

The Cloud, the Software-defined Enterprise and OpenStack

Christian Spindeldreher
Enterprise Technologist, Dell



Workshop della Commissione Calcolo e Reti dell'INFN
28.05.2014 | Catania, Italy

Cloud models & approaches

Deployment

Private

Operated and hosted by an enterprise IT department or external provider

- Exclusive use by the organization
- Host applications and IT services within a single-occupant datacenter

Hybrid

Link disparate clouds together

- Connect multiple clouds
- Move workloads between clouds
- Connect data across Clouds

Public

Open to multiple orgs and individual users on a shared basis

- Minimizes initial capital investment
- Blends agility & efficiency with massive scalability
- Computing resources are owned and operated by a third-party provider

Delivery

SaaS

Software as a Service

Target: End users – Collaborative, Engineering & Manufacturing Apps

PaaS

Platform as a Service

Target: Developers – database, middleware & infrastructure software

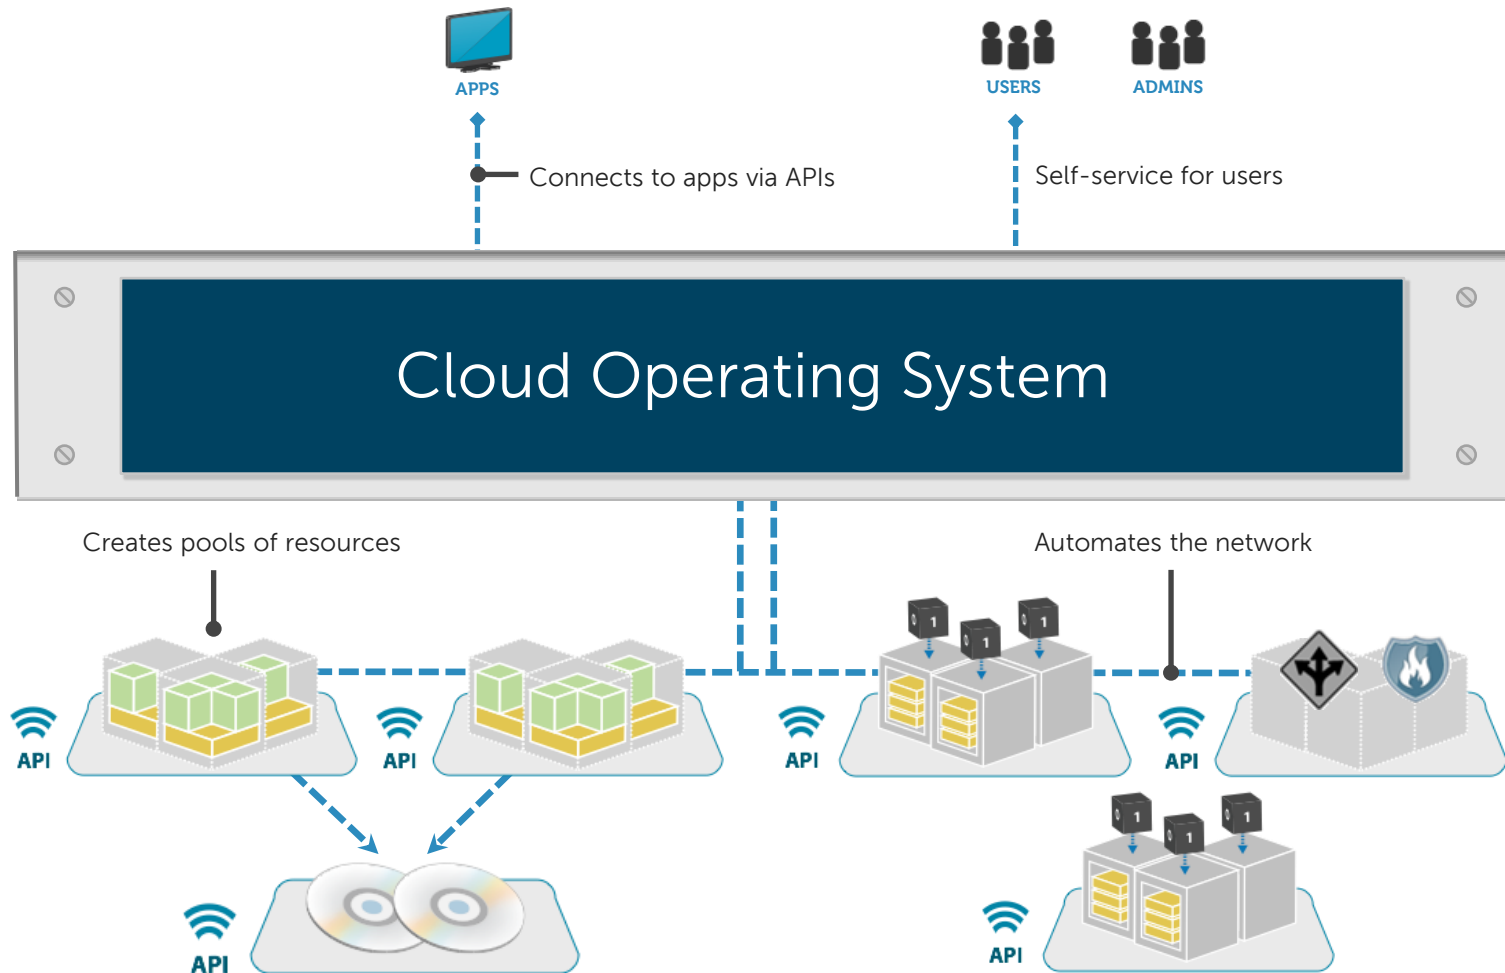
IaaS

Infrastructure as a Service

Target: Traditional IT – Server, Storage, Network, OS, Virtualization

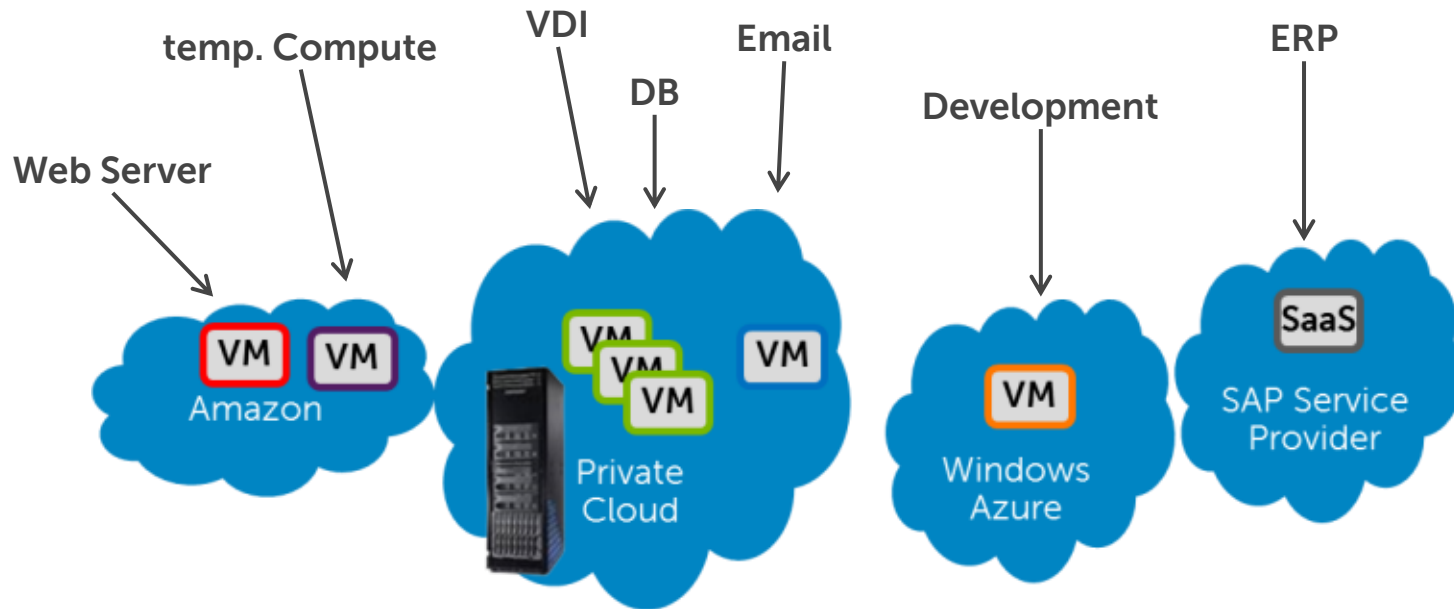
The Cloud Operating System

Manage the platform - not the VM...



Cloud Federation:

Making the Best of On- and Off-Premise


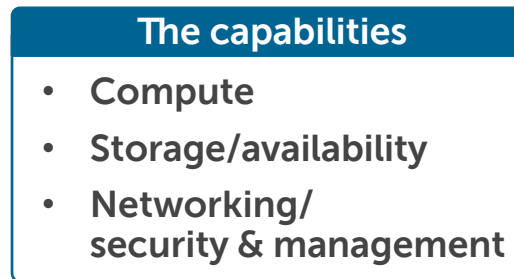
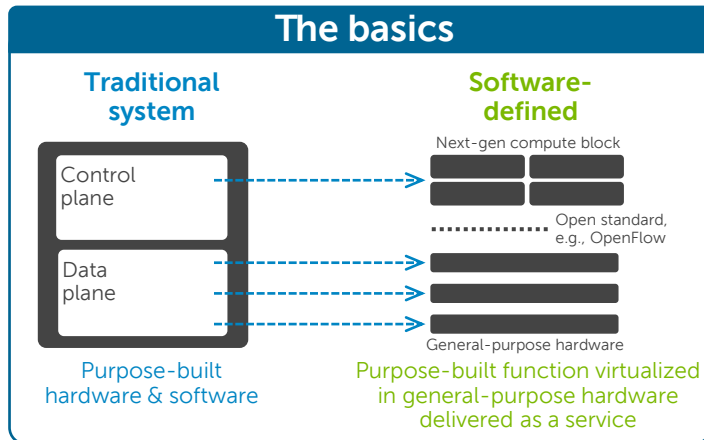
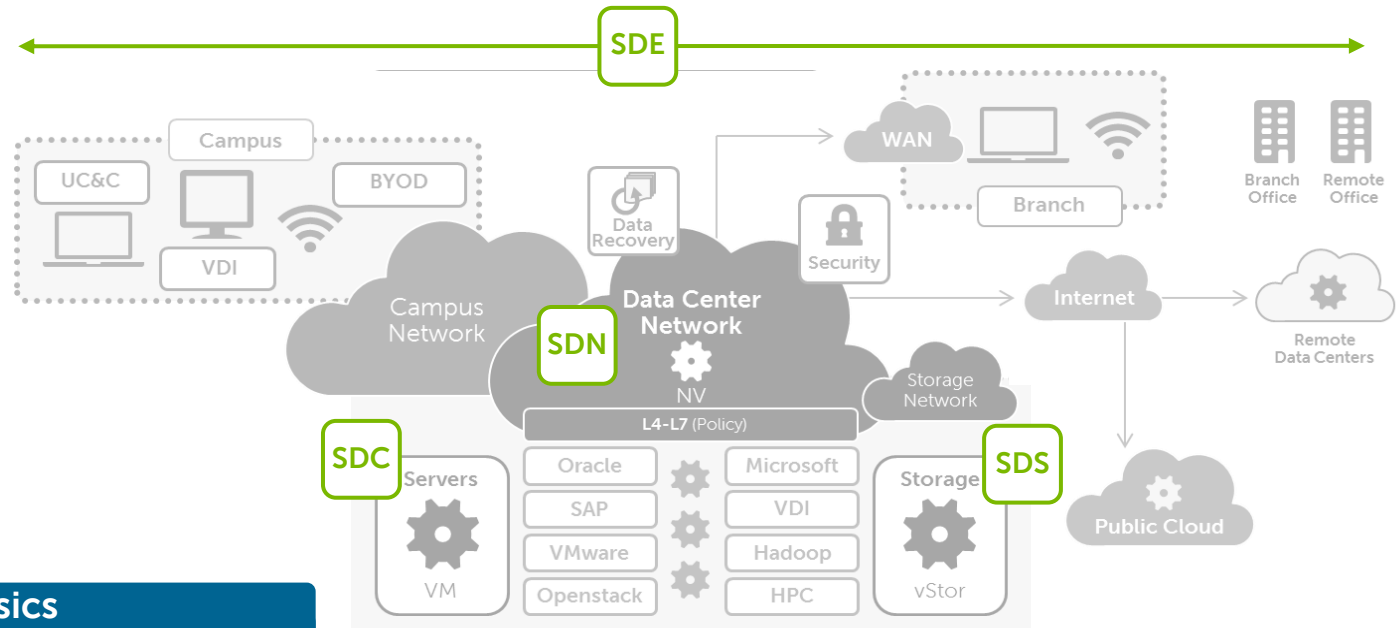


SDDC / SDE

Software-Defined DC


Software-Defined
Enterprise

Defining “software-defined”



Software-defined Compute?

Virtual Servers!



