

# Ceph In A Windows World

TUT-1121

Mike Latimer  
Engineering Manager  
SUSE / [mlatimer@suse.com](mailto:mlatimer@suse.com)

Alessandro Pilotti  
CEO / CTO  
Cloudbase Solutions / [apilotti@cloudbasesolutions.com](mailto:apilotti@cloudbasesolutions.com)

# Ceph In A Windows World

1. Overview of Ceph
2. Challenges with Ceph in Windows Environments
3. SUSE + Cloudbase Solutions
4. Sample use cases
5. Demos!

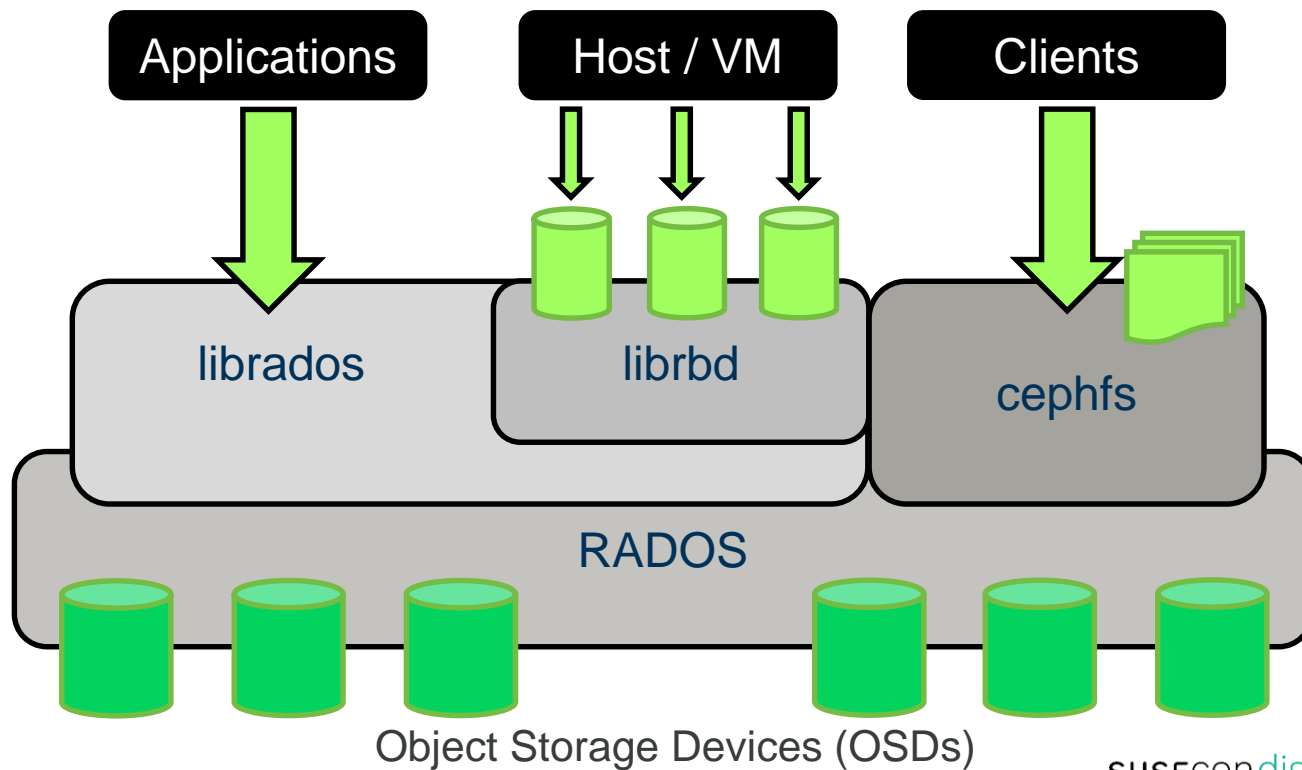


# Ceph Overview

Highly scalable and resilient, distributed storage system

- Based on Reliable Autonomic Distributed Object Store (RADOS)
  - Pseudo-random, algorithmic data distribution via Controlled Replication Under Scalable Hashing (CRUSH)
- Core services include:
  - Object storage
    - S3/Swift compatible gateway (and librados)
  - Block storage
    - Kernel module or librbd
  - File Storage
    - CephFS kernel module or FUSE client (libcephfs and librados)

