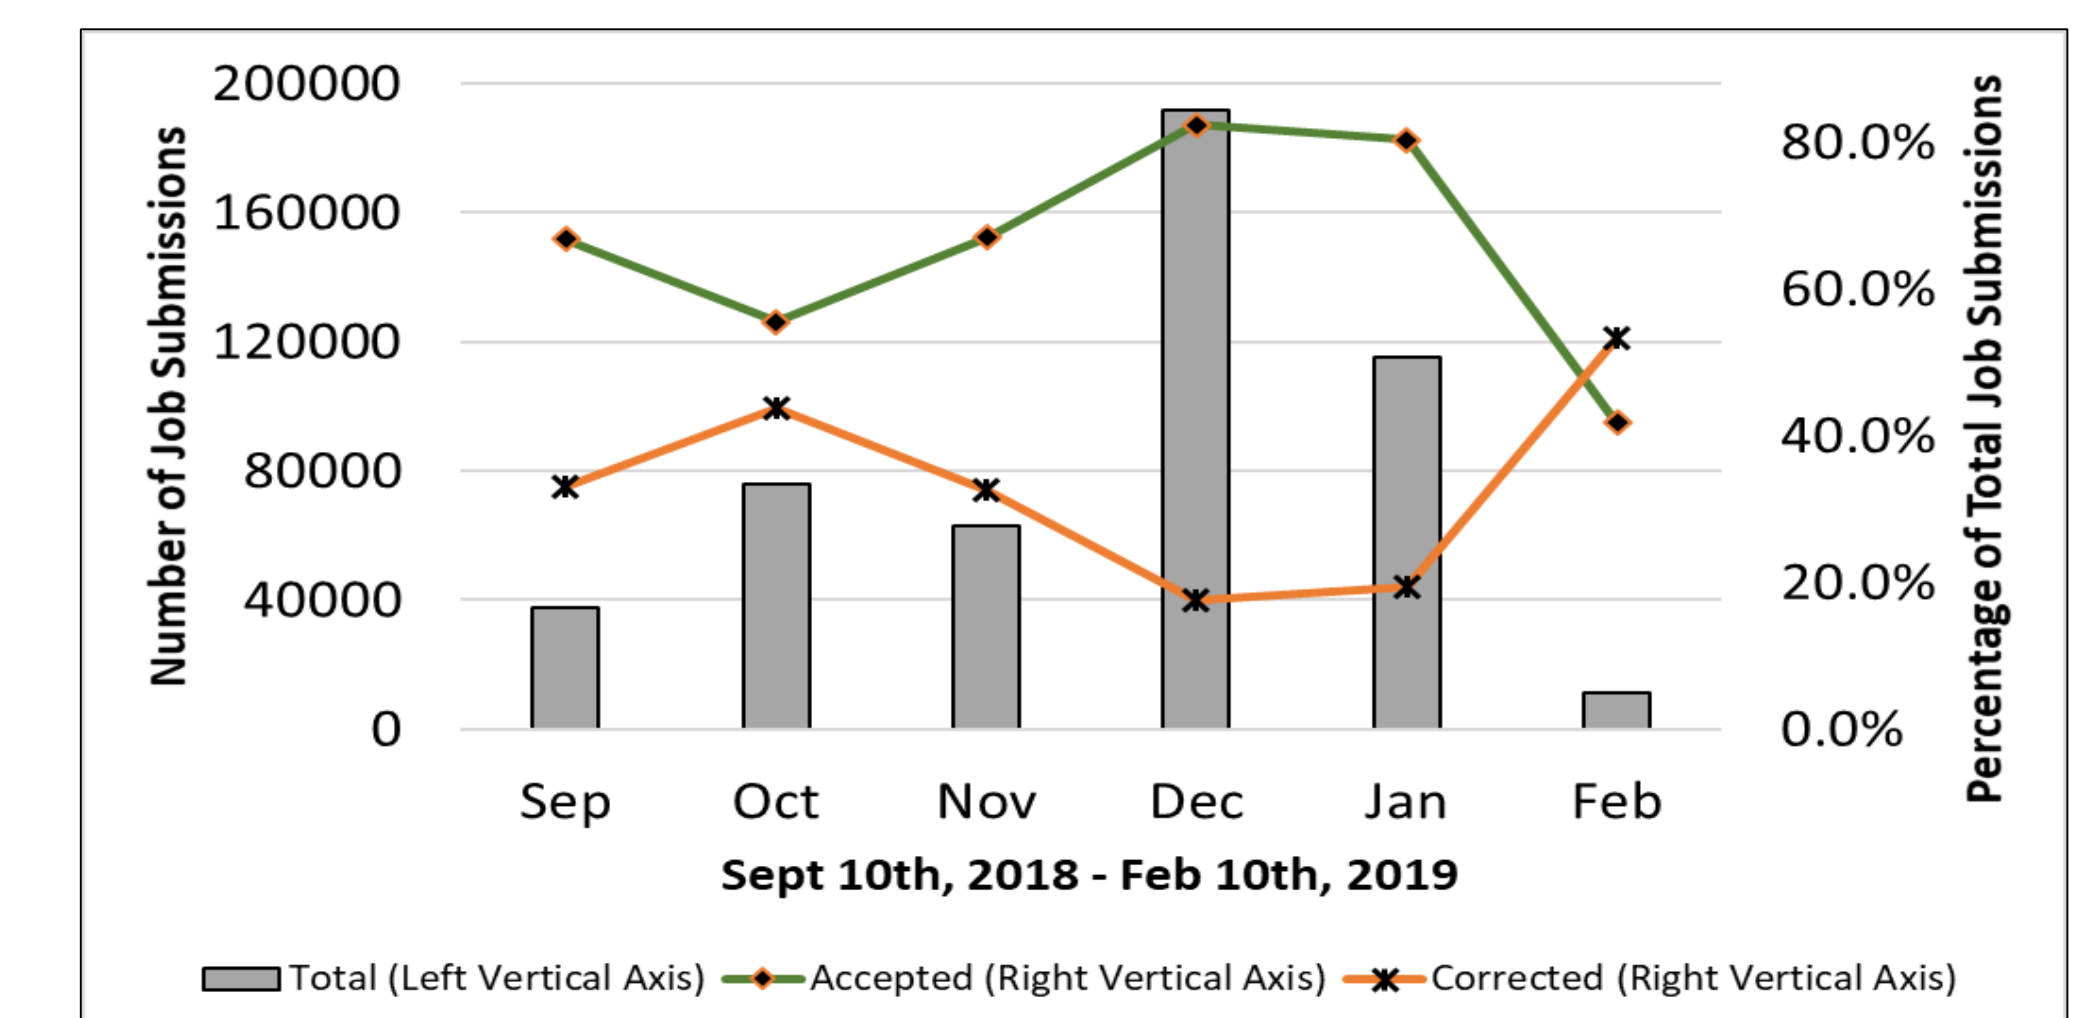