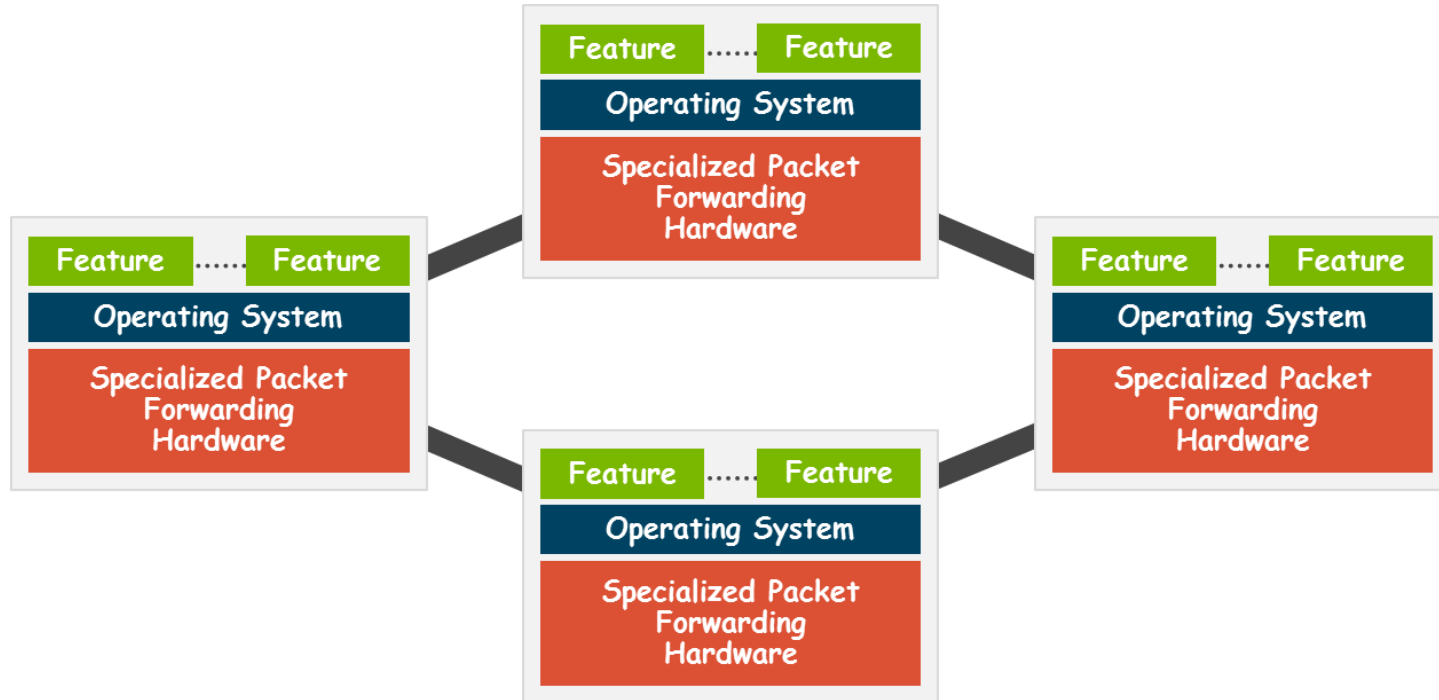
- ✓ Hardware abstraction for each server
- ✓ Better resource utilization for each server

SDN

Software-Defined Networking

So what is SDN really?

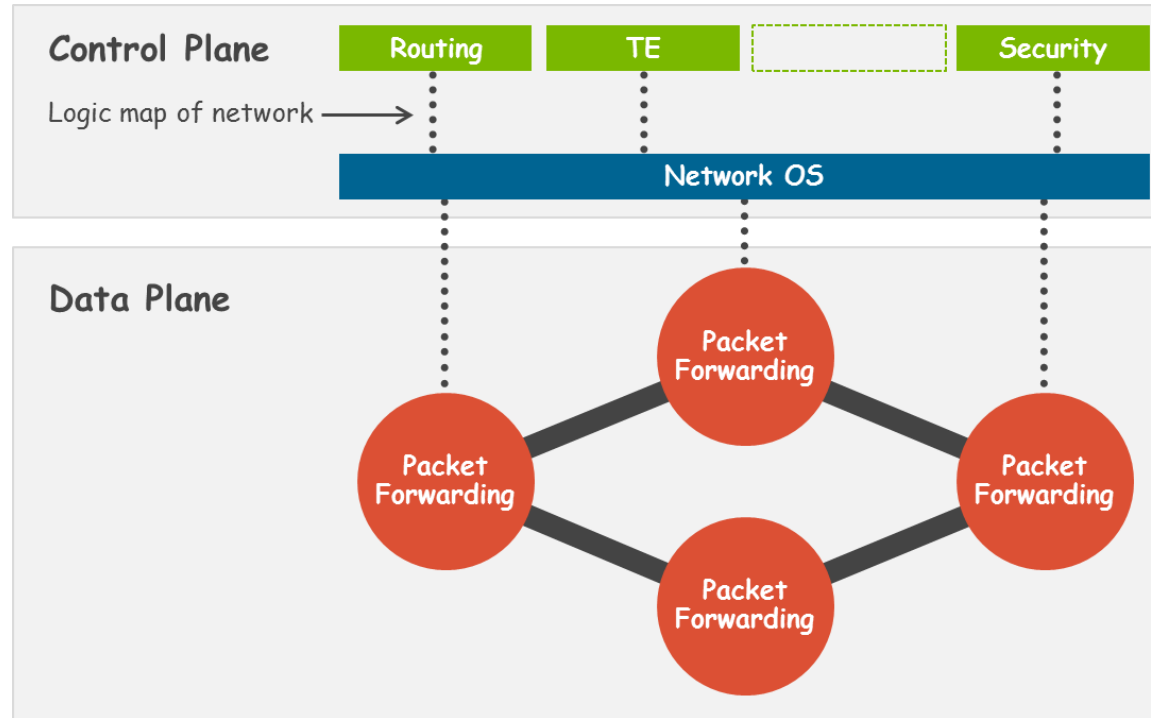
SDN according to Open Networking Foundation (ONF)



- SDN is a **new approach** to networking in which **network control is decoupled** from the **data forwarding** function and is **directly programmable**.
- The result is an extremely **dynamic, manageable, cost-effective**, and **adaptable** architecture that gives administrators **unprecedented programmability, automation, and control**.