# Ceph Architecture

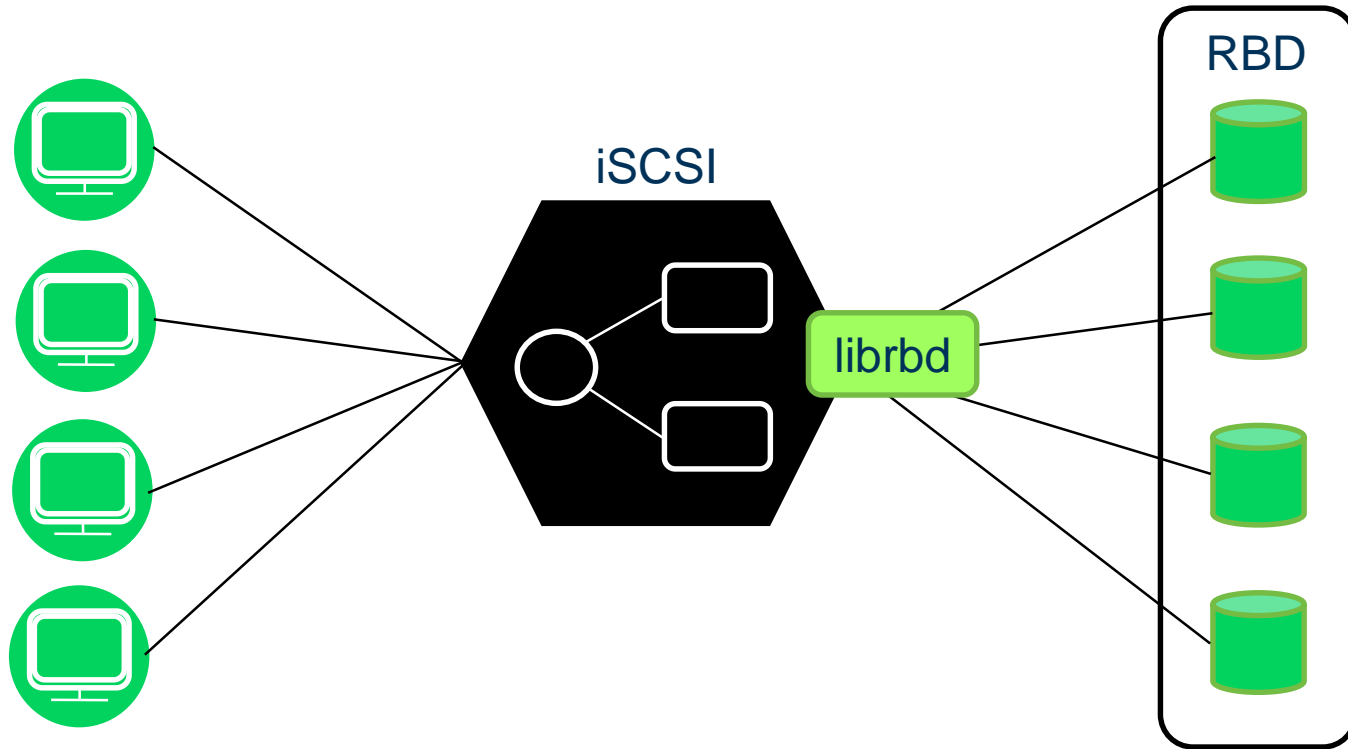





# Gateway vs. Native


- Gateway
  - iSCSI, Samba (SMB), S3
  - Client connects to gateway; gateway connects to OSD resource
    - All traffic flows through gateway(s)
- Native
  - Client uses CRUSH algorithm to determine where data is stored
    - Client performs I/O directly to individual OSDs
    - No I/O bottleneck, no single point of failure

# Gateway I/O Path



# Native I/O Path





*According to Gartner, Microsoft Windows Server accounts for 56.8% of the server market.*

Market Share Analysis: Server Operating Systems, Worldwide 2017



# Challenges With Ceph In Windows

- Windows environments currently access Ceph resources through:
  - iSCSI or Samba gateway
- iSCSI
  - Provides block device support
    - RBD image served as iSCSI target to client initiator
    - Requires at least one iSCSI gateway
    - No direct client to OSD connections!

# Challenges With Ceph In Windows

- Samba
  - Provides file access to Ceph storage (similar to CephFS)
    - Requires a Samba gateway
      - Multi-protocol support
      - Maps Windows SIDs to Linux UID/GIDs
        - Common use case is user home directories
    - No direct client to OSD connections!





# SUSE And Cloudbase Solutions

- Focused on bringing native RBD technology to Windows!
  - Remove gateway requirements
  - Enable Windows clients to connect directly to Ceph resources
- Development focus includes:
  - Porting librbd and librados to Windows
  - RADOS block driver (signed Windows storeport miniport driver)
- Everything is upstream and open-source!



