

SUSE® OpenStack Cloud 6 and Beyond

Roadmap and Priorities

Pete Chadwick

Senior Product Manager

pchadwick@suse.com



SUSE and OpenStack

2011	2012	2013	2014	2015



SUSE OpenStack Cloud 5

- Juno + SUSE Linux Enterprise 12
- SUSE Enterprise Storage

SUSE OpenStack Cloud 3 + 4

- Havana, Icehouse
- HA Control Cluster

SUSE OpenStack Cloud 2

- Grizzly
- Mixed hypervisor clouds

SUSE participates in creation of OpenStack Foundation

- Platinum Member
- Alan Clark elected Board Chairman

SUSE OpenStack Cloud 1

- 1st Enterprise supported OpenStack distribution
- OpenStack Essex release
- Focus on: Security and Improved Xen supports
- Deployment framework

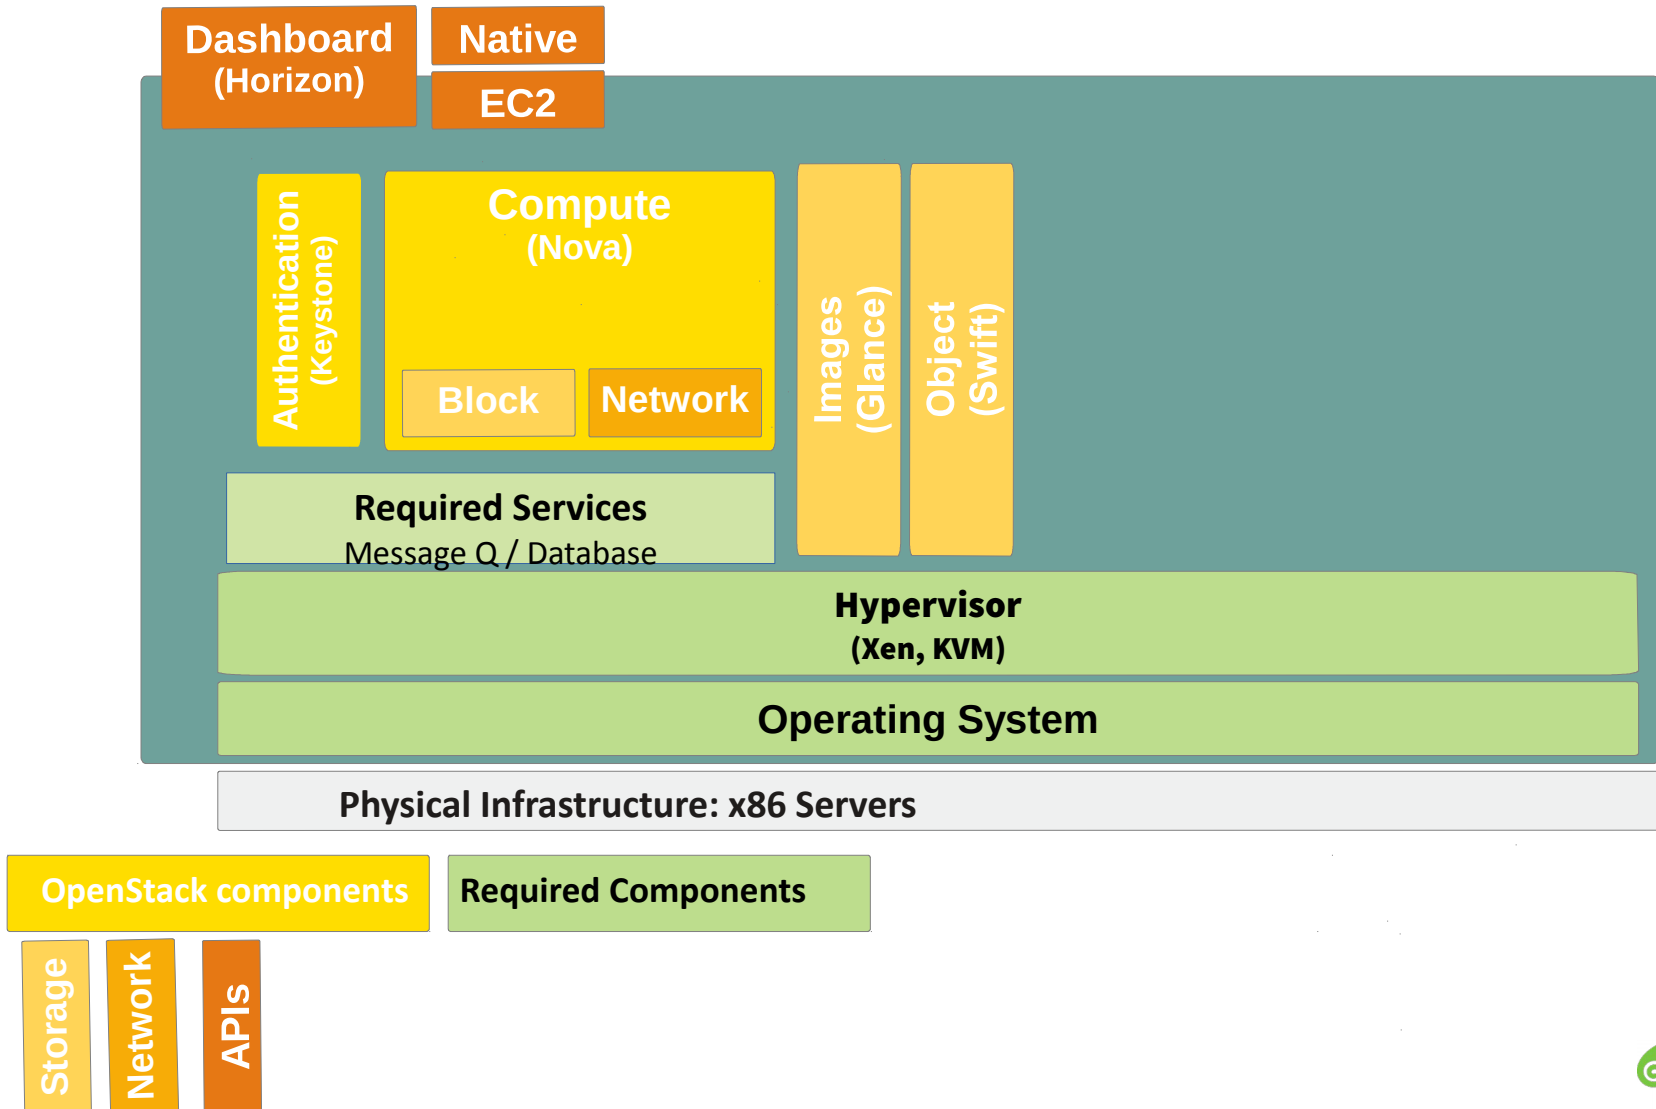
SUSE Joins OpenStack community

- Diablo based cloud appliance

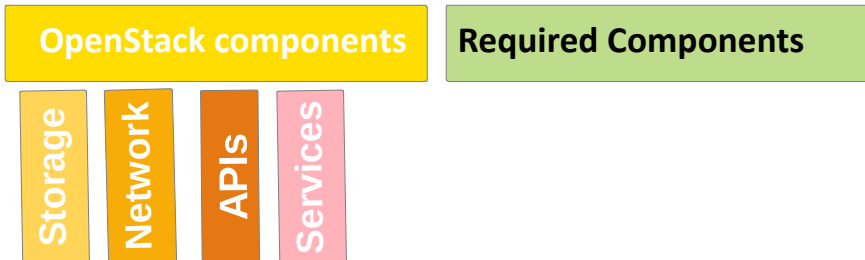
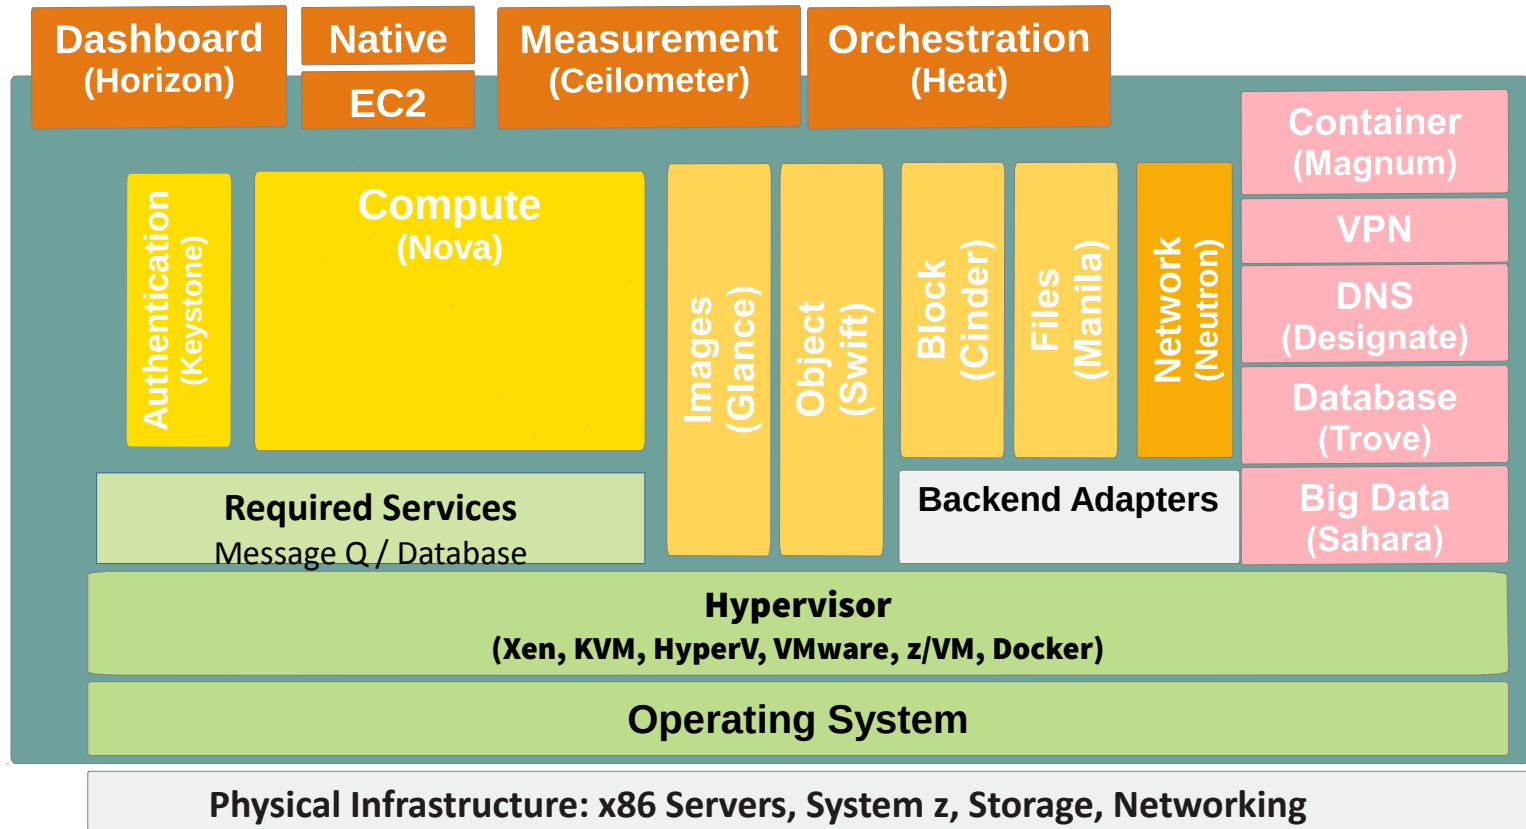
SUSE OpenStack Cloud Priorities

- Deliver Enterprise Stability
 - Fewer upgrades
 - Longer support cycle
 - HA compute nodes
 - Rolling release upgrades
- Stay aligned with upstream releases
 - Continued support of DefCore standards
 - High demand projects

OpenStack Then ... 2012



OpenStack Now



Theme Definitions

Scalability

Items that will impact the scale at which the service can operate

Resiliency

Items that will impact the high availability or ability to recover from failures for the service

Manageability

Items that improve the User Experience (UX), promote operational ease-of-use, or enhance the capabilities of the service