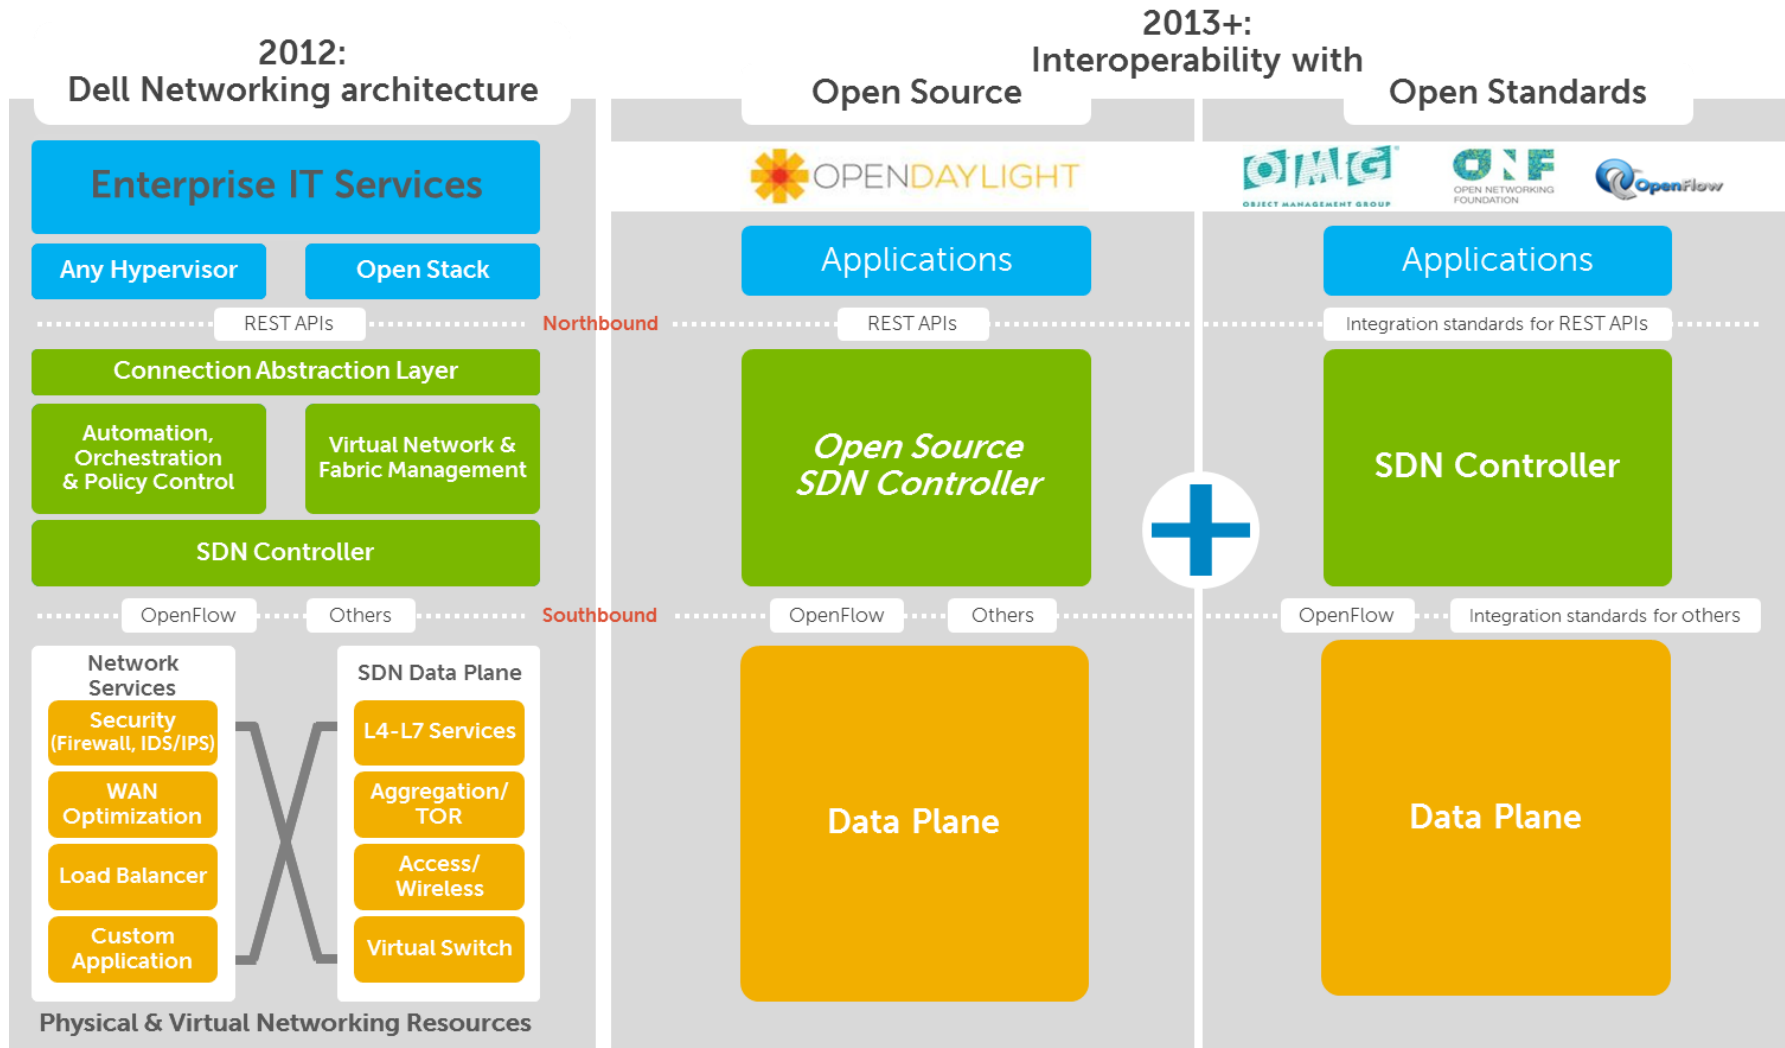
So what is SDN really?

SDN according to Open Networking Foundation (ONF)



- SDN is a **new approach** to networking in which **network control is decoupled** from the **data forwarding** function and is **directly programmable**.
- The result is an extremely **dynamic, manageable, cost-effective**, and **adaptable** architecture that gives administrators **unprecedented programmability, automation, and control**.

SDN – driven by Open Source & Open Standards



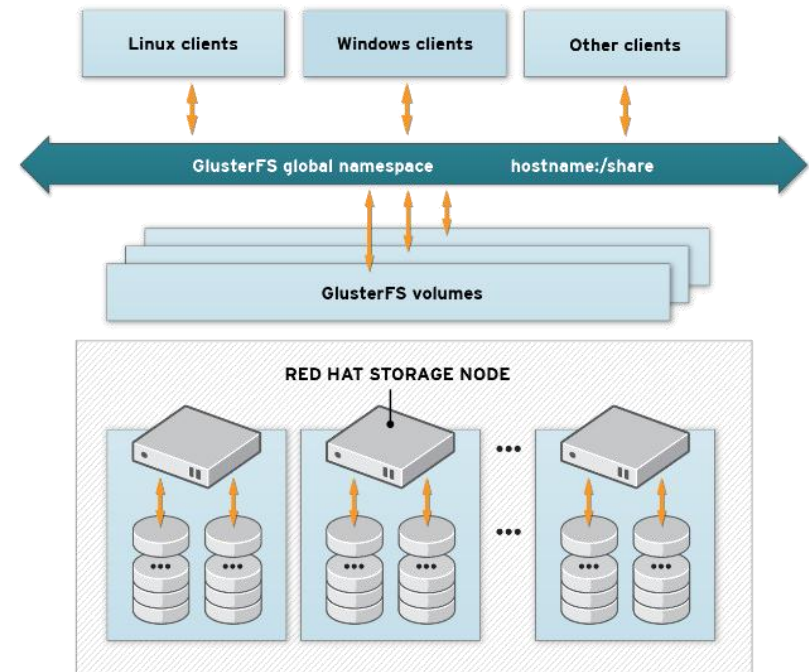
SDS

Software-Defined Storage

Software-Defined Storage (SDS)

