

# Introduction



## Managed SMB Support in Ceph: The New SMB MGR Module

Avan Thakkar  
@IBM



Avan - [avanthakkar](#) 

- Software Engineer at IBM
- 5 years of experience working with Ceph Storage, currently working with Samba team.
- Like to watch different Sports, reading HackerNoon articles in weekend afternoons.

# Ceph: Quick Intro



- Ceph is a distributed, software-defined storage
- It provides object, block, and file storage under a single system.
- CephFS is a POSIX-compliant, distributed filesystem, but it doesn't natively speak SMB—it needs a bridge.
- This is where **Samba comes in**: it enables SMB-based access to CephFS.

# How Samba Integrates with CephFS



- Samba connects with **CephFS through the `vfs_ceph` module**, allowing seamless SMB access.
- It enables Windows-compatible permissions (ACLs), authentication (Kerberos, AD), and locking mechanisms.

# SMB on Ceph – The Missing Piece



## We Have the Right Plugs, But...

- Samba + CephFS work, but you set up everything manually
- Complex setup – Manually defining Samba services, shares, authentication, so on..
- No built-in orchestration – Handling Active Directory, clustering, and failover can get tedious.

# Introducing SMB MGR Module



- A Ceph MGR module designed to natively manage SMB services within Ceph
- Provides centralized orchestration for Samba-backed SMB shares
- Seamlessly integrates with CephFS to simplify deployment and reduce manual setup
- Supports Active Directory
- Exposes commands (`ceph smb . . .`) to configure & monitor SMB
- Inspired by the NFS mgr module, following a similar approach to managing SMB clusters and shares.

# SMB Service Orchestration Workflow



- SMB service spec created in cephadm
- Spinning Up Init Containers - configs, setting up authentication, and ensuring required Ceph resources exist
- Deploying SMB Components as Daemons - sidecars
  - Cephadm orchestrates the deployment of key SMB daemons, such as:
    - **smbd** (handles file sharing requests)
    - **winbindd** (manages authentication and domain integration)
    - **CTDB** (ensures high availability and clustering)
    - **smbmetrics** (collects monitoring data for performance insights)

# Imperative and Declarative Management



- There are two ways to interact with the SMB module: imperative and declarative.
- The **imperative approach** relies on direct commands like `ceph smb cluster create ...` and `ceph smb share rm ...`, similar to how Ceph's NFS manager module is managed via the command line.
- The **declarative approach** uses `ceph smb apply` to process structured **resource descriptions** defined in YAML or JSON
- Compatible with each other



# Features



- Supports Active Directory integration
  - Join Authentication
    - SMB clusters can be joined to an Active Directory (AD) domain using `ceph smb join`
    - Allows authentication using domain credentials.
  - User & Group Management
    - Supports user/group-based access control.
    - Integrates with AD and local authentication methods.
- Allows defining share access restrictions
- Allows users to specify custom DNS settings
- Enables users to define custom global SMB configuration parameters (AYOR)
- Samba clustering support using CTDB
- Exposes Prometheus compatible performance & health metrics
- UI Management support via Ceph Dashboard

# Features: Earmarking Subvolume



- **Problem Statement**
  - Mixed protocol use within the same subvolume is unsupported and can lead to data corruption.
  - Users might accidentally expose the same subvolume via both SMB and NFS-Ganesha, causing protocol conflicts.
- **Solution: Protocol-Specific Earmarking**
  - Introduced an earmarking mechanism that restricts each subvolume to a single protocol (either SMB or NFS).
  - An earmark is stored as an extended attribute (xattr - `user.ceph.subvolume.earmark`) on the root directory of a subvolume
  - Dot separated format (`protocol.[optional_scope]`)
  - top level scope - nfs or smb protocol
- Inter & Intra Module Level Scope Handling
- commands exposed via cephfs to be shown in Demo



- Smb share management
- Earmarking
- Debugging & Troubleshooting
  - cephadm log management
  - RADOS cmds

# Future Plans



- Multiprotocol
  - Concurrent SMB/NFS support for the given volume/subvolume path
- Enhanced `ceph smb show` Command
- Offline Domain Join (ODJ) Support