Modularity

Changes that enhance the modularity of the service architecture and usually result in a more manageable code-base and/or remove duplicity of efforts

Interoperability

Items that enable the service to operate across multiple OpenStack clouds [federation], promote a common experience across separate OpenStack-Powered clouds [interop], or add dependency on another OpenStack service [service dependency], and/or backwards compatibility [compatibility]





openstack
CLOUD SOFTWARE

	Scalability Increase scale			Resiliency Availability or Durability			Manageability Operations and UX			Modularity Service/Component Modularity			Interoperability Interop, Federation, Compatibility		
	L	M	N	L	M	N	L	M	N	L	M	N	L	M	N
Ceilome	■	■	■	□	■	□	■	□	■	■	■	□	□	□	
Cinder	□	■	■	□	■	■	■	□	■	■	■	■	■	■	
Designate	■	■	■	□	□	■	■	■	■	■	□	■	■	■	
Glance	■	□	■	■	■	■	■	□	■	□	■	□	■	□	
Heat	■	■	■	■	■	□	■	■	□	□	□	□	■	■	
Horizon	■	■	□	□	□	□	■	■	□	■	■	■	■	□	
Ironic	□	□	■	□	■	□	■	□	■	■	■	□	□	□	
Keysto	■	■	■	□	□	■	■	■	■	■	■	■	■	■	
Kolla	■	■	■	■	□	□	■	■	■	■	■	□	■	■	
Kuryr	□	■	□	□	□	□	■	■	■	■	■	■	■	□	
Magnus	□	□	□	■	■	□	■	□	□	□	■	■	□	□	
Manila	□	□	■	■	■	■	■	■	■	■	□	■	■	■	
Neutron	□	□	■	□	□	□	■	■	■	■	■	□	■	□	
Novus	■	■	■	■	■	■	■	■	■	■	■	■	□	■	
Oslo	■	■	■	■	■	■	■	□	■	■	■	■	■	■	
Sahara	■	□	□	■	■	■	■	□	■	□	■	■	■	□	
Swift	■	■	■	□	■	■	■	■	■	□	□	□	□	□	
Triple O	■	□	□	■	■	■	■	■	■	■	□	■	□	□	
Trove	□	□	□	■	■	■	■	■	□	■	□	□	□	□	

