

SUSE® OpenStack and Ceph integration

Michal Jura

Senior Software Engineer
Linux HA/Cloud Developer
mjura@suse.com



Agenda

- Brief introduction to Ceph
- Brief introduction to OpenStack
- OpenStack and Ceph integration
- User experience in production
- Live demo presentation

Brief introduction to Ceph


From 10,000 meters

- Open Source Storage Distributed solution
- The best choice of distributed storage for OpenStack
- Lots of goodies
 - Distributed Object Storage
 - Redundancy
 - Efficient Scale-Out
 - Can be build on commodity hardware
 - Lower operational cost

From 1,000 meters

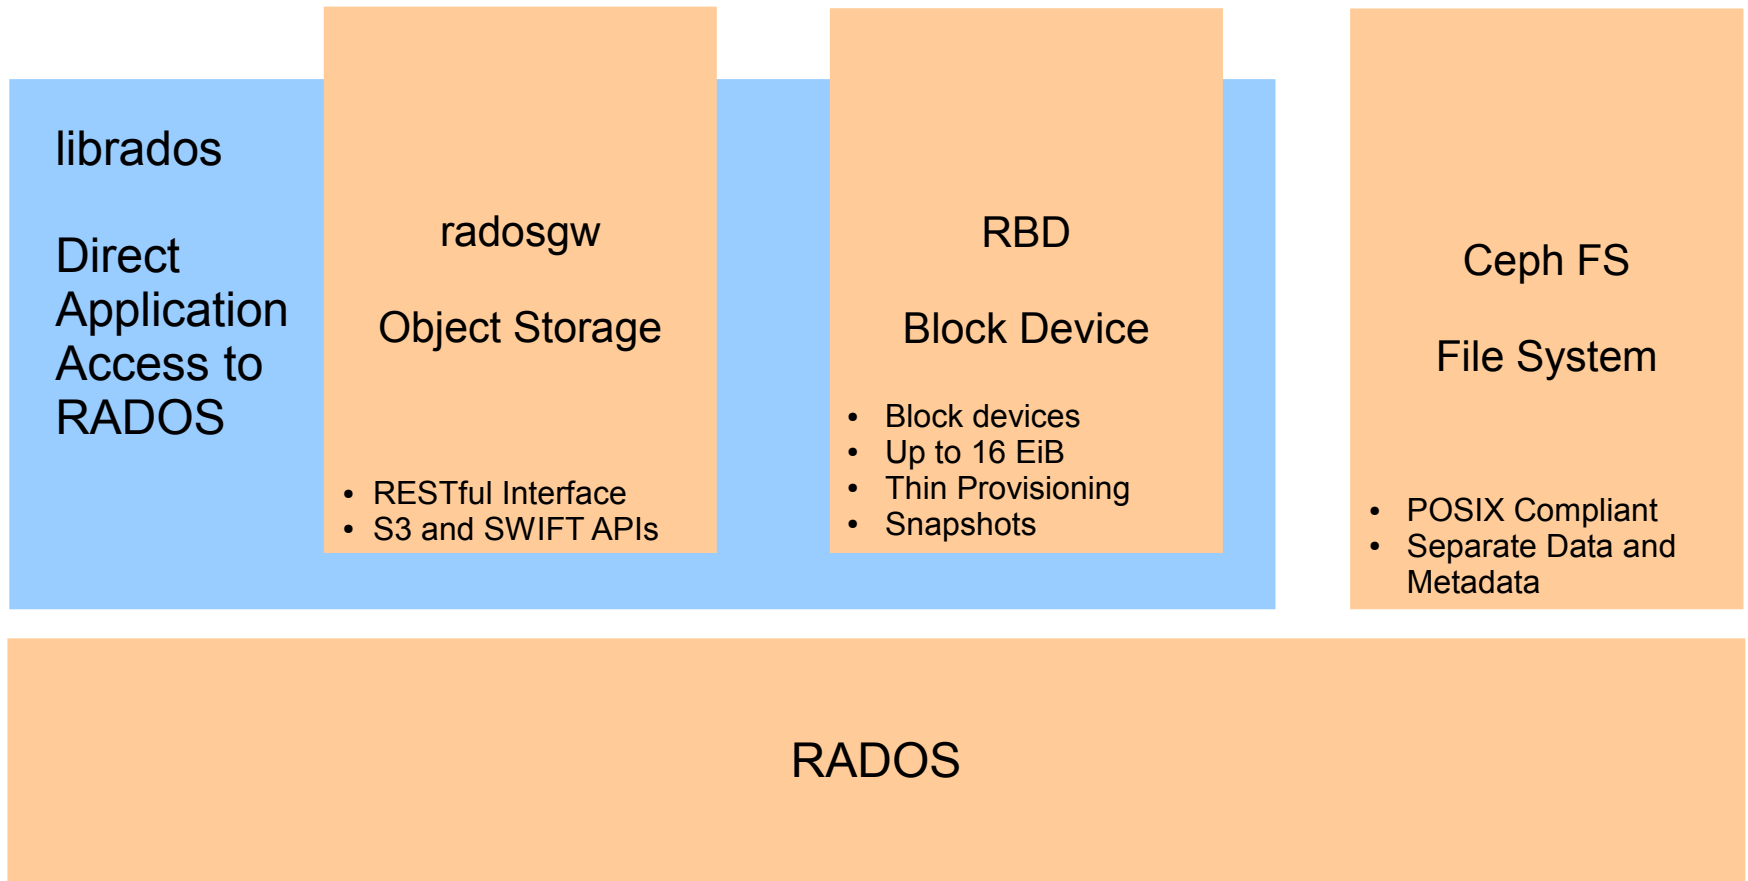
- Three interfaces rolled into one
 - Object Access (like Amazon S3)
 - Block Access
 - (Distributed File System)
- Sitting on top of a Storage Cluster
 - Self Healing
 - Self Managed
 - No Bottlenecks