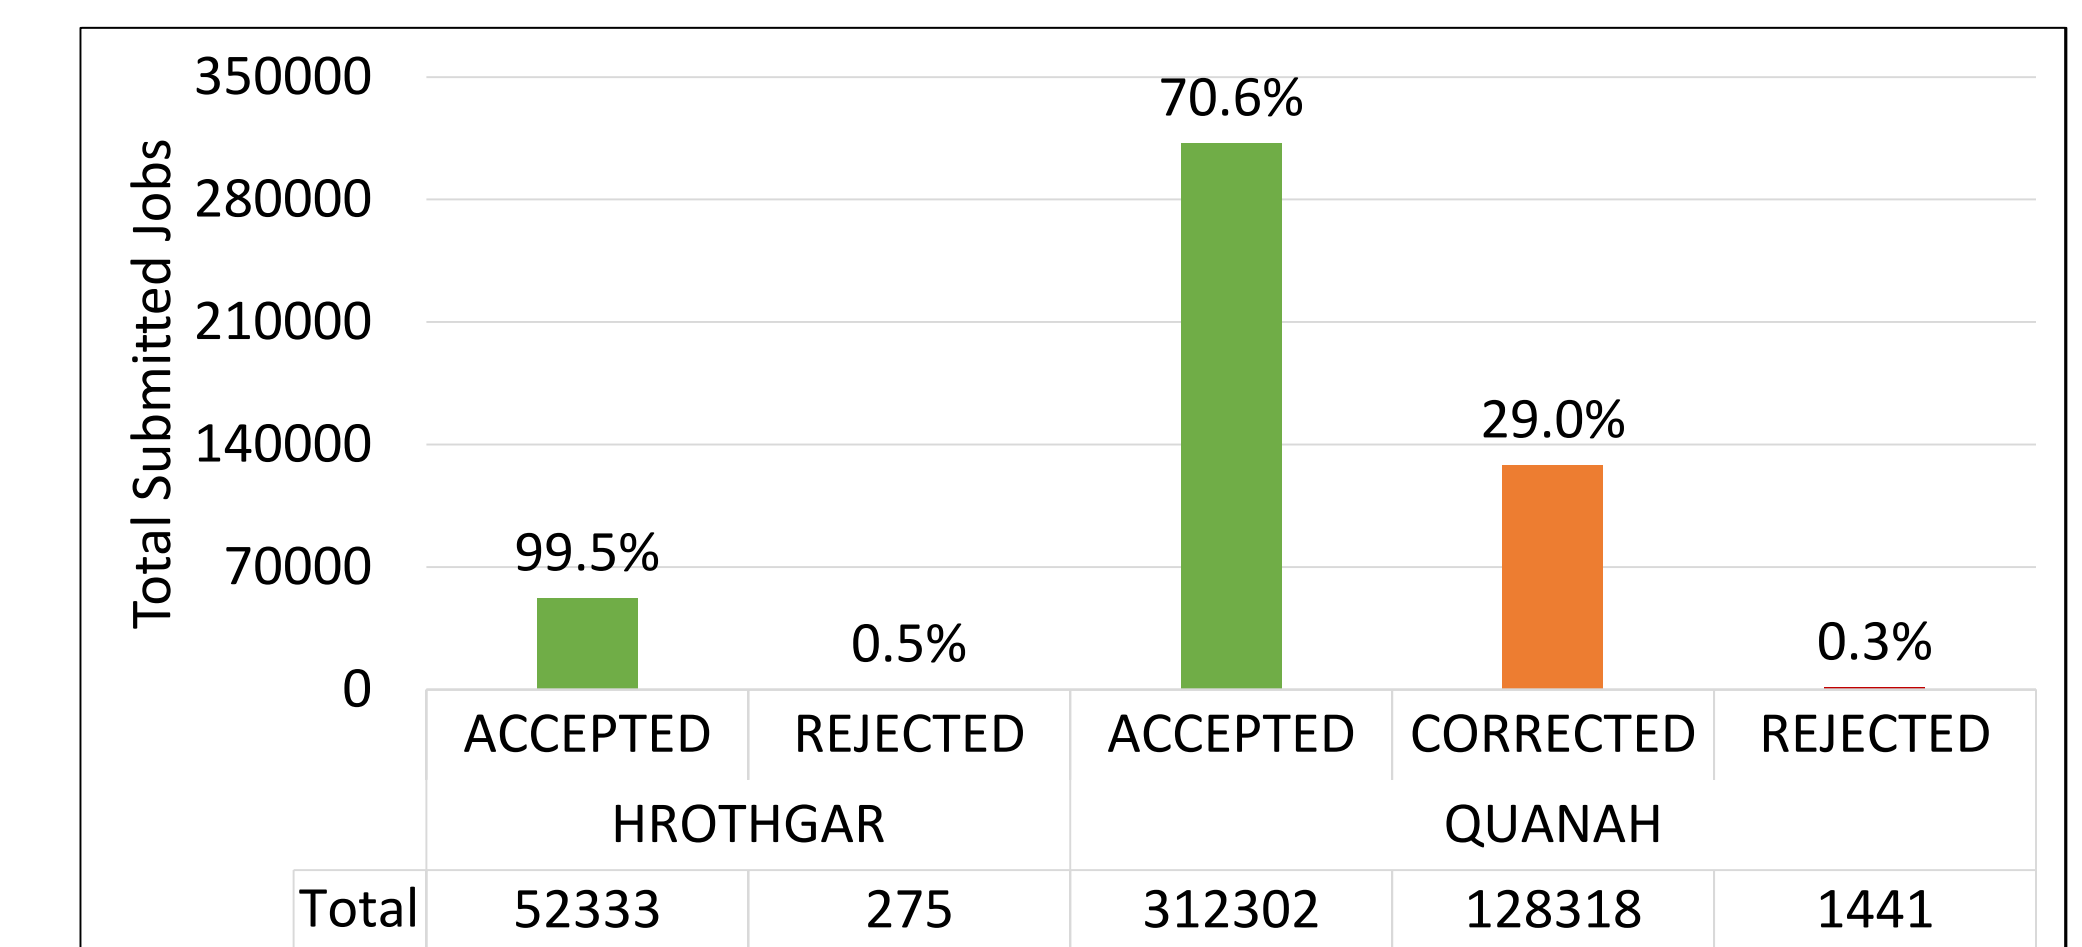
Job Submission Verifier (JSV)