■ Planned Work

□ No Information Provided

CLOUD SOFTWARE

SUSE OpenStack Cloud 6

- Platform Upgrades
 - OpenStack Liberty
 - SUSE Linux Enterprise Server 12
 - Installation Framework
- New Feature Highlights
 - New Compute Options
 - Docker
 - z/VM
 - Full support for Shared File Service
 - Extend HA to Xen/KVM Compute Nodes

Cielometer

- Resource metadata caching to reduce query load on Nova API
- Functional test improvements
- Declarative notifications with generic hardware sensor plug-in and generic snmp pollster, to facilitate easy addition of event exchanges
- Dispatcher for loose integration with Gnocchi (time-series DB service)
- Splitting of alarming code into separate repository

Cinder

- Image caching allows back-end to perform clone vs. copy over network
- Perform backups without detaching volumes
- Cloning support for Consistency Groups
- Nested Quotas Support
- Beginning work to separate Cinder Backup service
- Initiator/Target code from Nova/Cinder moved to os-brick
- Beginning to examine Cinder V2 API Capabilities for DefCore

Designate

- Secondary DNS Zones
- Ceilometer Events
- Federated Designate support
- Sharding of worker domains across multiple workers
- Active/Passive failover for designate-pool-manager
- v2 API

Glance

- Middleware support for server health
- Image signing and verification
- Re-use deleted image-member before creating a new one
- EXPERIMENTAL: V3 Artifacts API (generic data asset support)
- Default version for glanceclient is now Image V2 API
- HTTP Proxy Support for Glance S3 Driver

Heat

- 17 new resources from big tent projects aimed at users
- 9 new resources aimed at operators using heat to manage their cloud
- Conditional resources exposure
- REST/Client improvements including stack tags and preview
- Improved deprecation policy

Horizon

- Scalability improvement for Network Topology View (graph-based view)
- Theming support
- Plugin Architecture with Angular.js

Ironic

- Changed to SemVer/Independent Release (faster releases)
- Improvements to the driver API
- Added a new “enroll” state for nodes. Allows performing additional tasks before node is exposed to Nova.
- New State Machine (New processed: cleaning and inspection)
- Vendors have added new drivers and improved existing ones (AMT, iRMC, VirtualBox (testing only), iLO driver enhanced)

Keystone

- Dedicated Keystone auth library
- Stable driver interfaces
- Fernet token bug fixing
- Federation improvements
- Improve functional test suite and test coverage
- Streamline code base, remove unused code
- Configuration options for storing to SQL

Manila

- Expand/shrink share
- Consistency groups
- Oversubscription
- Mount automation
- Microversioned API
- Optional snapshots
- New Drivers: Hitachi HNAS, Windows SMB, Gluster Native

Neutron

- LBaaS reference implementation based on Octavia
- Pluggable IPAM with reference implementation
- QoS API
- RBAC for networks
- Improved quota enforcement
- Neutron ironic integration

Nova

- Scalability – Cells v2 is building our next generation scale out system (adding some initial supporting infrastructure)
- Resiliency – Improving migrations, resource tracking etc.
- Manageability – Including improved logging, host down handling and experimental support for running API services in Apache
- Modularity – Strongly versioned API between scheduler and rest of Nova, adopting Cinder's os-brick
- Interoperability – API V2.1 efforts, making progress towards Python3

OpenStack CLI

- Volume V2 API Support
- Image V2 API Support
- Cloud Configuration File Support

Sahara

- New plugins/supported versions
- Extra functionality (including HA) for CDH & HDP
- Bare-Metal (Ironic) Deployment Support
- Moved to using Heat as provisioning engine

Swift

- 1+ object server per disk
- async container listing updates
- Improvements in ring building to limit the amount of data moved in certain conditions
- Ring-builder-analyzer
- Bulk upload feature

Trove

Project Snapshot Database Service

of Contributors (Liberty):
47

of Companies (Liberty):
17

Liberty ([16](#) specs/blueprints)

- Support for Highly available MySQL clusters (Percona XtraDB/Galera Replication)
- Enabling Redis and Vertica clusters
- Support MongoDB user and database creation
- Better support for Fedora and Redhat-based linux variant.
- Refactor MySQL datastore and support MariaDB
- New UI panel added to Horizon
- New datastore management API for operators

Mitaka ([1](#) specs/blueprints as of 10/13/2015)

