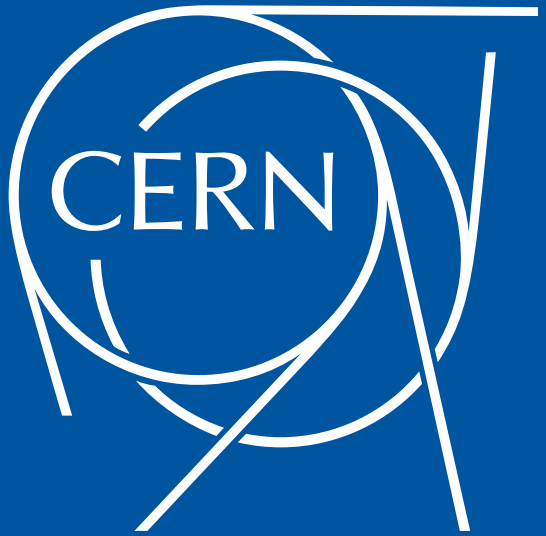
CERN

Introducing SDN in CERN Cloud

José Castro León
CERN Cloud Infrastructure

Outlines

- Introduction
- Network DC Architecture
  - Status & Plans
- OpenStack Network Status and Evolution
- Tungsten SDN region

European Organization for Nuclear Research

- World largest particle physics laboratory
- Founded in 1954
- 22 member states
- Fundamental research in physics

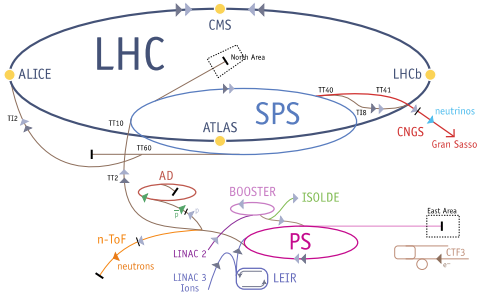

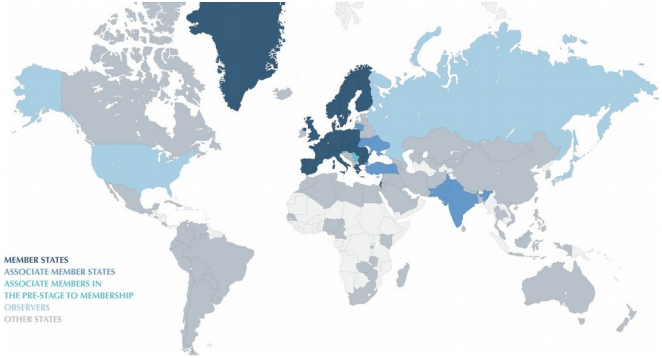

CERN Cloud Service

- Infrastructure as a Service
- Production since July 2013
- CentOS 7 based
- Geneva and Wigner Computer centres
- Highly scalable architecture > 70 nova cells
- Currently running Rocky release

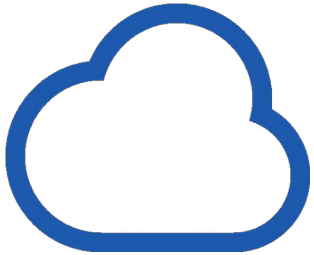

Cloud resources

Used	Available	Used	Available	Used	Available
307.4 K cores	300.0 K cores	858.7 TiB RAM	937.5 TiB RAM	10.3 PiB disk	15.6 PiB disk

Openstack services stats

Users	Projects	VMs	Magnum clusters	Hypervisors	Images	Baremetal nodes
3492	4526	36203	420	9212	3179	1294