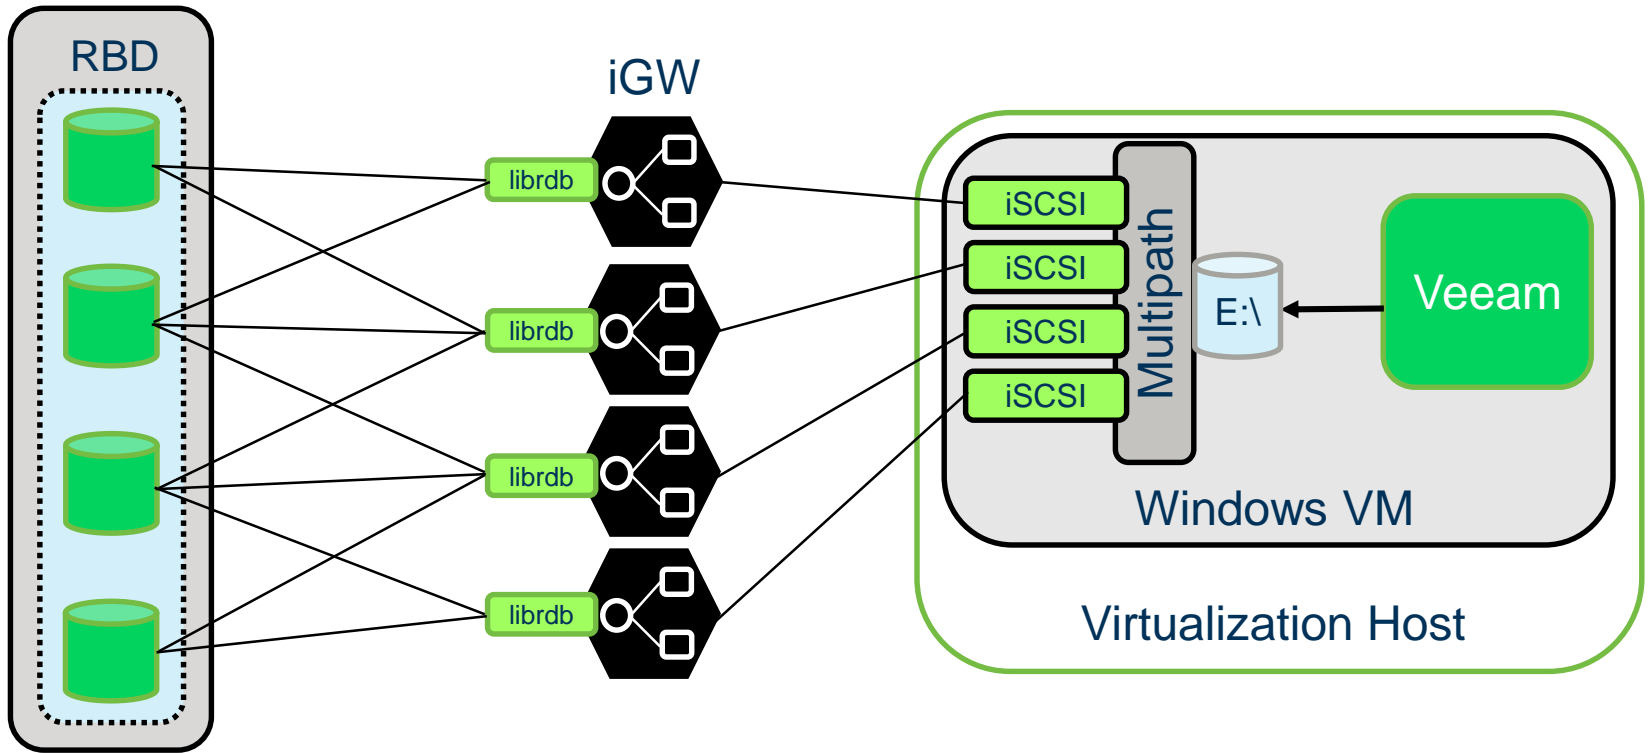
# Benefits Of Windows Native Driver

- Performance!
  - Replace gateway model with native Ceph
- Ease of deployment
  - Eliminate complexities of iSCSI setup
- Increased compatibility of Windows and Linux environments


# Target Use Cases

- Windows based backup solutions
  - Veeam
    - Block device with ReFS (for deduplication)
- Microsoft Hyper-V
  - 22.5% of market (37.9% of large companies)
- Microsoft Cluster Server (MSCS)
  - Shared data storage

