

The Ceph Samba Gateway

David Disseldorp
ddiss@samba.org

Samba

- File and print sharing
 - SMB / CIFS, SMB2 and SMB3+ dialects
- Authentication
 - NTLMv2 and Kerberos
- Identity mapping
 - Windows SIDs to *uids* and *gids*
 - Active Directory domain member or domain controller



Clients

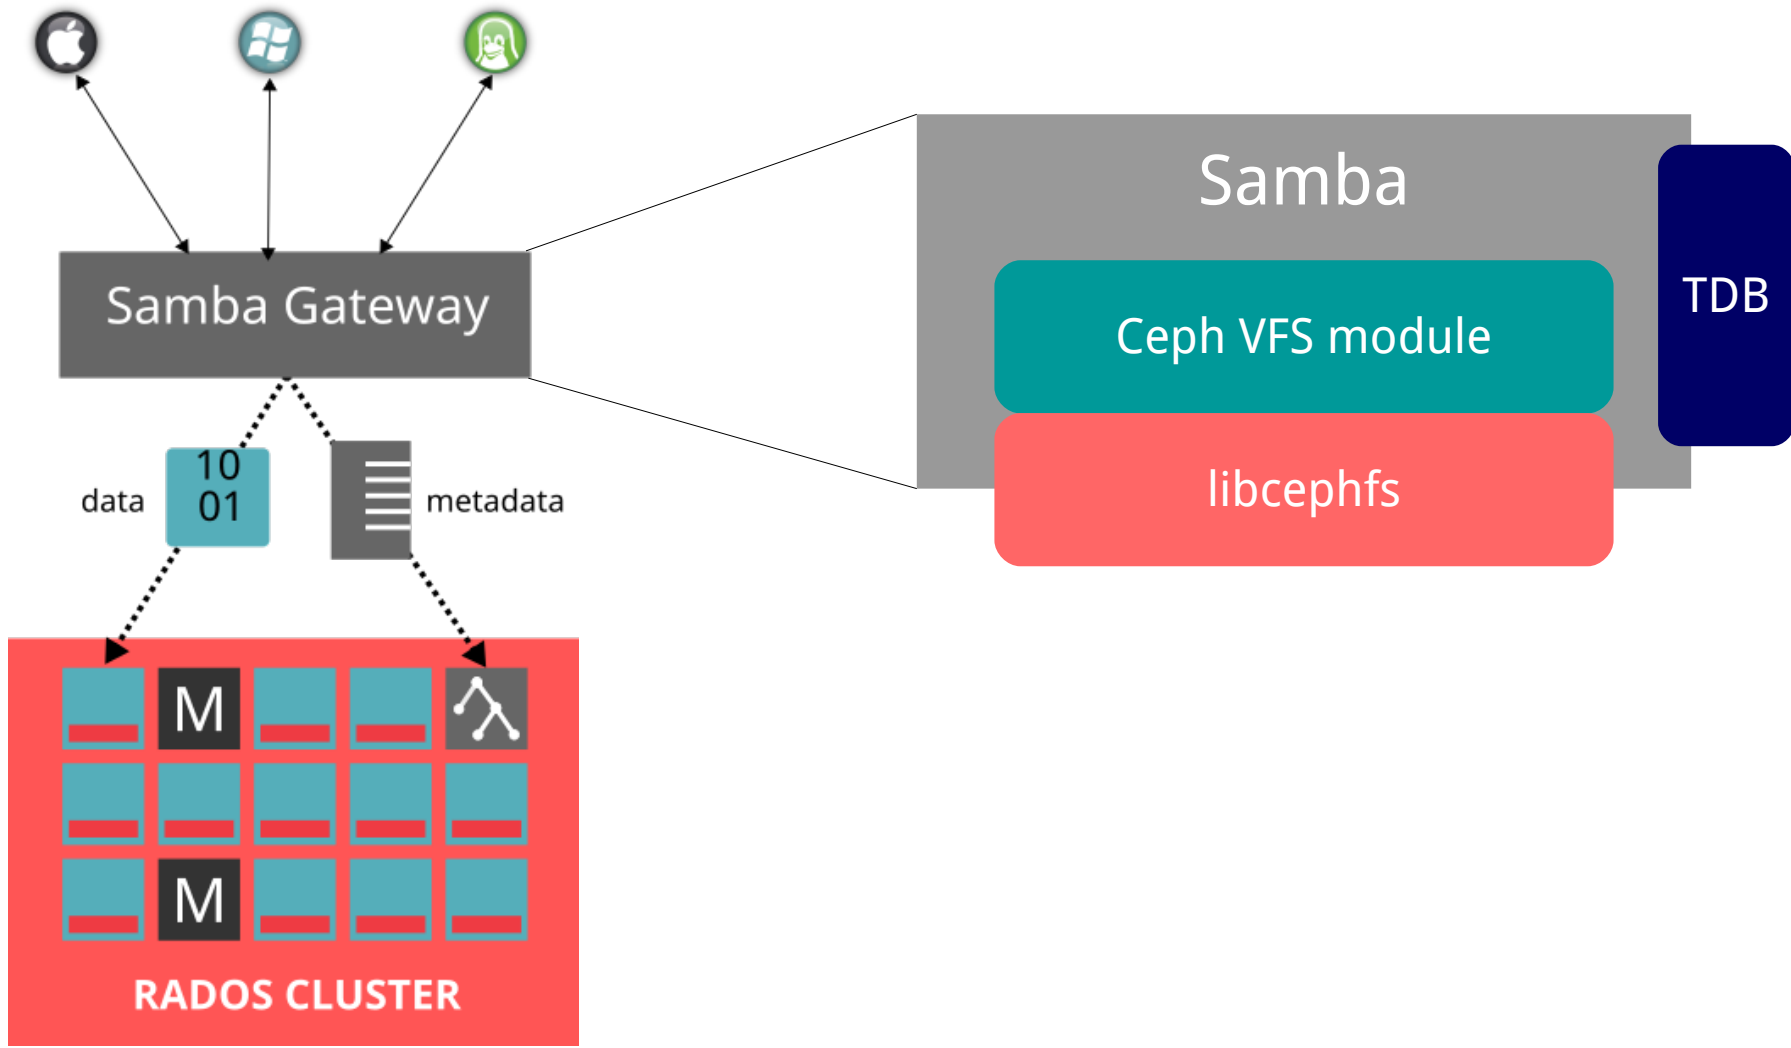
- Windows
 - Reference client
 - Protocol specification publisher
- macOS
 - AFP replaced by SMB2 as default for Mavericks (2013)
- Linux
 - CIFS kernel module and Samba smbclient