- Self-healing cluster for MySQL
- Cassandra clusters and enhancements

“N” Release

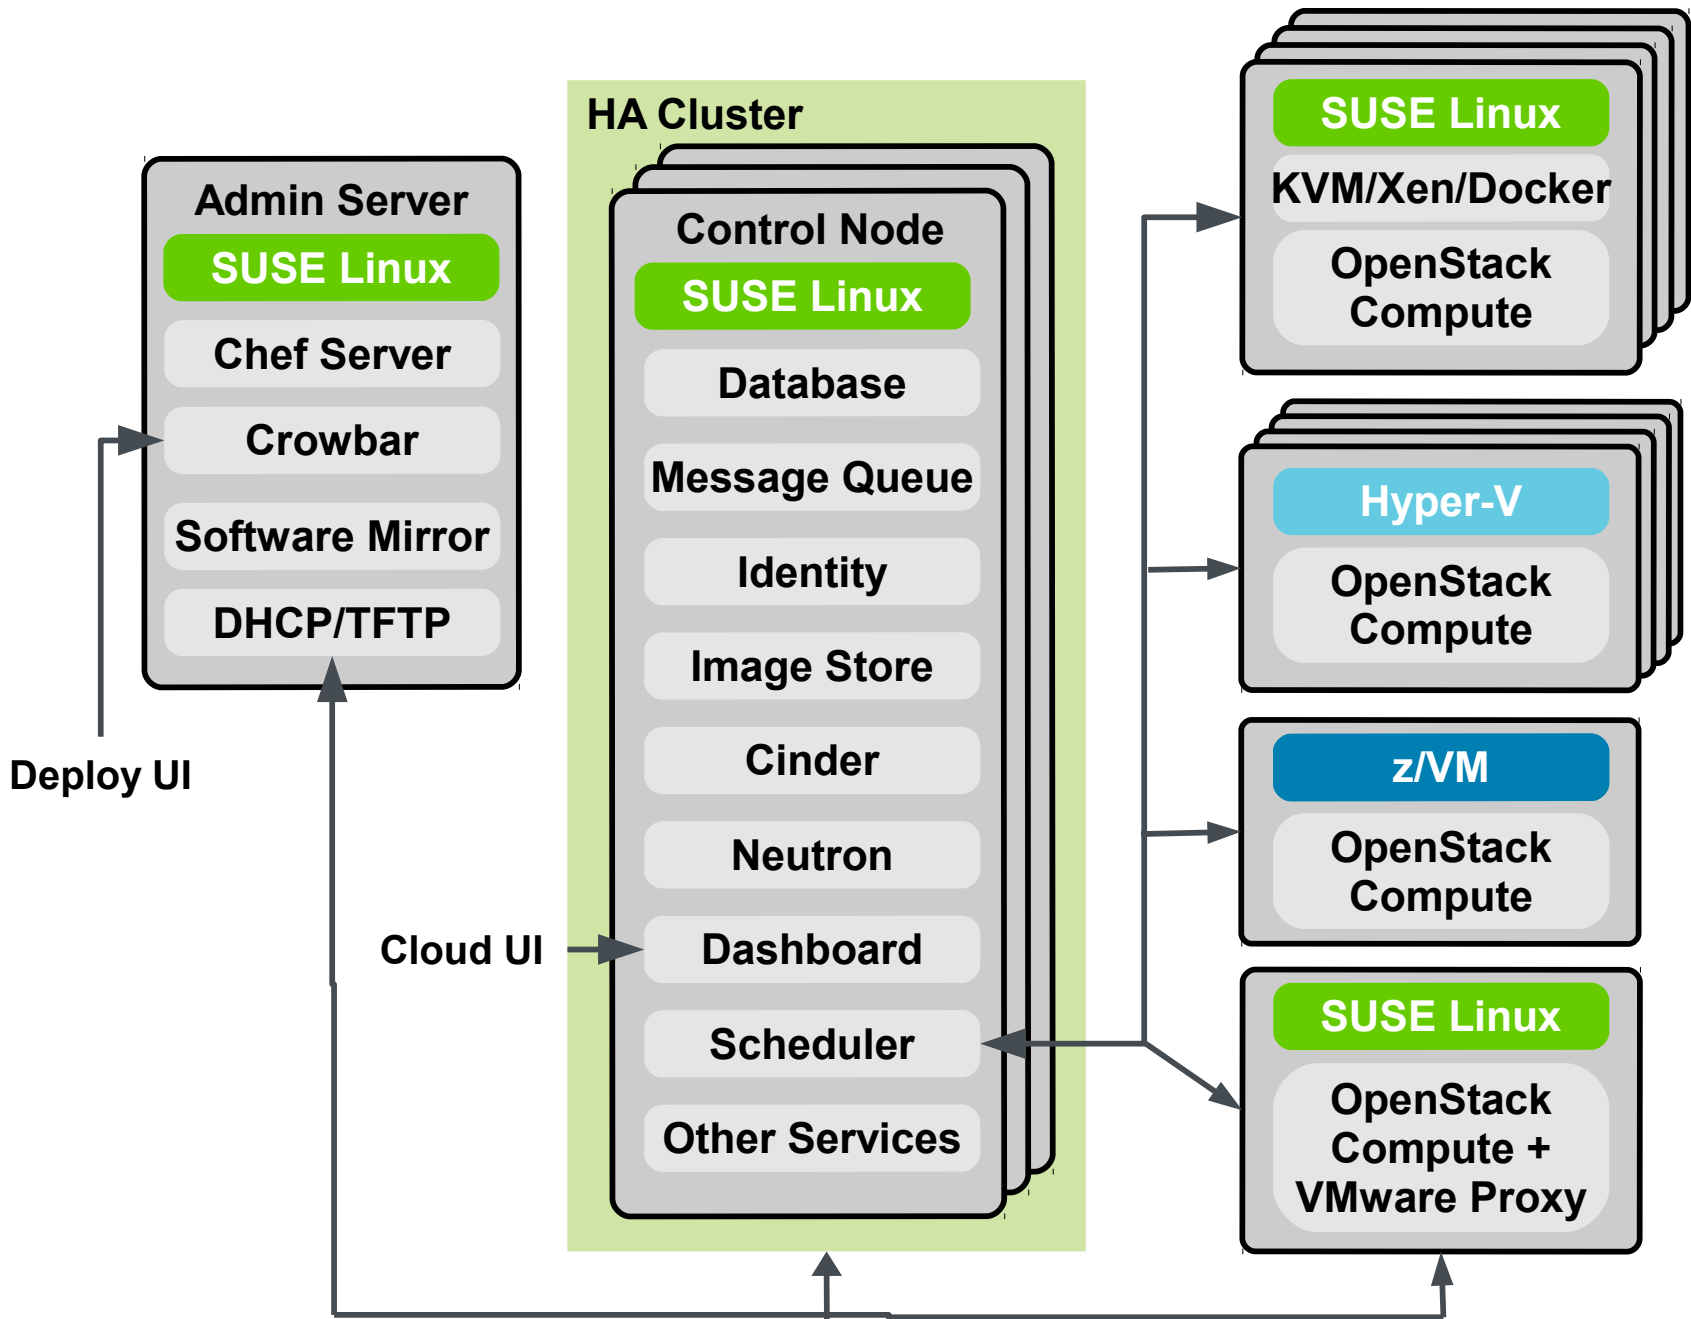
- TBD, themes are not prevalent yet either

Install Framework Upgrades

Install Framework Upgrades

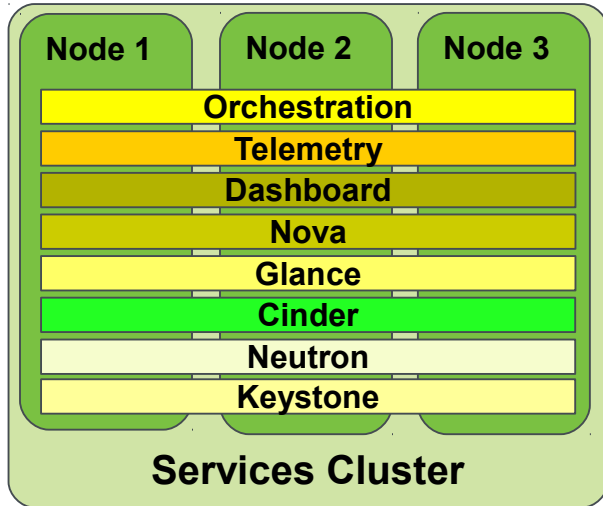
- Increased scalability
- Configuration support:
 - Shared File Service
 - Docker
 - HA Compute Nodes

Mixed Hypervisor Support

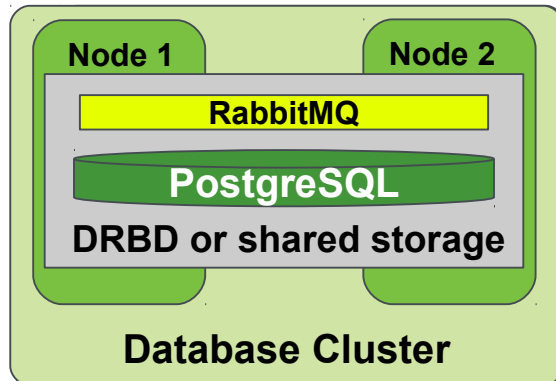


HA Compute Nodes

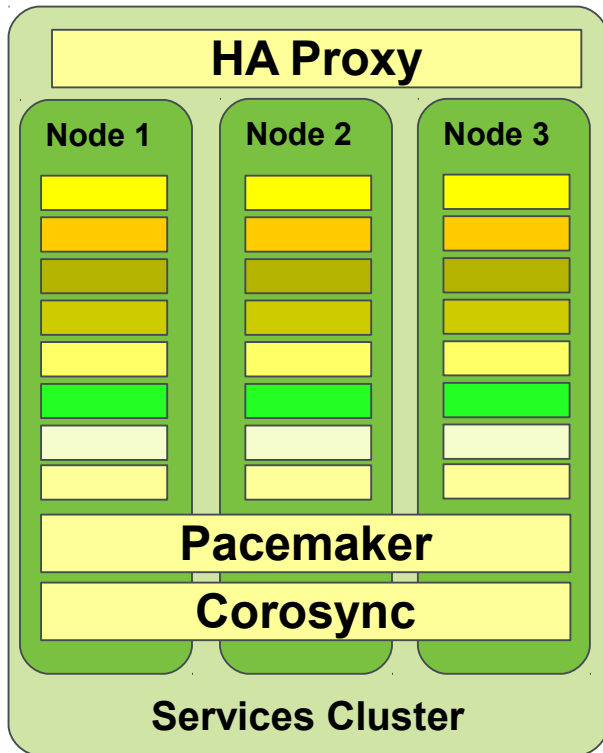
Example HA Control Cluster



- Maintain cloud uptime
- Automatic Restart of Cloud Services
- Active/Active configuration with load balancing