Ceph architecture



<http://docs.openstack.org/juno/config-reference/content/ceph-rados.html>

Ceph Technical Overview



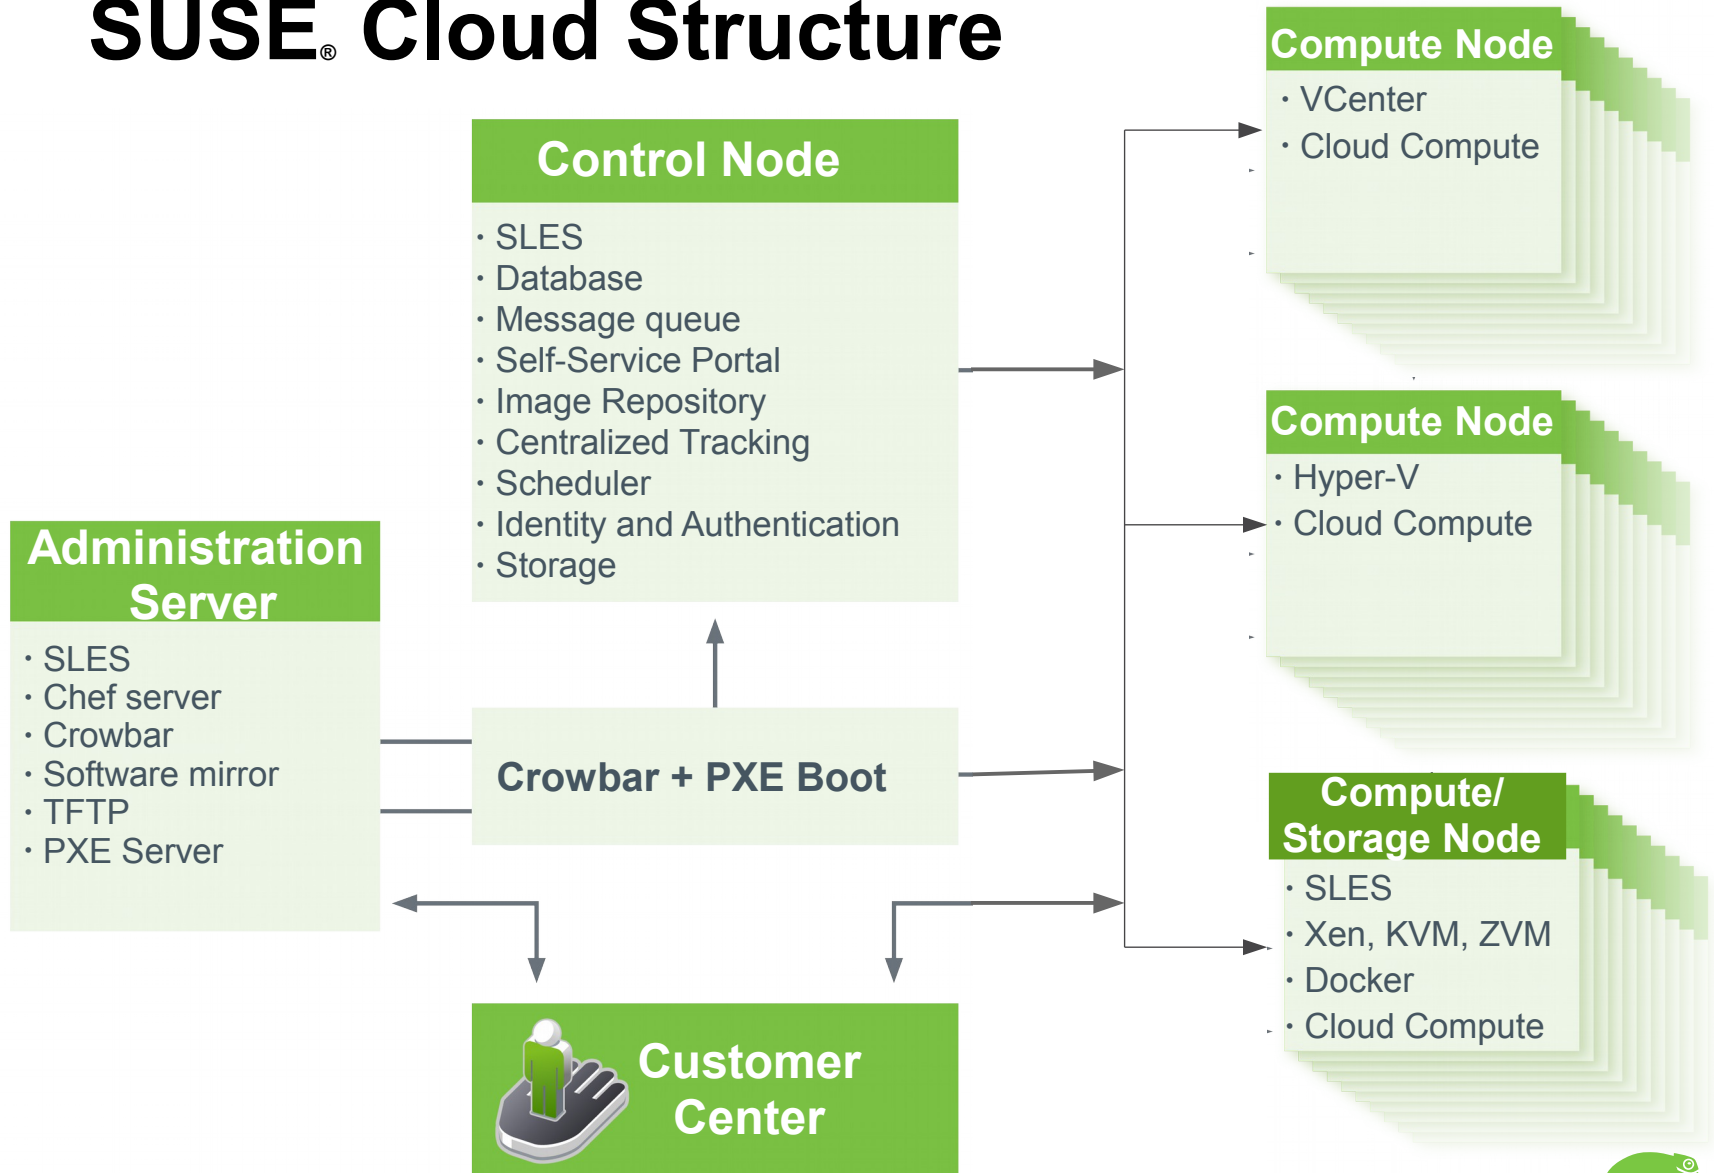
Brief introduction to OpenStack

Project History




- Provide components for Infrastructure-as-a-Service
- Started by Rackspace and NASA in July 2010
- Currently used for example by CERN
- Today: more than 500 companies involved in the OpenStack ecosystem (including SUSE)
- Eleven releases so far (Austin, Bexar, Cactus, Diablo, Essex, Folsom, Grizzly, Havana, Icehouse, Juno, Kilo)
- Next release: Liberty
- Under development: Mitaka