- ACCEPT:** allow the job to be received by UGE components
- CORRECT:** modify the users' requests based on determinable criteria
- REJECT:** drop the potential faulty jobs with a wrong and uncorrectable request format



Evaluation

Sept 10th, 2018 – Feb 10th, 2019



Acknowledgment

- HPC resources that have contributed to this research study, were provided by High Performance Computing Center (HPCC) at Texas Tech University
- This research is supported in part by the National Science Foundation under grant CNS-1338078, CCF-1718336, OAC-1835892, CNS-1817094.

References:

- Li, H., Groep, D., Wolters, L. and Templon, J. 2006. Job Failure Analysis and Its Implications in a Large-Scale Production Grid. 2006 Second IEEE International Conference on e-Science and Grid Computing (e-Science'06) (Dec. 2006), 27–27.
- Anglano, C., Brevik, J., Canonico, M., Nurmi, D. and Wolski, R. 2006. Fault-aware scheduling for Bag-of-Tasks applications on Desktop Grids. 2006 7th IEEE/ACM International Conference on Grid Computing (2006), 56–63.
- Fadishei, H., Saadatfar, H. and Deldari, H. 2009. Job failure prediction in grid environment based on workload characteristics. 2009 14th International CSI Computer Conference (Oct. 2009), 329–334.
- Engineering, U. 2016. Univa Grid Engine Documentation Univa Grid Engine Administrator's Guide. Univa Corporation.
- Univa Grid Engine: <http://www.univa.com/>.