



Ceph-SMB: Updates and future roadmap

Sachin Prabhu - sprabhu1@ibm.com



- IBM/Red Hat
- Ceph Team - SMB Service
- <https://github.com/samba-in-kubernetes>
 - Samba Container
 - SambaCC
 - Samba-Build
 - Samba Integration Testing

- Ceph
 - Software defined storage platform
 - CephFS - Clustered filesystem built on Ceph
- SMB
 - Network file sharing protocol primarily used by Microsoft Windows and Mac OS X. Linux client.
 - SMB v 3.1.1
 - Samba - open source SMB file server
 - <https://www.samba.org/>

- Ceph SMB manager module
 - Adds the "ceph smb" command
 - Calls ceph orchestration to deploy the service.
 - reads the service spec
 - spins up containers
 - Init containers - runs checks and setup.
 - Smbd container - fileserver
 - CTDB - clustering
 - Winbind - id mapping
 - Prometheus - metrics

- Creating a Ceph-SMB service - Imperative method
 - `ceph smb cluster create test1 active-directory \`
`--domain-realm=MYDOM.EXAMPLE.ORG \`
`--domain-join-user-pass=Administrator%Ph0nyPassw0rd \`
`--custom-dns=192.168.76.210`
 - `ceph smb share create test1 pics cephfs \ --subvolume=company/photos`
`--path=/ \`
`--share-name="Company Photos"`

- Creating a Ceph-SMB service - Declarative method

```
resources:  
  - resource type: ceph.smb.cluster  
    cluster id: mysmb  
    auth mode: user  
    user group settings:  
      - source type: resource  
        ref: ugl  
  
  - resource type: ceph.smb.usersgroups  
    users groups_id: ugl  
    values:  
      users:  
        - name: test1  
          password: x  
        - name: test2  
          password: x  
      groups: []  
  
  - resource type: ceph.smb.share  
    cluster id: mysmb  
    share id: share2  
    intent: present  
    name: share2  
    readonly: false  
    browseable: true  
    cephfs:  
      volume: mycephfs  
      subvolume: smb2  
      path: /  
      provider: samba-vfs
```

- `ceph smb apply -i mysmb.yml`

- user/group based access control.
- active directory integration
- clustering support
- Go-ceph go bindings available
 - <https://github.com/ceph/go-ceph/tree/master/common/admin/smb>
- <https://docs.ceph.com/en/latest/mgr/smb/>

- I/O Rate Limiting
 - Implemented using the new Samba vfs module - `vfs_aio_ratelimit`
 - Per share setting
 - Applicable cluster wide
 - Limit read/write I/Os - `read_iops_limit`, `write_iops_limit`
 - Limit read/write bandwidth - `read_bw_limit`, `write_bw_limit`
 - Supports burst read/writes - `read_burst_mult`, `write_burst_mult`
 - <https://docs.ceph.com/en/latest/mgr/smb/#qos-parameters>
 - Samba XP talk on implementation of rate limiting
<https://www.youtube.com/watch?v=68bl6CtBqgU>

- SMB per share metrics
 - Accessed from the dashboard.
 - Data now available on per share basis.
 - Monitor throughput and iops.
 - Implemented using Samba profiling counters.
 - <https://www.ibm.com/docs/en/storage-ceph/9.9.0?topic=clusters-smb-overview-dashboa rd>

- SMB3 multichannel support
 - Uses multiple network connections
 - Aggregates network bandwidth
 - Improves fault tolerance
 - <https://www.ibm.com/docs/en/storage-ceph/9.9.0?topic=management-smb-multichannel>

New Feature

- IP Based access control for SMB Shares
 - Setup on per share basis.
 - Documented under hosts_access for share resource-
<https://docs.ceph.com/en/latest/mgr/smb/#share-resource>

- Bring Your Own Key - BYOK
 - Encryption at rest.
 - Key lifecycle controlled by cluster owner.
 - Keybridge config per cluster.
 - Fscrypt setting per share.
 - Documented under keybridge resource and fscrypt for share resources - <https://docs.ceph.com/en/latest/mgr/smb/>

- Remote Control using GRPC
 - Remotely control SMB service using GRPC client.
 - GRPC client expected to be replaced by custom tool
 - Documented under remote_control - <https://docs.ceph.com/en/latest/mgr/smb/>
 - The API documented at - <https://github.com/samba-in-kubernetes/sambacc/blob/master/sambacc/grpc/protobuf/control.proto>
 - Command examples:
 - Status, ConfigDump, ConfigSharesList
 - CloseShare, KillClientConnection
 - CTDBStatus, CTDBMoveIP
 - SetDebugLevel, GetDebugLevel

- SMB-RGW export
 - Uses a new Samba VFS module - `vfs_ceph_rgw`
 - Supports: Directory listing, File reads, File writes(sync only)
 - Inherits librgw limitation. No directory renaming. No advanced ACL support. No random writes.
 - Under Development - Samba MR:
https://gitlab.com/samba-team/samba/-/merge_requests/4381
 - Samba XP presentation: <https://www.youtube.com/watch?v=AMiNrYmZ4pc>

- MacOS support
 - Sets up required Samba modules - vfs_fruit, vfs_catia, vfs_stream_xattr
 - Enables the SMB AAPL extension
 - Improves directory lookup time
 - Allows Apple allowed file characters(:, ?, *, <, >, |)
 - Enables streams. Used to store file metadata.
 - Enables other apple specific features: copy progress bars, server side copying.



Thank You

Sachin Prabhu
sprabhu1@ibm.com