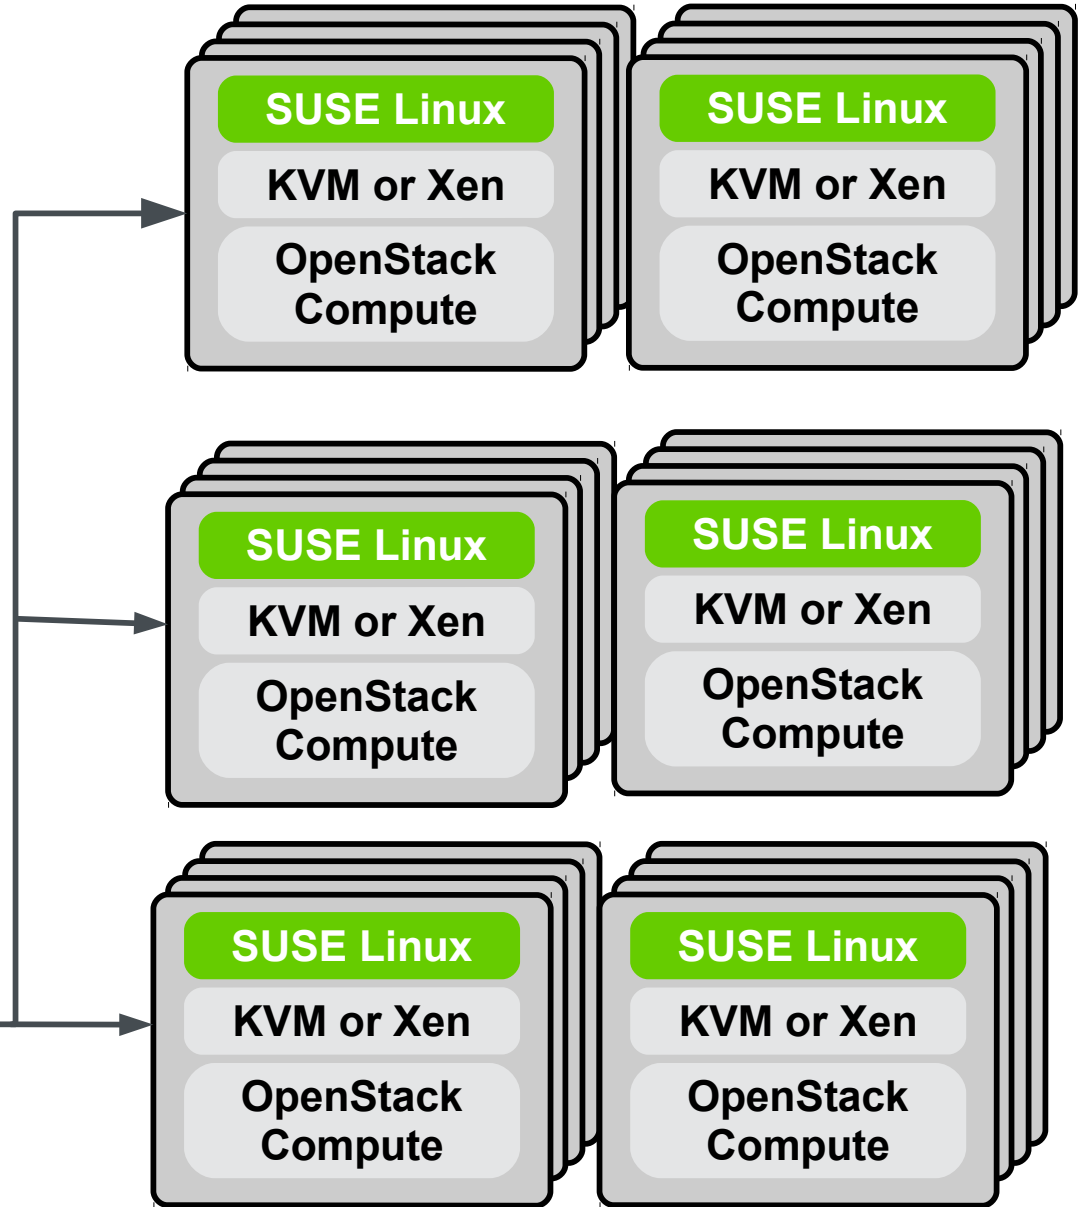
Under the covers



- HA Proxy distributes service requests
- Pacemaker – server and service monitoring and control
- Corosync – membership, messaging, quorum

But what I really want to do is keep my workloads up!