SUSE® Cloud Structure

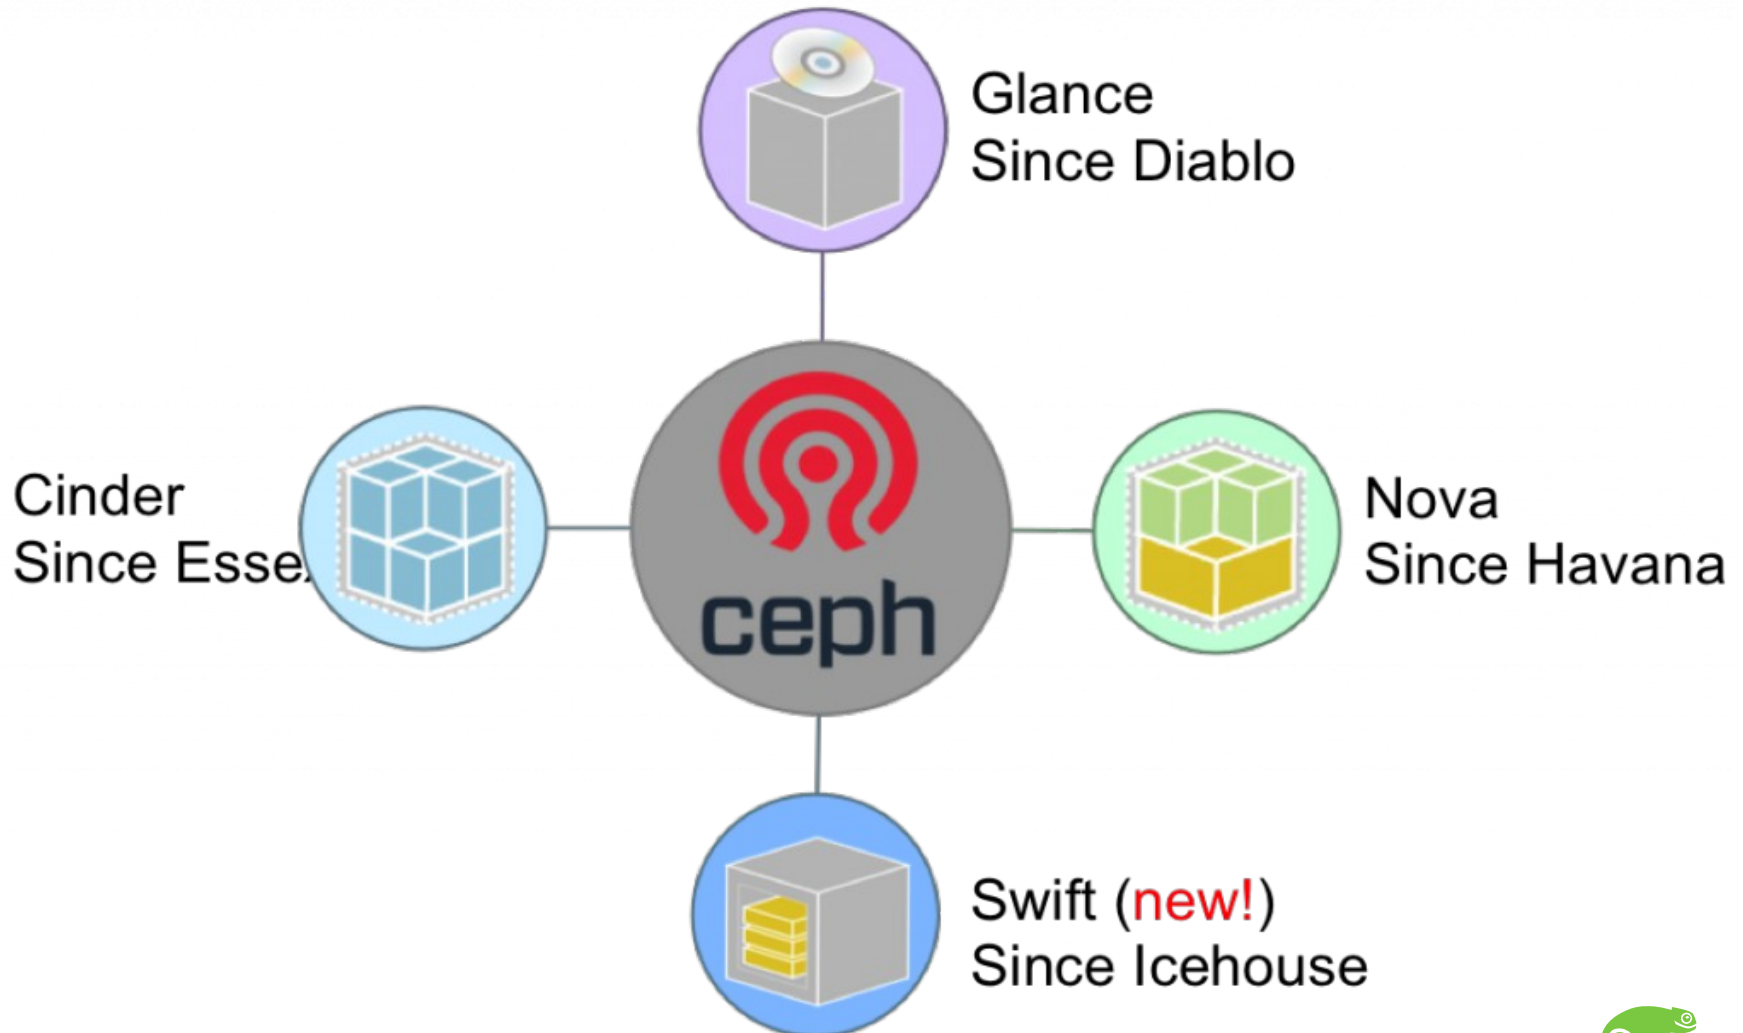


OpenStack and Ceph integration

OpenStack and Ceph winning pair



Ceph and OpenStack ecosystem



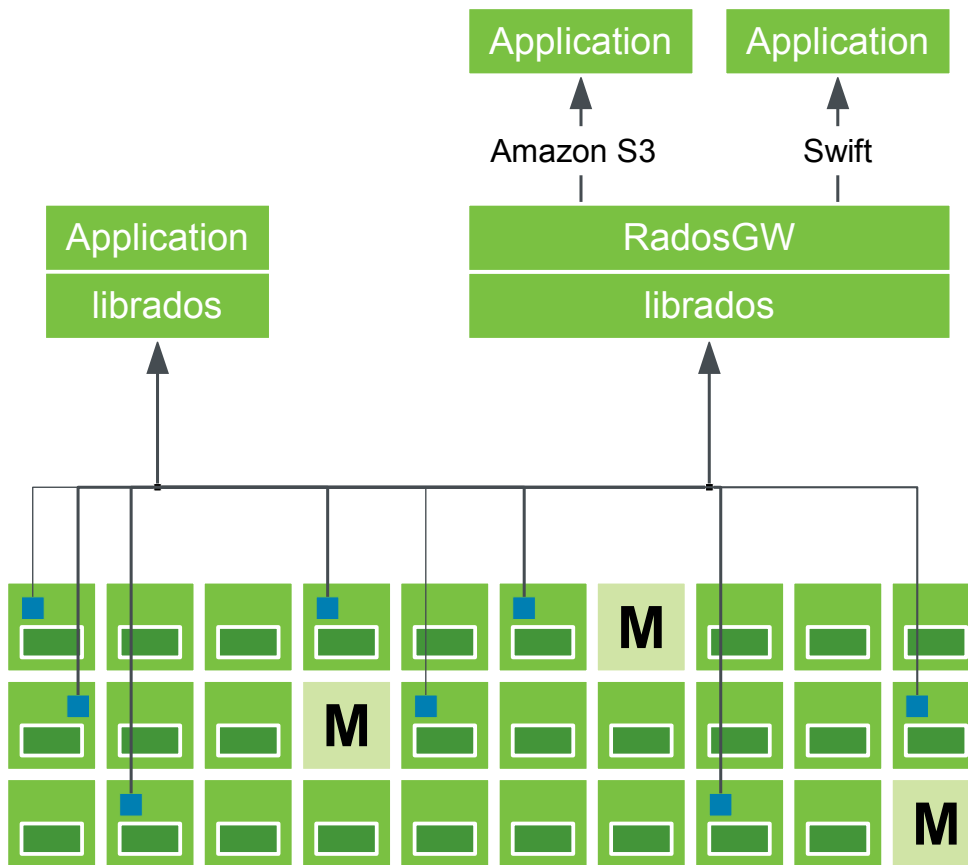
<http://blogs.rdo-project.org/7056/back-from-the-juno-summit-ceph-integration-into-openstack>



Ceph Object Gateway

- HTTP interfaces for RADOS
 - S3-compatible – subset of Amazon S3 RESTful API
 - Swift-compatible – subset of the OpenStack Swift API
- Ceph Object Gateway daemon (radosgw)
 - Apache2 with FastCGI module
 - Embedded CivetWeb webserver from Hammer release
- Swift and S3 use the same HTTP methods but differ in their use of HTML headers

SUSE Storage: RadosGW



librados features:

- ▶ C, C++, Java, Python, Ruby, PHP, etc... bindings

RadosGW features

- ▶ Amazon S3 compatibility
- ▶ OpenStack Swift compatibility

OpenStack Glance and Ceph

- Supported since OpenStack Diablo release
 - Implemented by RBD driver
 - Can be used also by Swift interface
- OpenStack Glance Images
 - Create and store
 - Upload and download
 - Status and update
 - Resizable images
 - Revert to snapshots

