Wilson security issues

Glenn, Irwin, Joe B, Mine, Tony T

27/8/2020

Our current situation/Some definitions

- A system which is accessible from the grid (aka for non-FNAL badged people) is in the Open Science Enclave (OSE). All the Fermigrid and T1 worker nodes for instance.
- A system that allows **interactive access** for FNAL badged people is in the General Computing Enclave (GCE).
- You could have a machine that was in the GCE for some period of time, move to the OSE, and then back to the GCE. Month time periods ok; week ok; doing it per job is not ok (not a technical issue, could you walk through the computer room and point at the machines in the GCE?).

Two topics of concern on Wilson

1. Is a shared filesystem which is readable, writeable, and executable across the worker nodes and the interactive nodes with both OSG and FNAL badged jobs running a significant security risk?
   - The feeling is that this is an acceptable practice. There is authentication to run a job on the cluster and those jobs will be run under appropriate unix uids so that is sufficient to not worry about trojans.
2. Can we allow interactive access to the worker nodes on Wilson via the Slurm srun command while OSG jobs are running.
   - This is a problem because machines would flip between GCE and OSE.

Slurm use cases on Wilson

1. A batch system (sbatch command)
   - When used in a non-interactive fashion by everyone there is no problem. We could have OSG jobs and FNAL badged jobs running at the same time.
   - The Wilson worker nodes would be in the OSE
   - The Wilson login node(s) would be in the GCE
   - The Wilson CondorCE receiving OSG jobs would be in the OSE.
2. As an interactive debugging (srun command) platform
   - If there is interactive access to the worker nodes for FNAL badged users then those worker nodes are in the GCE. As previously mentioned switching back and forth between GCE and OSE on a per job basis isn't defendable under current policy.

The future

- Security team understands that the computing world is becoming very elastic and containerized
- Usage of hardware resources is getting mixed quickly by containerization and with orchestration like kubernetes to make maximum use of hardware
- It will not be quick but the security team agrees that the security policies need to be updated to allow this sort of intermixing of usage. We should be able to change the policies to allow Wilson to be used by both OSG jobs and FNAL badged jobs at the same time.

What are our options in the short term?

1. We could allow only FNAL badged users to run jobs on Wilson. It is a small cluster relatively.
2. We can allow both FNAL badged users and OSG users on the cluster at the same time but we'd have to disable interactive use (srun command) somehow. Not sure that's possible.
3. We could keep it covid only for a while but the policy changes won't happen quickly
4. We can partition up the cluster via slurm administrative config to have some worker nodes in the GCE for FNAL badged jobs and other worker nodes in the OSE for OSG jobs (27 nodes have 16 cpus and 4 gpus, 100 additional nodes have 16 cpus only).

Decison

- From Projects meeting on Sep 3, 2020.
  - Choosing option 4
  - Setup all the non-gpu and 25% of the GPU nodes for FNAL badged people
  - Setup 75% of the GPU nodes for OSG usage
  - We can re-evaluate over time and perhaps shift the partition layout
  - This clearly delineates which machines are in the GCE and which are in the OSE so this complies with our current policies without changes. We will NOT move systems between the two enclaves on a short time period basis.
  - We will setup a new node as the CondorCE for the OSG jobs to flow into the cluster as it's currently on a GCE machine and must be moved to an OSE machine.