HA Cluster



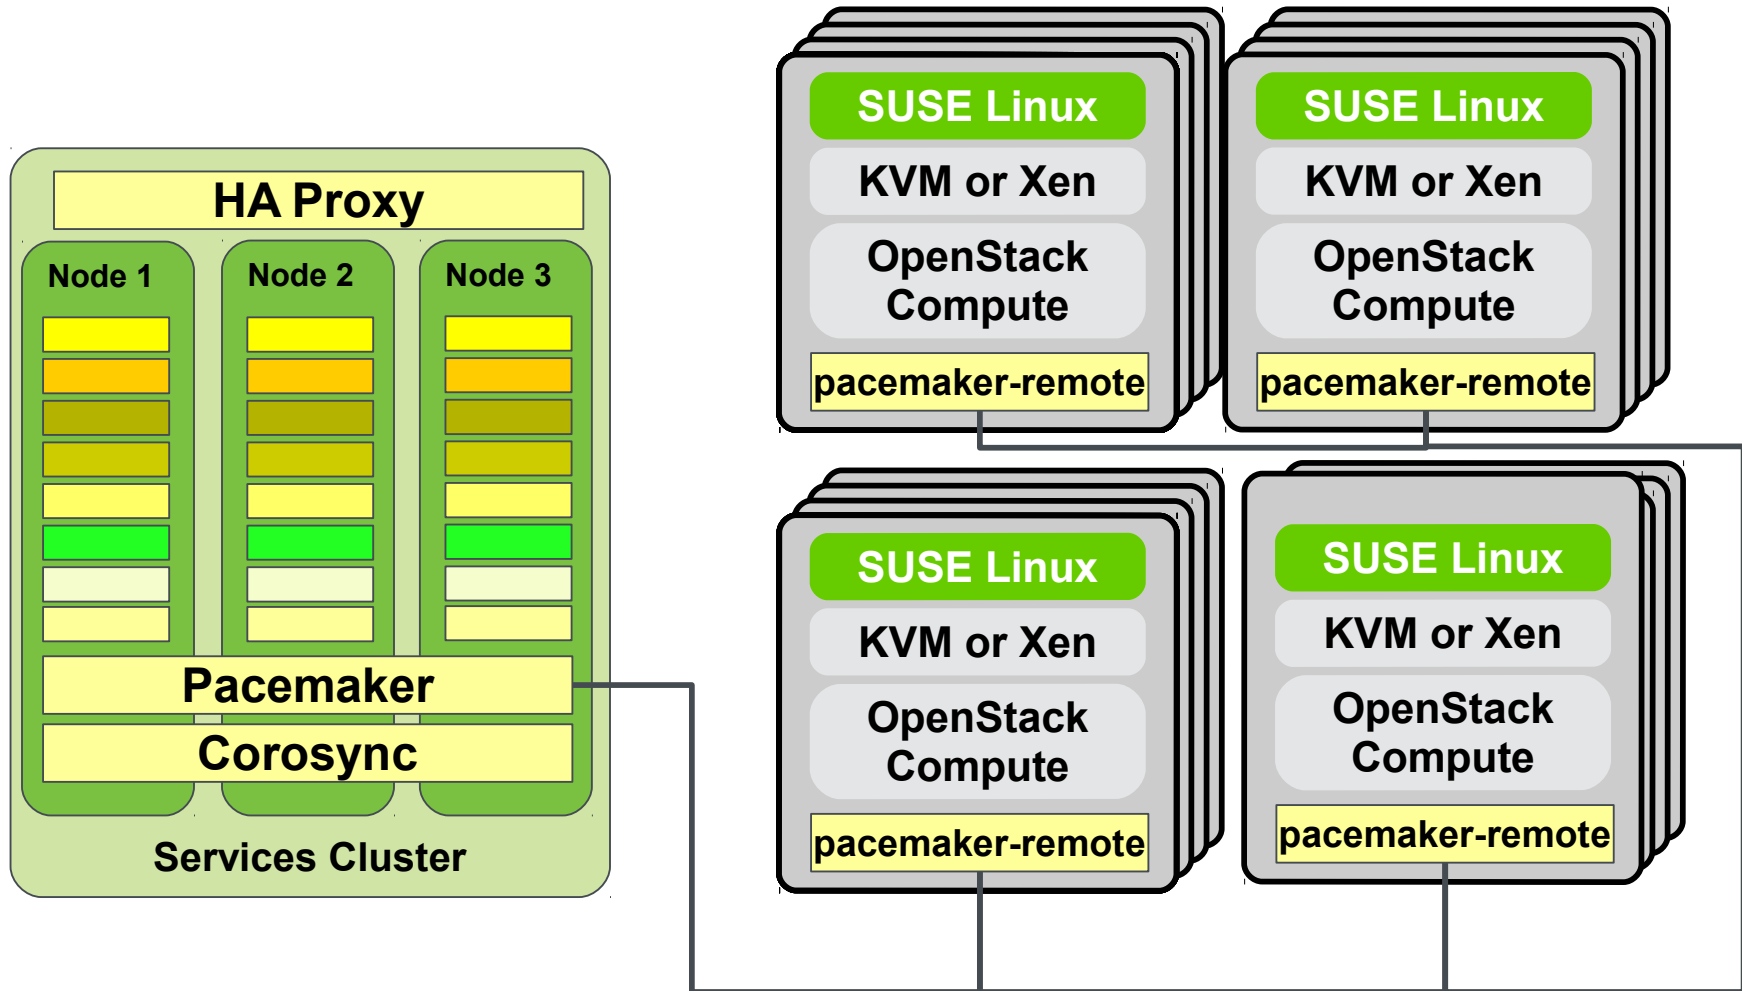
Different Scale

- Corosync limit of 16 nodes - internode communication
- But many more compute nodes
- Alternatives
 - Multiple compute clusters
 - VM level HA
- Neither are easy
 - Complex set up
 - Difficult to manage
 - Limits configuration options