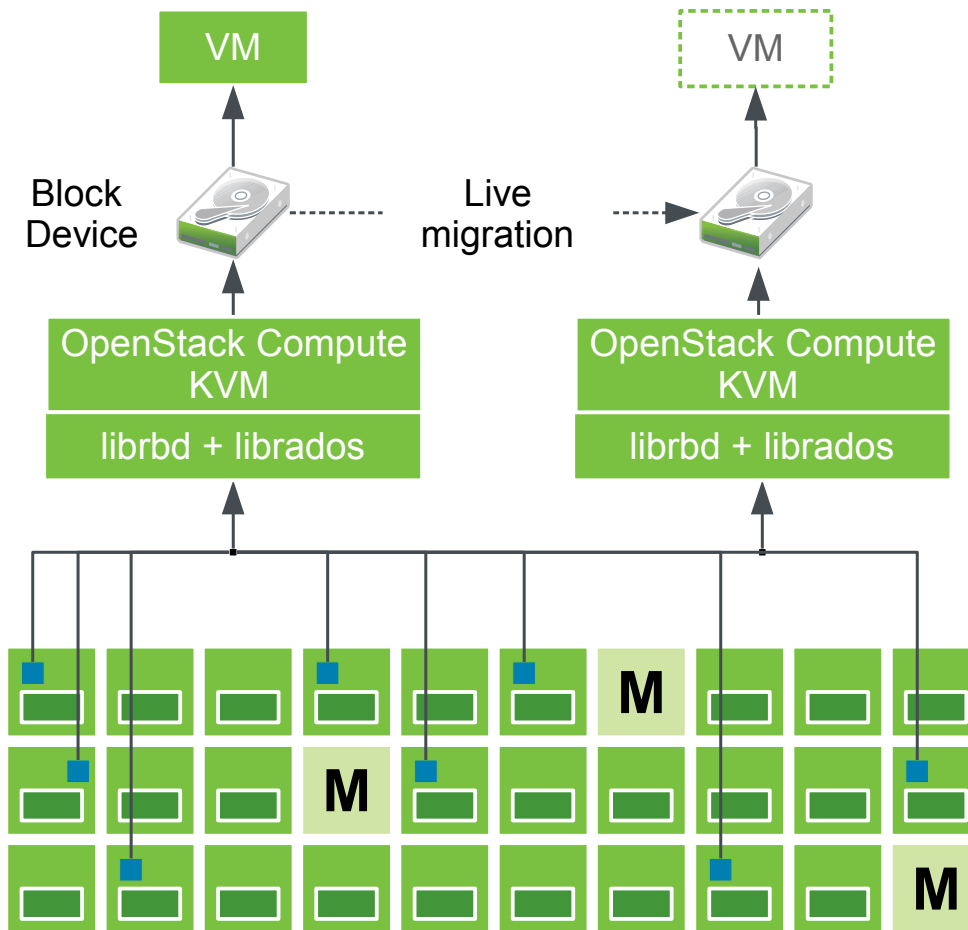
OpenStack Cinder and Ceph

- Supported since OpenStack Essex release
 - Implemented by RBD driver
 - Multi RADOS backend
- OpenStack Cinder
 - Create volumes in RADOS
 - Create boot volumes from images
 - Online attaching and detaching volumes
 - Resizing volumes
 - Copy-On-Write cloning of images to volumes and instances

OpenStack Nova and Ceph

- RBD driver for OpenStack make libvirt configure the QEMU interface to libvirt
- Can be accelerated by enabling RBD cache
- Full support for Live-migration and Evacuate
- Ephemeral storage support in Nova