New Horizons

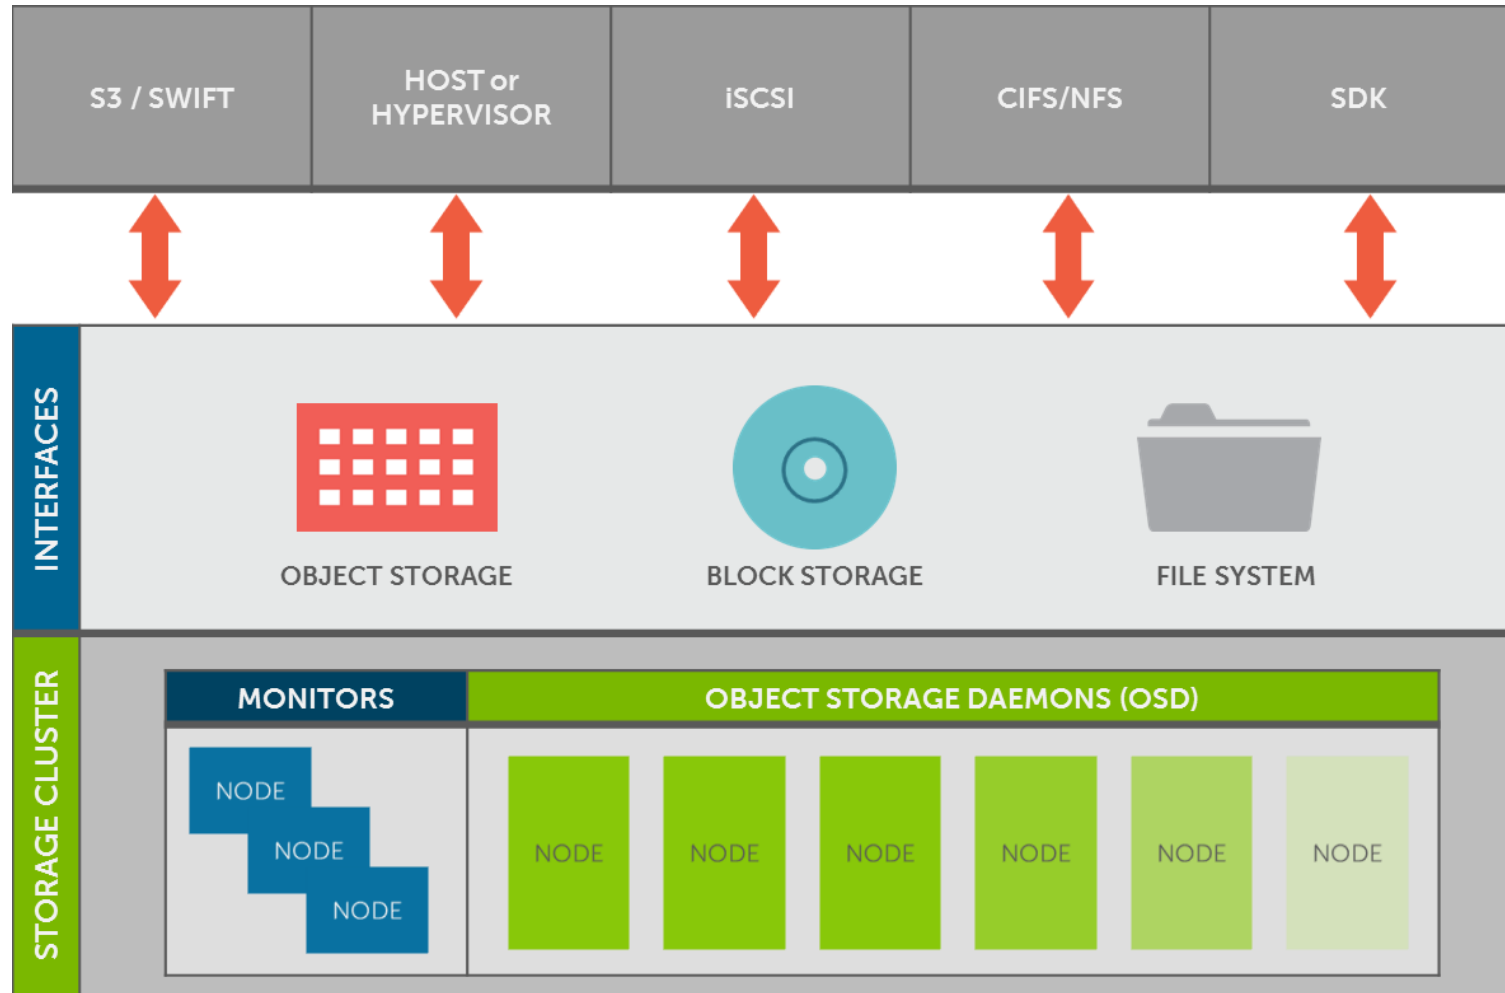
- From modern **local** file systems...
 - ZFS, btrfs, XFS,...
- ...via **distributed** file systems...
 - Lustre, GlusterFS, HDFS,...
- ...to **object** storage systems
 - OpenStack Swift, Scality, Ceph,...
- Infrastructure
 - **Commodity** servers & storage modules
 - **Commodity** network (1GbE, 10GbE)



Picture courtesy of Red Hat

Ceph – a Distributed Storage Platform

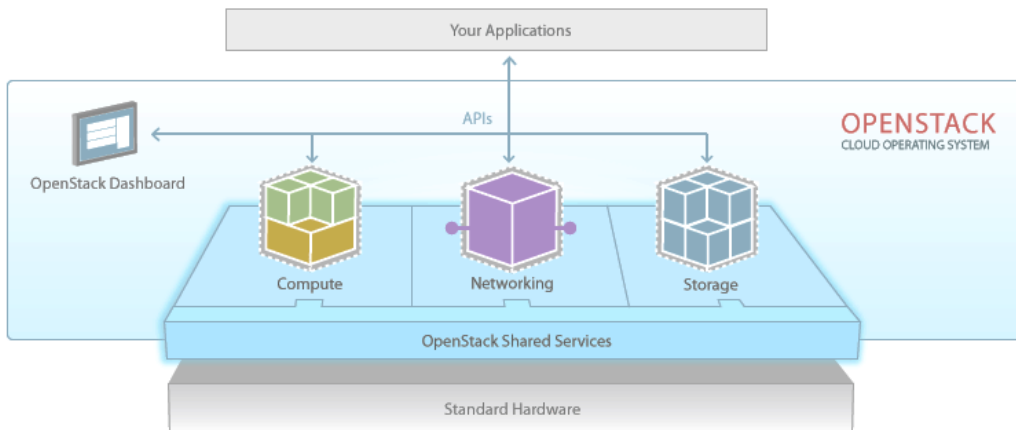
Development/support by **inktank** (now Red Hat)