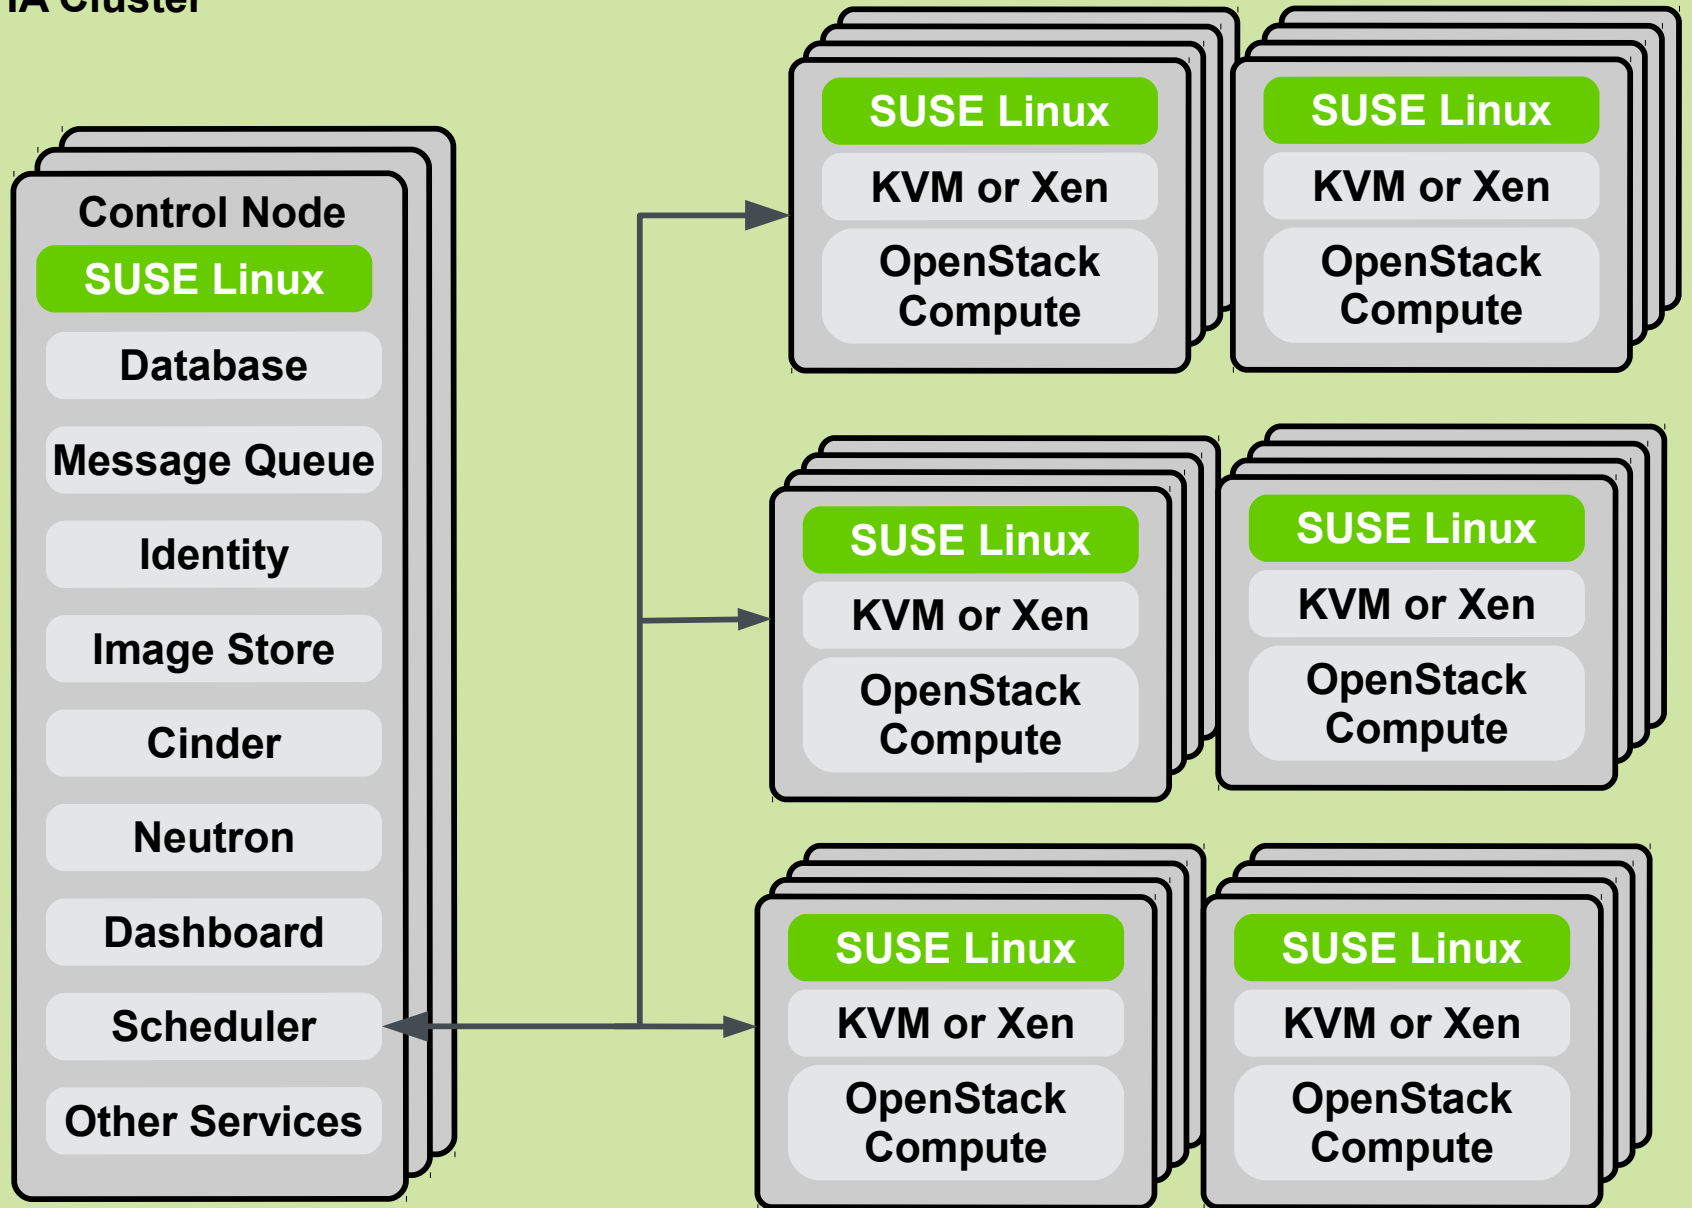
Pacemaker Remote

- Daemon which runs on compute nodes
- Extends an existing Pacemaker Cluster
- All intelligence on the cluster
- No inherent limit on scale

Under the covers



HA Cluster

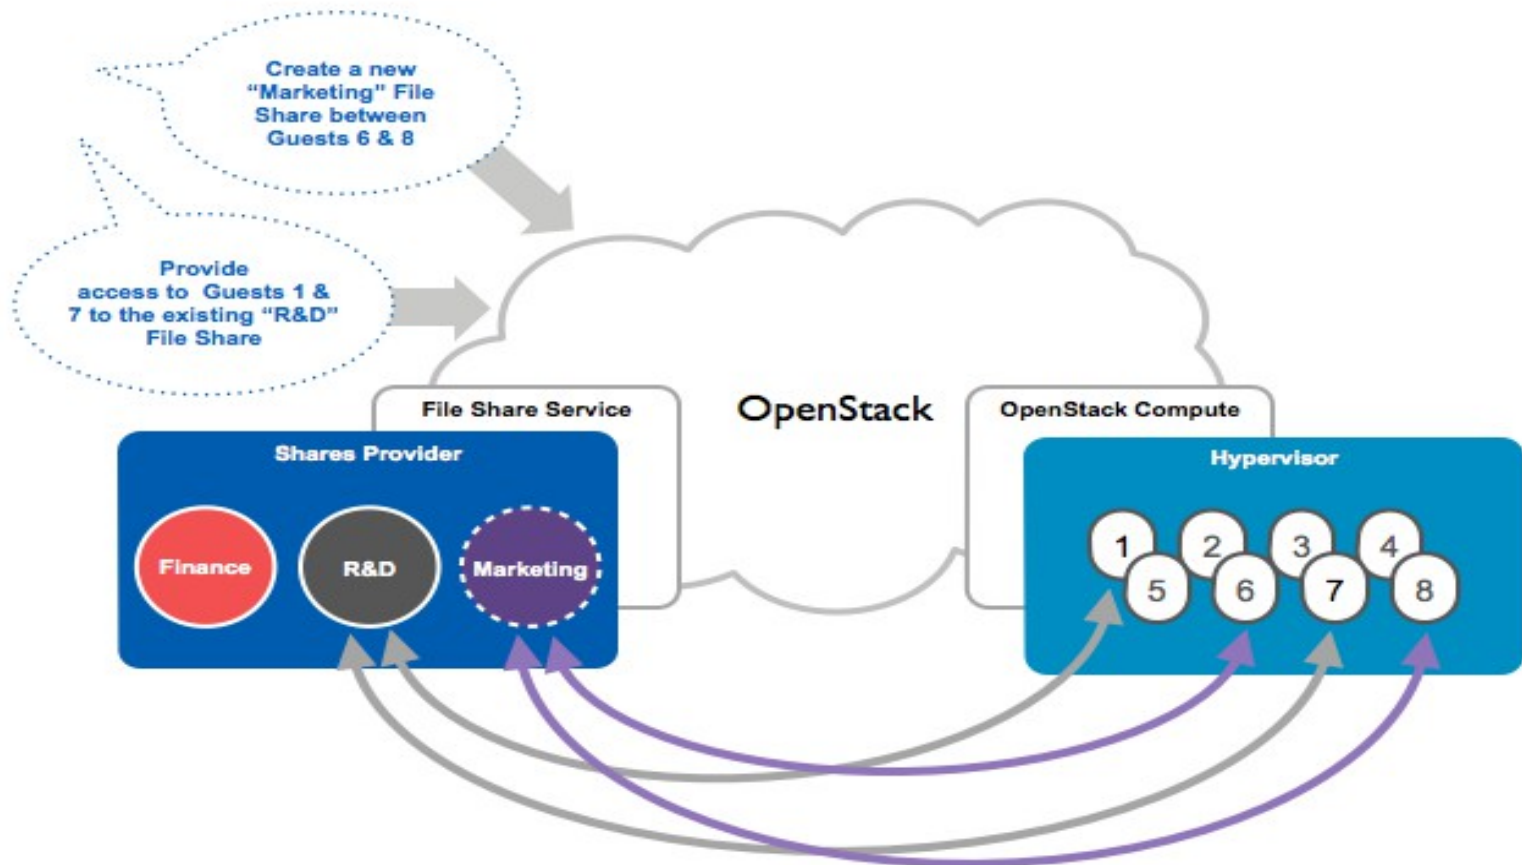


Capabilities

- Increase availability of compute nodes
 - Detect failed nodes
 - Automatic recovery of nodes
- Coordinate with control plane
 - Trigger node evacuation
 - Migrate VMs
- Isolate failing compute nodes
 - STONITH extends to remote nodes