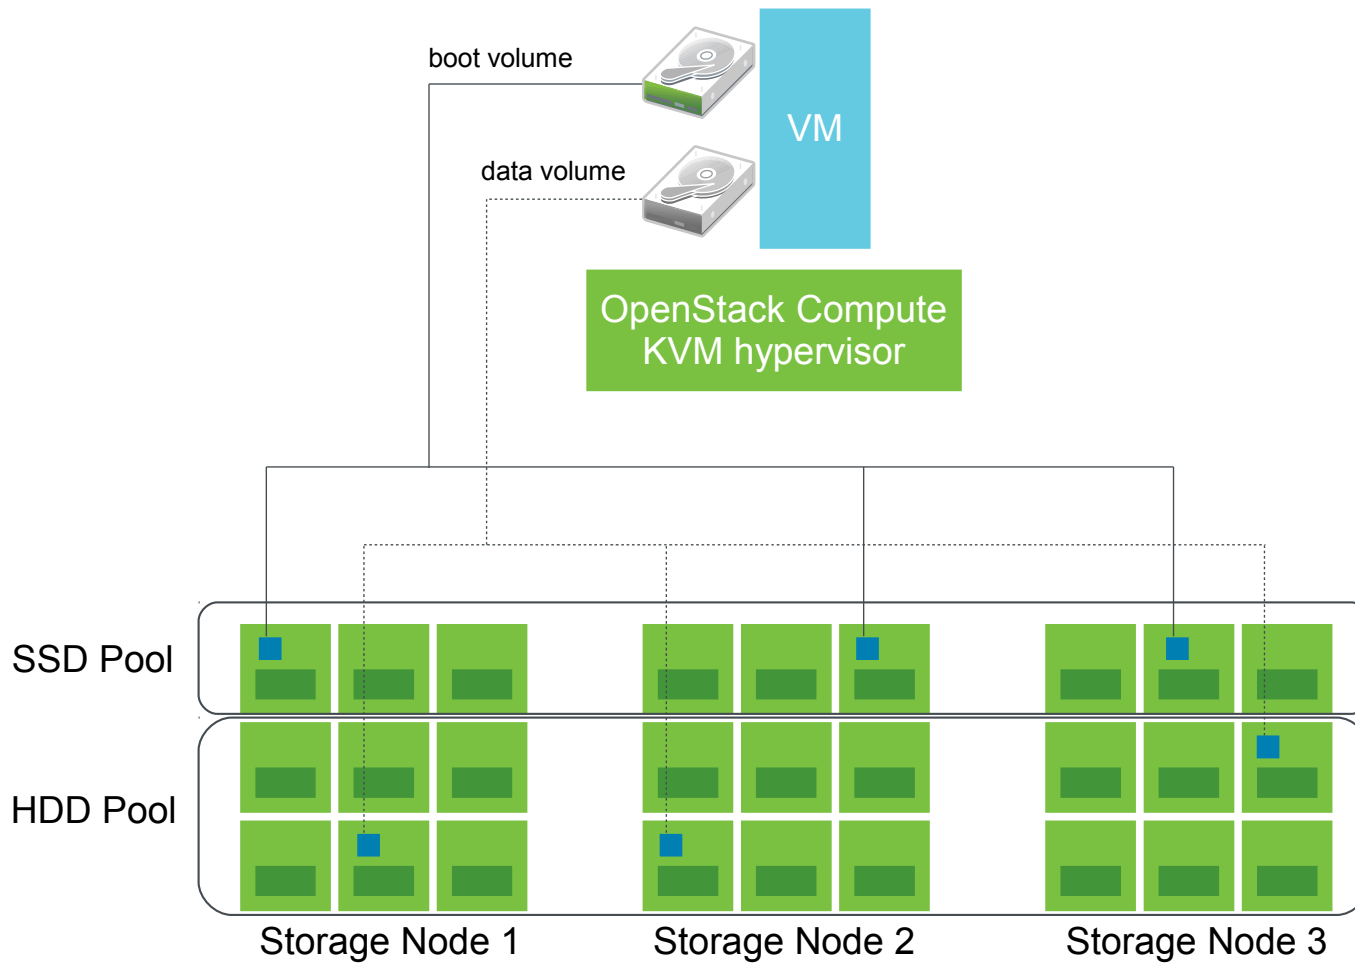
SUSE OpenStack and Storage



Features:

- Thinly Provisioned
- Resizable images
- Image import/export
- Image copy or rename
- Read-only snapshots
- Revert to snapshots
- Ability to mount with Linux or QEMU KVM clients!