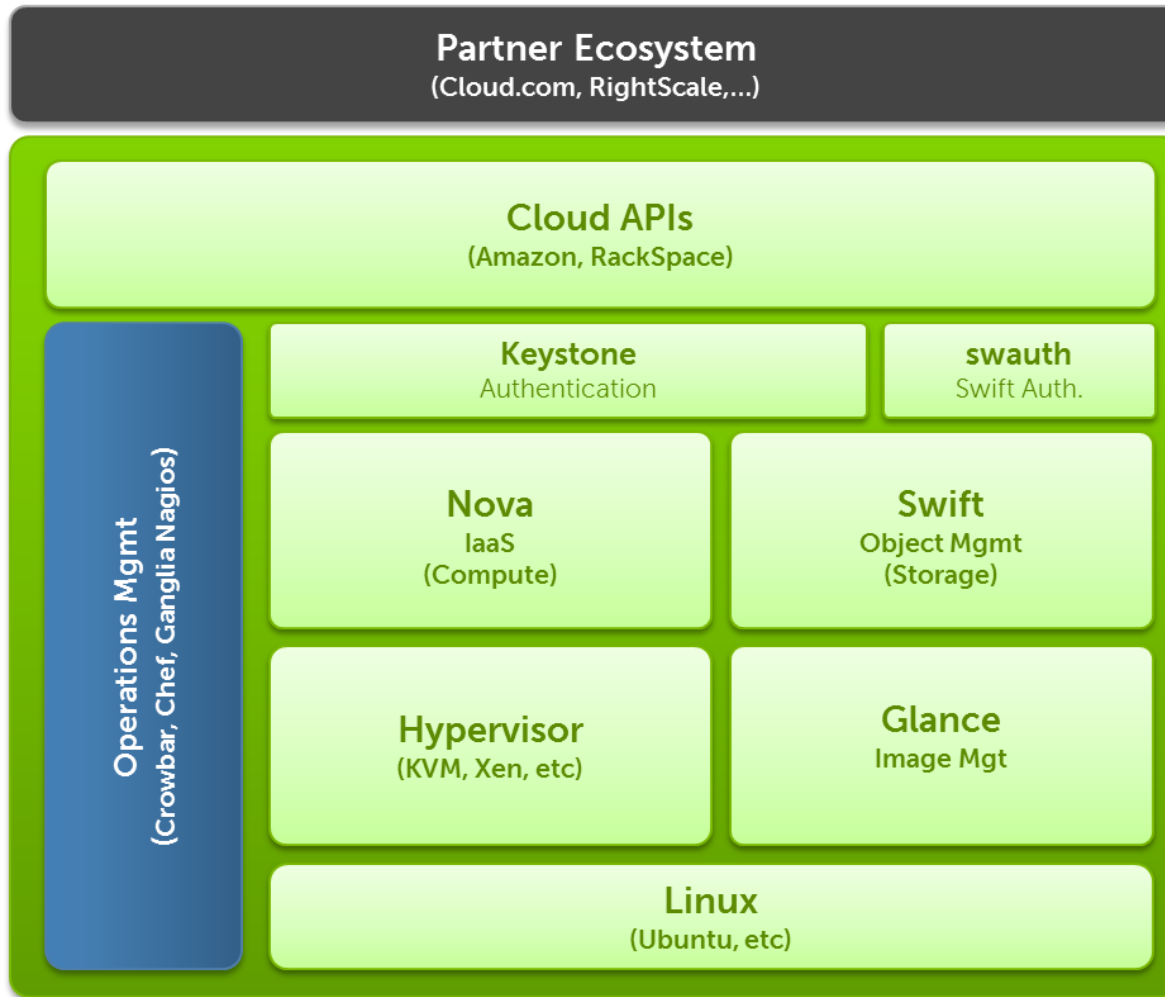
OpenStack

What is in OpenStack?

- A collection of projects developed and maintained collaboratively by a large, active community; e.g.:
 - **Compute** (Nova)
 - **Storage** (Swift & Cinder)
 - **Glance** (Imaging)
 - **Horizon** (Dashboard)
 - **KeyStone** (Authentication)
 - **Quantum/Neutron** (Network service)
 - **Heat** (Orchestration)
 - **Ceilometer** (Metering)
- Basic requirement: "it must be **simple to implement** and **massively scalable**."



OpenStack Technology Overview



Dell OpenStack-Powered Cloud Solutions

Proven solutions

Elastic , vastly scalable and designed to handle massive data loads

- Tested, validated, and innovative designs in infrastructure, software, and services
- Develop, deploy, and deliver your cloud environment
- Quickly offer new cloud services, lower software licensing costs, and help mitigate the risks of cloud computing



"Dell's commitment to OpenStack and their team with deep expertise in Data Center Solutions is the **foundation for a great partnership.**"

Ben Cherian, General Manager of Emerging Technologies, DreamHost

Proven components

OpenStack Cloud operating system

Dell-developed Crowbar Software*

Dell PowerEdge C-Series and R-Series servers

Force10 or PowerConnect Switches

Reference Architecture

Deployment Guide

Dell Service and Support



Dell | Red Hat Enterprise Cloud

Powered by OpenStack



Make
OpenStack
the defacto
OPEN
Cloud Solution
in the
Enterprise

- Address **the gaps for OpenStack** in the **Enterprise** via a jointly engineered solution
- **Deliver enterprise-grade, certified**, highly-scalable, secure, and **fully supported** cloud infrastructure solutions
- **Provide a fulfillment, deployment, and support experience** like commercial software
- Support **specific enterprise use cases** for cloud applications, software defined storage, development and test



Dell | Red Hat Enterprise Cloud

Powered by OpenStack



All benefits of community **OpenStack ...plus:**

- Enterprise **hardened** and **integrated** with Red Hat Enterprise Linux
- Provides **enterprise software lifecycle**
- World-class **global support**
- Worlds **largest OpenStack partner ecosystem**
- OpenStack **training** and **certification**
- Integrated with trusted Red Hat stack
 - Red Hat **CloudForms** (Hybrid cloud management)
 - Red Hat **OpenShift** (PaaS)
 - Red Hat Enterprise Virtualization
 - › **Docker** (containers vs. VMs)
 - Red Hat Storage: **Gluster & Ceph**
- Additional integration
 - ASM, Dell Cloud Manager, EQL, CML, SonicWall, Quest, Secureworks, OpenDaylight,...



Big Data & HPC in the Cloud

Scale-Out Use Cases

HPC/BigData in the Cloud – Why?


- Share Resources
 - protect investment
 - increase utilisation
- Multi-Tenancy
 - offer services for multiple customers/departments
- Agility / Flexibility
 - dynamically setup/expand/shrink a cluster – quickly & easily
- Multi-Purpose/Workload
 - use one platform for multiple use cases / workloads
 - › HPC (3D rendering, CFD,...)
 - › Big Data (Hadoop, noSQL,...)
 - › IaaS, PaaS,...
- Automation is Key!

Hadoop vs. HPC Clusters

Two Forms of Distributed Computing


HPC Cluster

- Parallel File System
 - High throughput
 - All nodes can access all data
 - **Compute-centric workloads**




Hadoop Cluster

- Distributed File System
 - Global namespace (ingest!)
 - Nodes just work on local data
 - **Data/IO-centric workloads**



Big Data in the Cloud – e.g. Hadoop

- Pros
 - no hardware dedicated to Hadoop
 - › very flexible usage of resources
 - › potential use of hybrid clouds
 - **good if** the analytic workload is **just temporary**
 - › typical data node: 2x 6-core CPU, 32-64GB RAM
- Cons
 - Hadoop is designed for local disks (JBOD)
 - › Shared local disks
 - potential performance loss
 - **large data can't be moved easily** (e.g. vMotion)
 - › SAN / NAS / Object Store / Parallel FS
 - **network IO will increase** a lot (no local processing anymore)
 - Hadoop is rack-aware
 - › availability concept (replication, node failure,...) needs to be reflected in cloud



HPC in the Cloud

- Network / Latency
 - std. clouds usually have 1/10 GbE
 - low-latency might be necessary
 - › Infiniband, RDMA over Converged Ethernet (RoCE)
- Compute Efficiency
 - heavily depends on compute model (i.e. workload!)
 - › hardware-specific programming (drivers,...)
 - › parallel programming / userspace communication
 - applications that need to bypass the OS kernel and communicate directly with remote user processes
 - GPGPU Computing
 - › pass-through with Xen HVM hypervisor & NVIDIA GPUs today
 - VMware, KVM, etc.: not yet (or tech previews)
 - › check **white paper** from **Indiana University** and **Information Sciences Institute (ISI), University of Southern California**:
[Enabling High Performance Computing in Cloud Infrastructure using Virtualized GPUs](#)
 - test using OpenStack is planned



HPC in the Cloud – Storage Considerations

- Parallel file access / high IO bandwidth necessary
 - Lustre, Gluster, pNFS,...
 - › mounted locally in VM
 - fast storage network!
 - › again: IB, RoCE (10/40GbE)
 - potential architectural overhead when block and/or object storage is needed as well
- ...so: this HPC cloud will probably look very different from a standard cloud...