# Sample Veeam Configuration

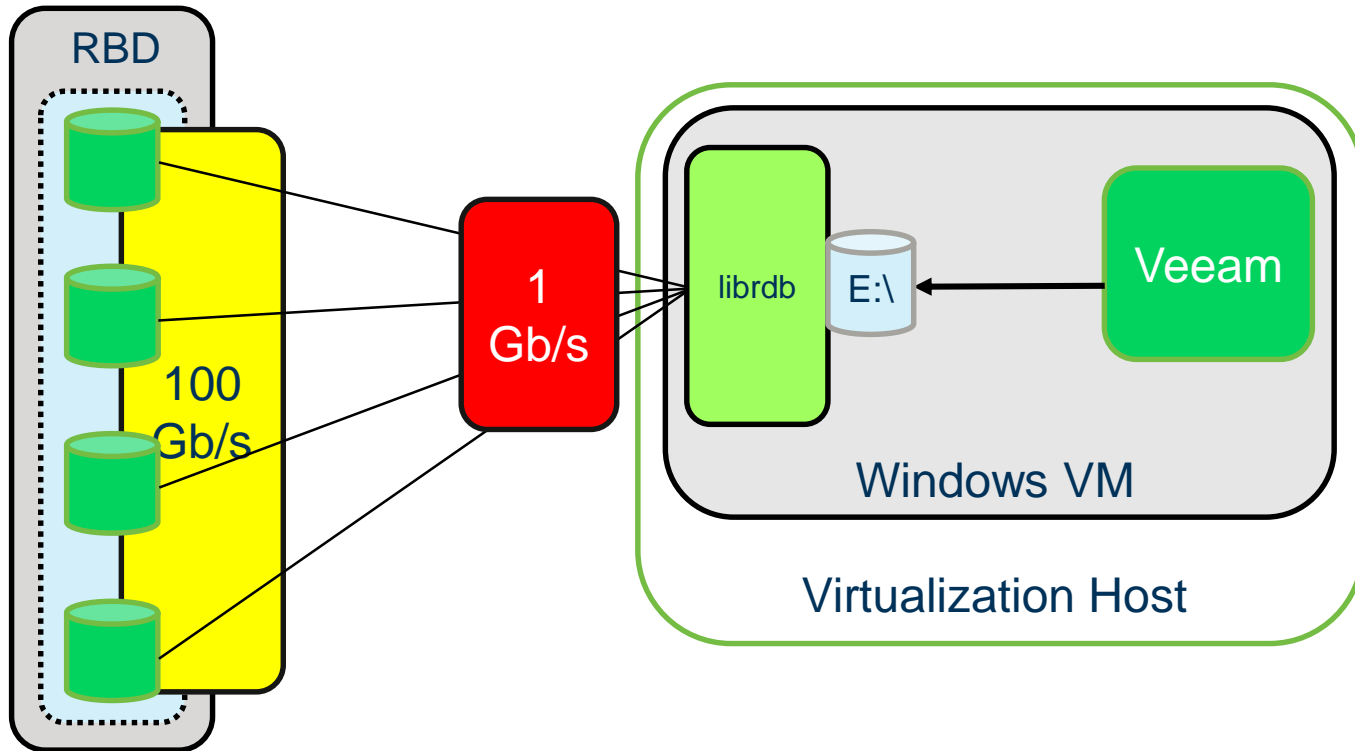


# Sample Veeam Configuration w/librdb





# Network Speeds Still Matter!





# Demos!

- RBD driver
- Cephfs
  - Windows And Linux Interoperability
- Using RBD with Hyper-V
  - Convert an existing qcow2 image to RBD

# Work In Progress

- Functionality and Performance
  - File locking and ACLs
  - Code-level tuning and optimizations
- Driver delivery/installation mechanism
  - Graphical installation tool vs. basic driver + inf file?
- Pricing and Availability
  - Details are still being finalized
    - Significantly less than traditional storage model!



# Thank You

## General Disclaimer

This document is not to be construed as a promise by any participating company 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. SUSE 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 SUSE products remains at the sole discretion of SUSE. Further, SUSE 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 SUSE marks referenced in this presentation are trademarks or registered trademarks of SUSE, LLC, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

The logo features the text "SUSEcon digital '20" in white. "SUSE" is in a bold, uppercase sans-serif font, while "con digital" is in a lighter, lowercase sans-serif font. The "'20" is contained within a white, rounded, teardrop-shaped icon. The background is a vibrant green with a faint, white, grid-like pattern that resembles a leaf's vein structure.

SUSEcon digital '20