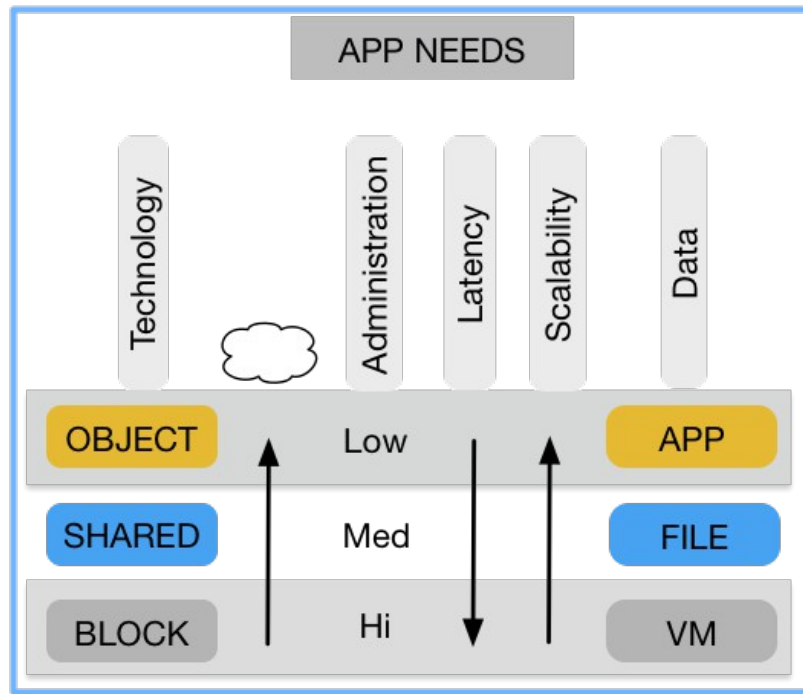
Shared File Service

Manila Shared File Service



Manila Shared File Service

- Lower Cost Heterogeneous Infrastructure
- Application porting - Moving Workloads to OpenStack



Manila in SUSE OpenStack Cloud

- The Manila service is tech preview in SUSE Cloud 5
- Fully supported in SUSE OpenStack Cloud 6
- Crowbar deployment tool integration
 - Controller HA
 - NetApp driver
 - Custom driver possible



Futures

SUSE OpenStack Cloud

	Q4 2015	Q1 2016	Q2 2016	Q3 2016	Q4 2016	Q1 2017
SUSE Cloud	β	GA				7

SUSE OpenStack Cloud 6

- OpenStack Liberty (+ DefCore)
- SUSE Linux Enterprise 12 SP1
- **Docker**
- **Non-disruptive upgrade**
- **High Availability Compute** Nodes
- z/VM Compute

Queued

- Cloud Foundry Integration
- Ironic
- VPNaaS
- z/VM Control

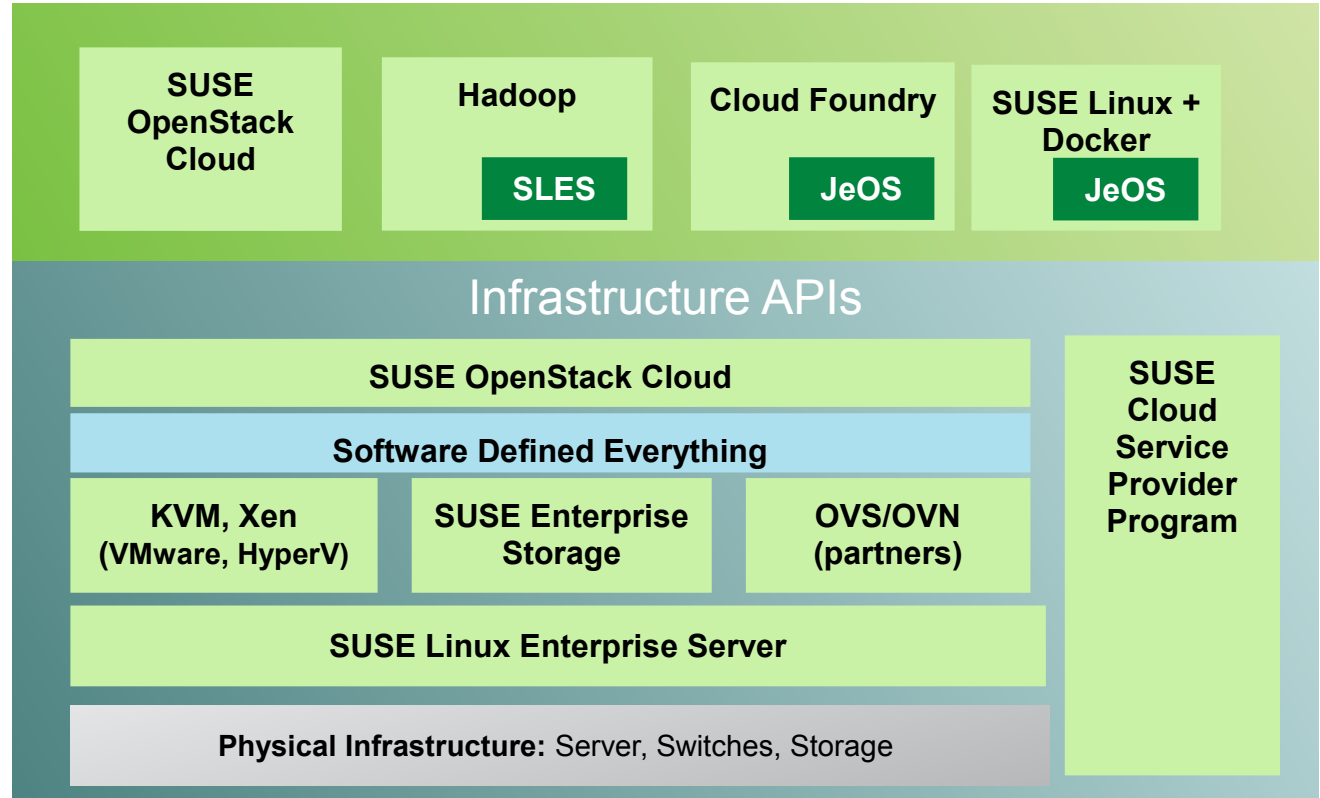
SUSE OpenStack Cloud 7

- OpenStack “N”
- SUSE Linux Enterprise 12 SP2
- **Workload automation**
 - Cloud Foundry
 - Magnum (Docker)
- **Scalability** improvements
- Manila – CephFS integration

SUSE Portfolio for the Future

Management

- SUSE Manager
- SUSE Studio
- Kiwi
- Machinery
- SaltStack
- SUSE OpenStack Cloud





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.