The Dell | Cambridge HPC Solution Centre

- Established in 2010
- Proof-of-Concepts
- Services
 - Implementation, Support,...
- Best Practices



UNIVERSITY OF
CAMBRIDGE

HPC at University of Cambridge

- One of the largest HPC sites in EMEA
- HPC cloud for both Internal and external users
- First NVIDIA CUDA centre of excellence in the UK (2009)
- Strong growth in both academic and commercial programmes
- OpenStack to come this year!





Thank You!

Christian Spindeldreher

Enterprise Technologist

christian_spindeldreher@dell.com



Backup Slides

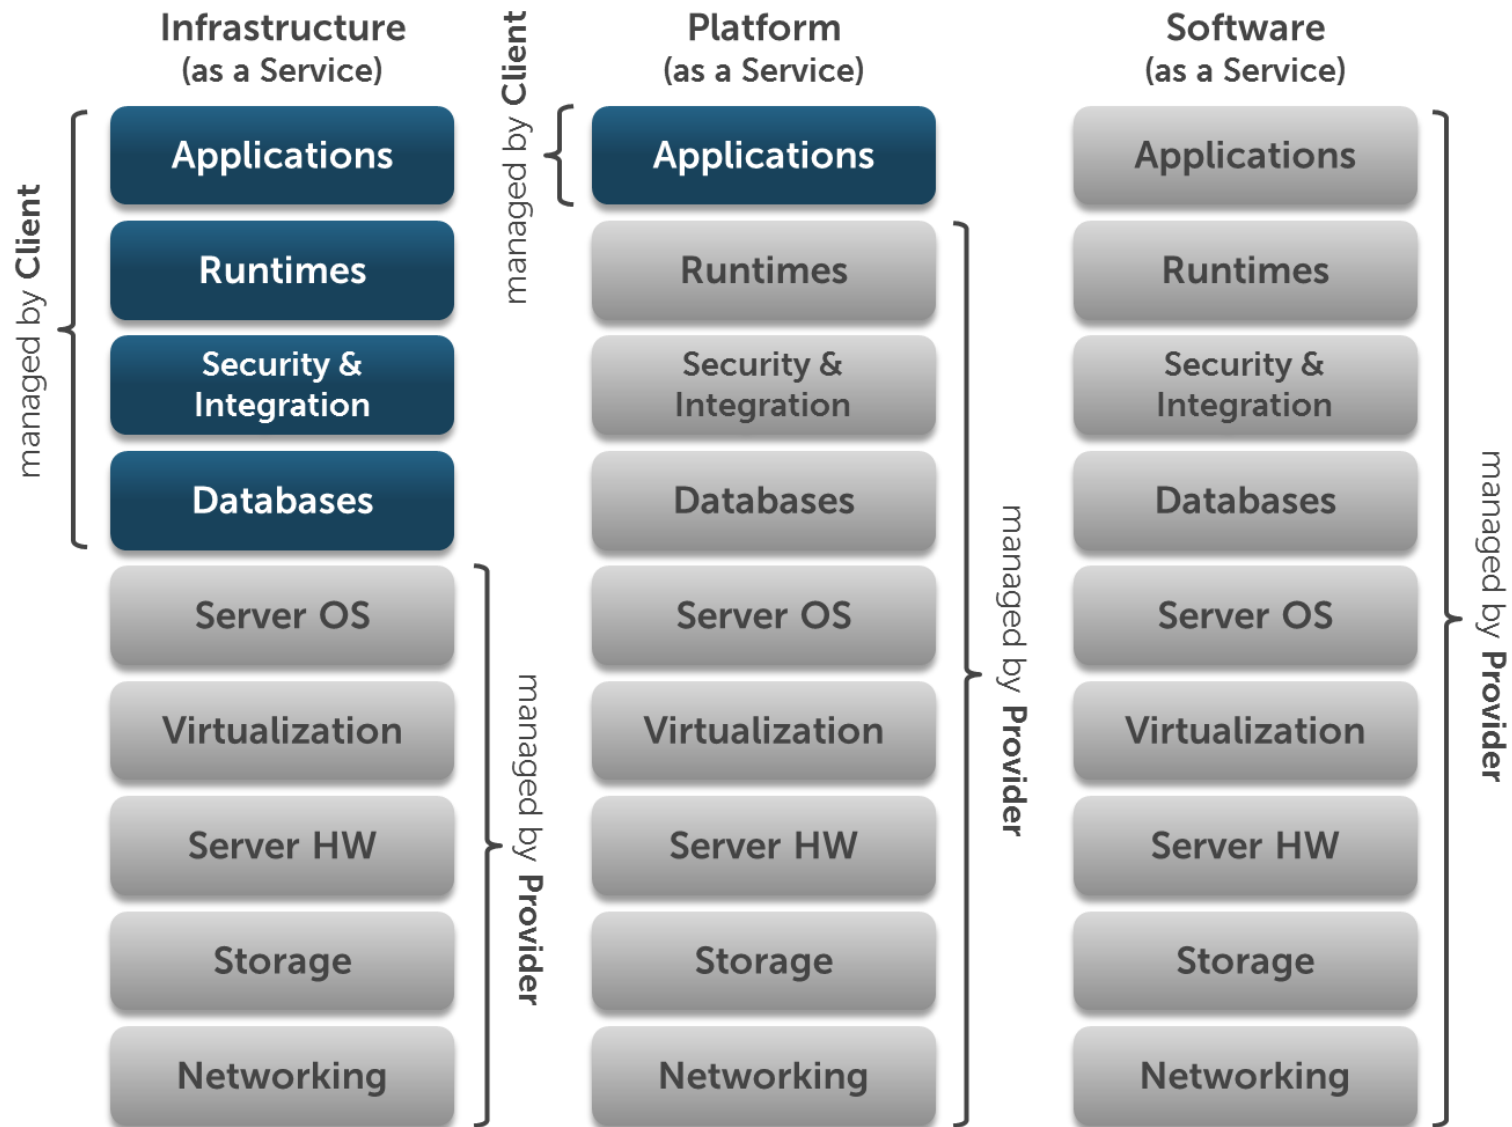


Cloud Computing Defined

Cloud computing is a model for enabling ubiquitous, convenient, **on-demand** network access to a shared pool of **configurable computing resources** (e.g., networks, servers, storage, applications, and services) that can be **rapidly provisioned** and released with **minimal management effort or service provider interaction**

