

Infrastructure Issues

1. ~~Make script to test c2e on local cluster~~
2. ~~Fix volume c2e tests so they can pass consistently~~
3. Investigate how to test under different environments such as selinux support and fibre channel availability, etc

High Level Overview

Test cases	AWS EBS	GCE PD	OpenStack Cinder	Azure File	NFS	iSCSI	Glusterfs	Ceph RBD	Ceph FS	Fibre Channel	Notes
SELinux and FSGROUP	N	N	Y	N	Y	Y	Y	Y	Y	N	
Attach/Detach		Y									Only GCE PD has tests
PV/PVC					Y						Only NFS has tests
Dynamic Provisioning	Y	Y	Y								Only AWS/ GCE/ Cinder have tests

Test cases:

1. [SELinux and FSGroup](#)

- Functional tests:
 - Volume(all kinds) mount read/write permissions, FSGroup, supplemental group, owner and selinux testing
 - Verify that correct user could access the mount directory
 - Verify that user with correct group id could access the mount directory
 - For block devices, verify the mount directory is created with correct uid/gid.
 -
 - Verify the files created under the mount directory has correct uid/gid
 - Verify directory without correct selinux label can not be accessed.
- Special case tests
 - If FSGroup doesn't match that in the containerized volume server, access to the volume should expect denial

2. Persistent Volume

- Function tests:
 - [Persistent_volumes.go](#) currently only tests NFS PV with recycle policy. Other tests cases are to create PVs with all retention policies and use all volume plugins (but with focus on Cloud providers). Verify the PVCs can be bound and after PVCs are deleted, recycling is supposed to fail for all but NFS and hostPath.
 - Create multiple PVs with different sizes and multiple PVC also with different sizes. Verify when PVC is bound, the match PV's size is no less than claim size.
 - Create one PV and multiple PVCs. Verify exactly one PVC bound while other PVCs are pending
 - Create PVCs with different access modes. Verify PVCs are bound to the PVs with the matching access modes.
 - Create Pod that uses multiple PVCs. Verify all PVCs if successfully bound, should be mounted.
 - Create PV that requests a specific PVC. Create similar PVC to make sure bind is respected (new design may differ..but should test strict binding of PV/PVC)

- Delete PVC and make sure PV changes to proper state given retention policy defined in the PV
 - Create PVCs in multiple namespaces, verify they can bind to PV that matches access/size criteria
3. [PD Attach/Detach](#)
- Test detach/attach on Cinder and RBD: ensure only one can attach as read/write
4. Bug fixes:
- Test safeMountAndFormat on rbd, aws, gce pd, cinder, iSCSI
 - Test iSCSI portal without specifying default port
 - Test Ceph and Gluster with invalid gluster endpoints and Ceph mon
 - Multiple iSCSI LUN rw and ro and ensure the access behavior is correct
 - [AWS hosts support at most 39 EBS volumes](#)
 - [EBS should be deleted after Pod exits](#)
 - [Test NoDiskConflict on pd/ebs/rbd](#)
5. TBD:
- Add test cases after binder refactoring is complete
 - Add test cases after attach/detach controller is complete
 - Add test cases after dynamic provisioning is complete
 - Load test: Generate concurrent PV claim and delete requests