Clustered Trivial Database (CTDB)

- Samba persistent state stored in TDB key-value store
- Share some of this state across multiple nodes
 - Cluster consistent database
 - Reliable messaging
- HA features bolted on
 - Monitoring and failover





Samba Ceph Integration

- Tech-preview in SES5
 - SLE-HA provided for scale-out clustering
- CephFS module for Samba: *vfs_ceph*
 - Maps SMB file and directory I/O to libcephfs API calls
- Active Directory membership
- POSIX Access Control Lists (ACLs)



Samba Ceph Integration

- Ceph RADOS clustered mutex helper for CTDB
- Static cephx credentials
 - Regardless of Samba authenticated user
- Ceph librados service integration (*coming soon*)



Testing

- Samba smbtorure
- cifs.ko fstests
- Interoperability
 - MacOS, Windows, Hyper-V, etc.



Current and Future Challenges

- Cross-protocol client support
 - Shared ACL model
 - Coherent client caching
 - Map leases to CephFS *FILE* and *AUTH* capabilities
- Unified authentication and user mapping
 - Use Kerberos / AD for Samba gateway



Current and Future Challenges

- Multichannel support
 - Experimental in upstream Samba
- Automated deployment
- Change notifications



Future Challenges

- Witness protocol
 - Advertise Samba cluster state to clients
 - Transparent client failover
 - Load balancing

- Persistent handles



Future?


- Ceph backed key-value store
- Replace or modify CTDB
 - Samba database API demanding
 - Multiple processes and writers
 - Record locking and transactions
- Object gateway



References

- Samba: <https://samba.org/>
- CTDB: <https://ctdb.samba.org/>
- SMB 3.1.1 encryption: [https://technet.microsoft.com/en-us/library/dn551363\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/dn551363(v=ws.11).aspx)
- Multichannel deployment: [https://technet.microsoft.com/en-us/library/dn610980\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/dn610980(v=ws.11).aspx)
- Witness Protocol:
http://www.sambaxp.org/archive_data/SambaXP2015-SLIDES/wed/track1/sambaxp2015-wed-track1-Guenther_Deschner-ImplementingTheWitnessProtocolInSamba.pdf
- Samba Multichannel Blocker Bug: https://bugzilla.samba.org/show_bug.cgi?id=11897
- CephFS cache flags:
<https://jtlayton.wordpress.com/2016/09/01/cephfs-and-the-nfsv4-change-attribute/>





Join Us at www.opensuse.org



License

This slide deck is licensed under the Creative Commons Attribution-ShareAlike 4.0 International license. It can be shared and adapted for any purpose (even commercially) as long as Attribution is given and any derivative work is distributed under the same license.

Details can be found at <https://creativecommons.org/licenses/by-sa/4.0/>

General Disclaimer

This document is not to be construed as a promise by any participating organisation to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. openSUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for openSUSE products remains at the sole discretion of openSUSE. Further, openSUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All openSUSE marks referenced in this presentation are trademarks or registered trademarks of SUSE LLC, in the United States and other countries. All third-party trademarks are the property of their respective owners.

Credits

Template

Richard Brown
rbrown@opensuse.org

Design & Inspiration

openSUSE Design Team
<http://opensuse.github.io/branding-guidelines/>