US National Institute for Standards and Technology

IaaS, PaaS & SaaS

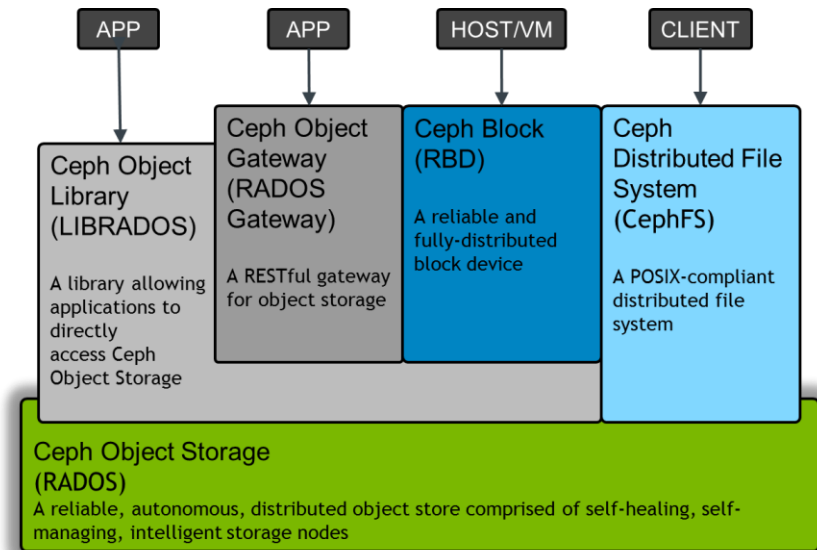


Ceph – a Distributed Storage Platform

Development/support by **inktank** (now **redhat**!)



- Object Storage
 - S3 & Swift APIs
- Block Devices
 - vLUNs
- File System
 - Global namespace
 - Metadata server cluster
 - Linux client
- Redundant shared-nothing cluster architecture
 - no SPOF
- Data distribution with **CRUSH**
 - **C**ontrolled **R**eplication **U**nder **S**calable **H**ashing
 - object location can be calculated
 - storage nodes can be grouped in failure domains
 - replication for data protection



- Roadmap
 - Erasure Coding v1.2 (May 2014)
 - Cache Tiering
 - CephFS v2.0 (Sep 2014)
 - Hyper-V support
 - VMware support later (2015)
 - DeDup



Dell Cloud Manager

Once known as Enstratus



Dell Cloud Manager

SaaS or On-Premises

Deploy & manage enterprise applications
across private, public & hybrid clouds
with Agility, Governance and Choice

Dasein Cloud Abstraction Layer



Off-Premises Public Clouds




On-Premises Private Clouds

Applications in the Cloud

Composite Applications

Multi-Cloud

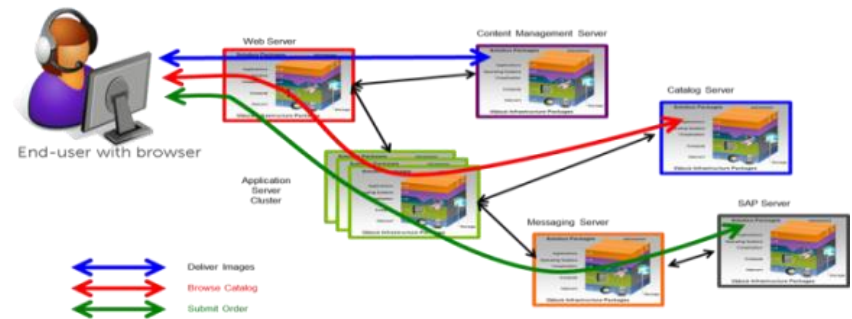
Our customers have workloads running in many locations.



+

Composite Applications

Applications are becoming complex combinations of interconnected software components.



Composite Applications

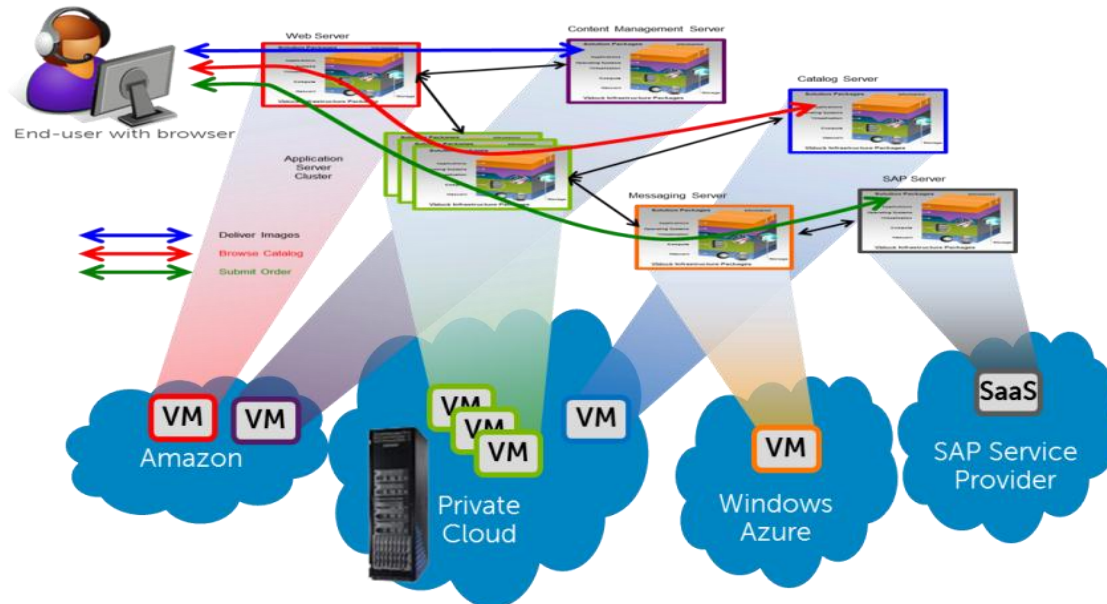
Multi-Cloud

+

Composite Applications

Our customers have workloads running in many locations.

Applications are becoming complex combinations of interconnected software components.



= a **Composite Application** running **across Multiple Clouds**

What are the right workloads for the Cloud?


- Identify applications that benefit from cloud characteristics
 - horizontal scalability / cloud bursting
 - very dynamic / on-demand
 - › batch, fluctuating traffic,...
 - › projects, launches, temporary needs,...
 - Disaster Recovery (DR) – depends on details(!)...
- What might be less suitable?
 - high performance requirements
 - › esp. storage
 - high availability requirements
 - applications that just scale vertically
 - applications that work with confidential data
 - applications with large data capacities
 - applications that are hardware-dependent/-optimized



**needs further
investigation**


Modern Environments Can Be Complex

User Experience? Latency? Performance Bottlenecks?



Dell Foglight APM

Application Performance Management



- User Experience Monitoring, TransactionDNA Technology, Analytics,...

Dell Foglight Performance Monitoring

End-to-End View



Application-Centric Enterprise Performance Monitoring

Maximize team collaboration based on a common platform and single version of the truth

Application Performance Monitoring

Most integrated customer and transaction-centric APM

Database Performance Monitoring

Deepest DB monitor and workload analysis for multiple platforms

Infrastructure Performance Monitoring

Virtualization

Enterprise virtualization management with analytics, advice and intelligent automation

Storage

Deepest storage monitor with analytics for virtual and physical environments