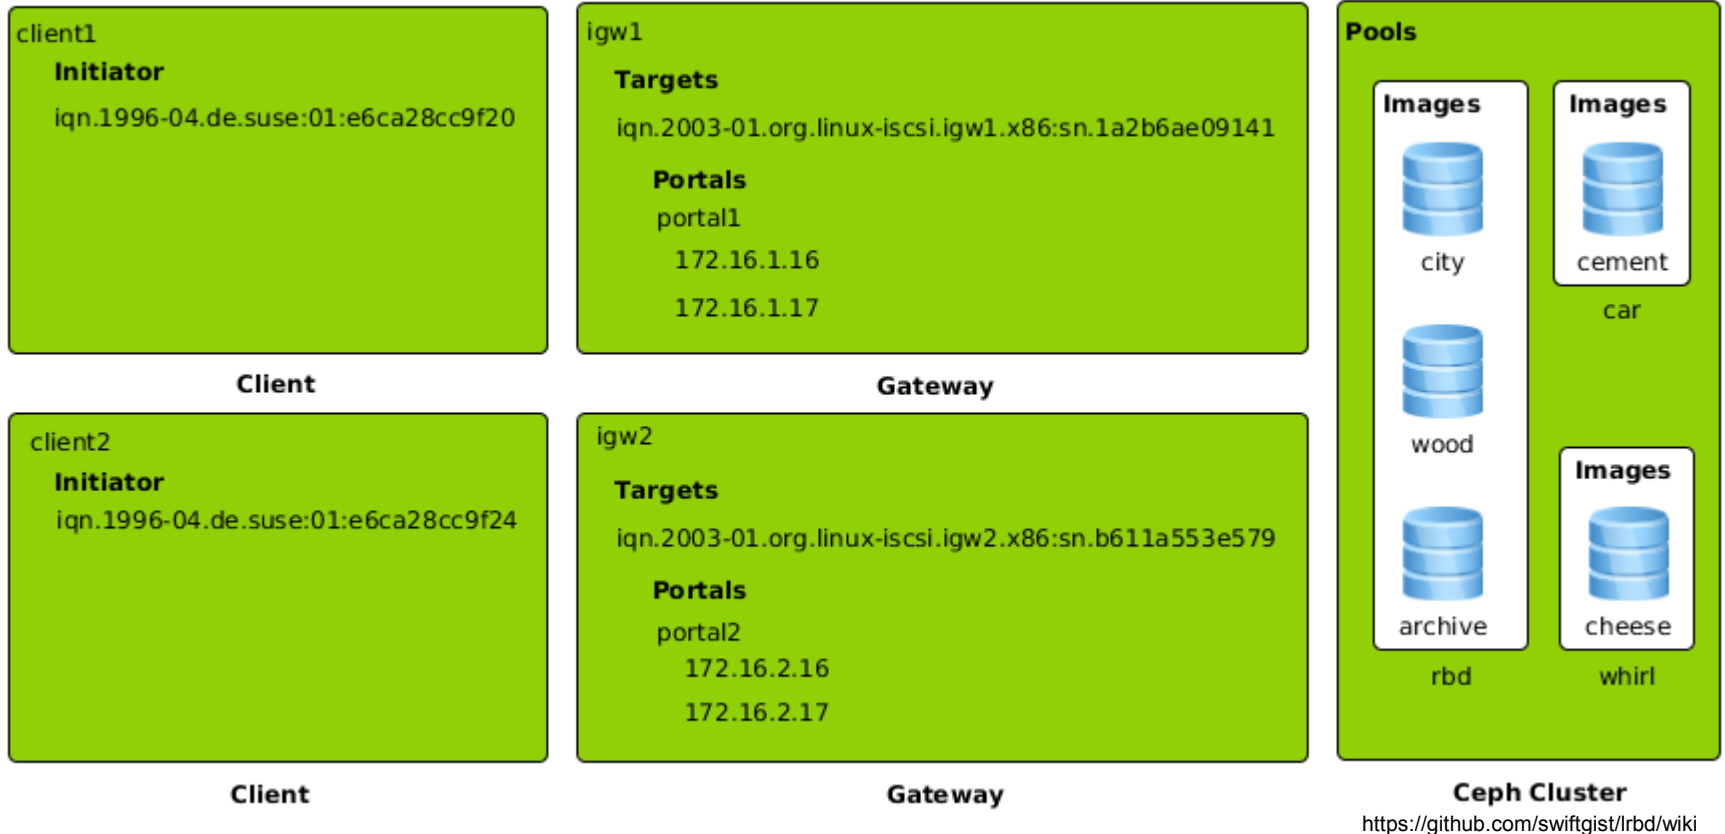
OpenStack multi-backend support



iSCSI RBD Gateway

- Support for RBD in STGT
 - Linux SCSI target framework
 - iSCSI target implementation in userland
 - RBD support implemented in kernel
- Perfect solution for hypervisors without librados support
 - VMWare ESXi
 - Hyper-V

iSCSI RBDGW - Configuration



Live demo

Questions ?



Corporate Headquarters
Maxfeldstrasse 5
90409 Nuremberg
Germany

+49 911 740 53 0 (Worldwide)
www.suse.com

Join us on:
www.opensuse.org

Unpublished Work of SUSE LLC. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE LLC. Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

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 Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