Volumes	Volume size	Fileshares	Fileshares size
6093	1.90 PiB	248	215 TiB

Resource overview by time

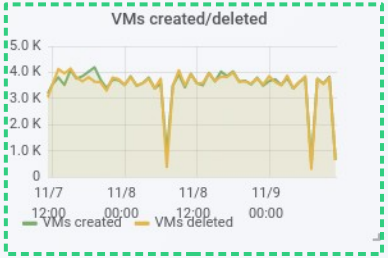

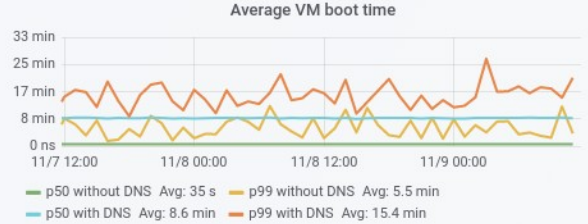

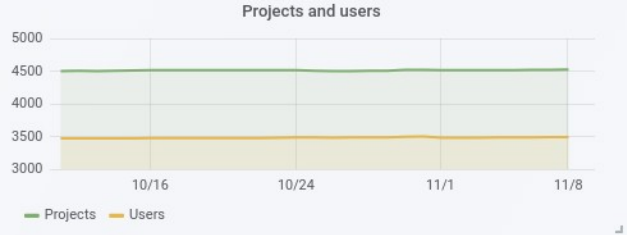

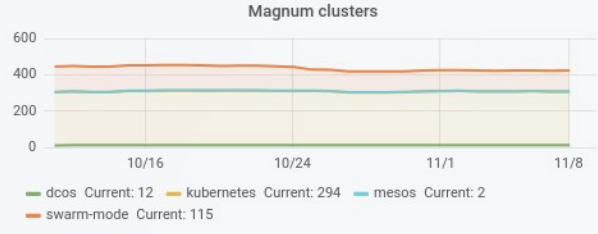

CERN Cloud Infrastructure

IaaS+

Orchestration

heat

Container Orchestration

magnum

Automation

mistral

Web UI

horizon

IaaS

Network

neutron

Compute

nova

ironic

Storage

cinder

manila

glance

Identity

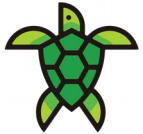

keystone

Key manager

barbican

Datacentre Networking

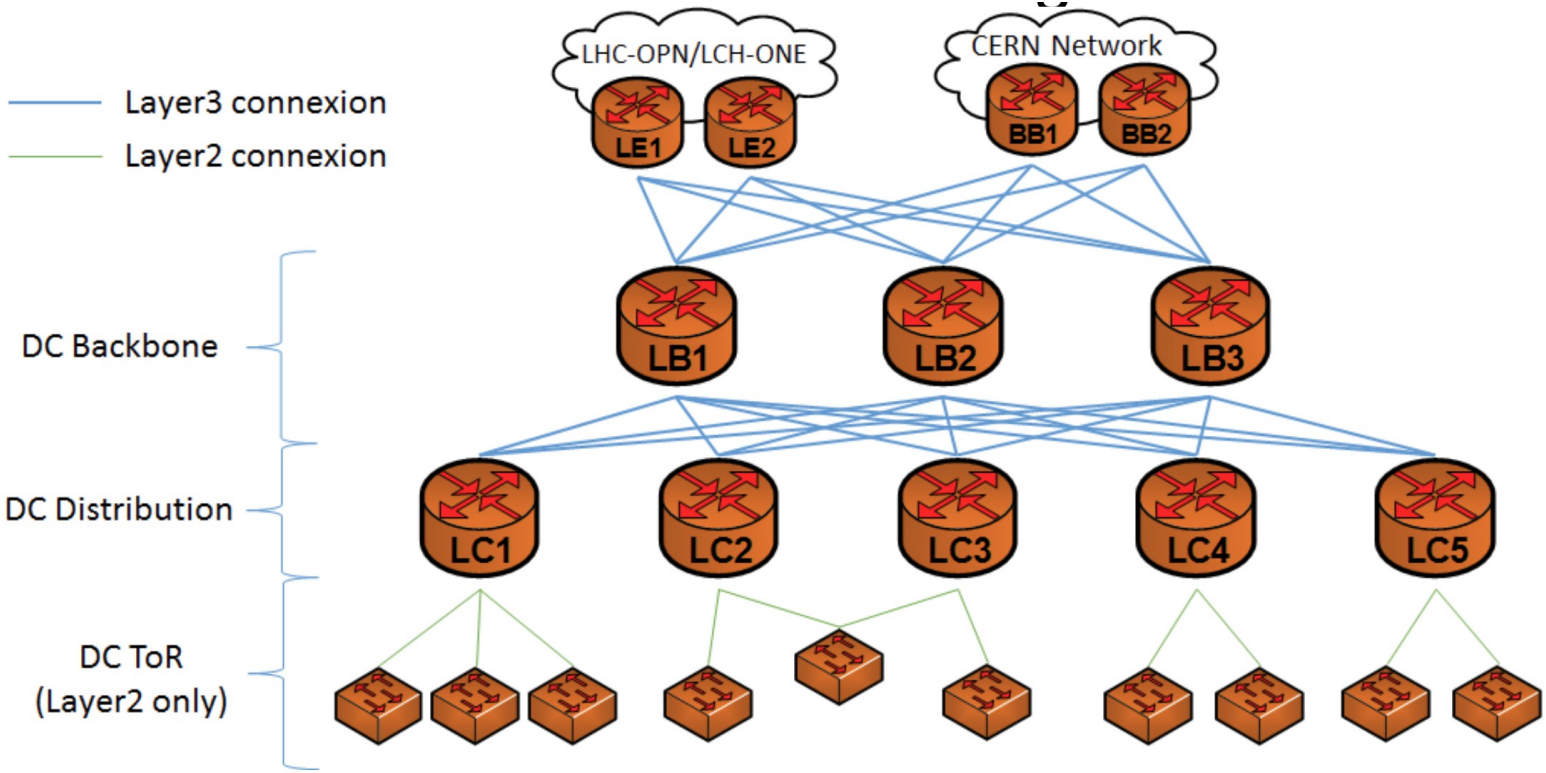

Limitations of current setup

- Limited IP Mobility
  - Segmented broadcast domains
  - Live migration limited to single cluster
  - Ad-hoc tunnels for hardware retirement campaigns
- Hardware Repurposing
  - Multiple network domains (General, Services, ...)
  - Services dedicated to a single domain
- No Floating IPs
- No Tenant/Private Networks

OpenStack networking

- CERN OpenStack networking recently migrated to Neutron
- Linuxbridge, Flat / Provider networks
- Better integration using ML2, mechanism driver and extensions
  - Quickly became possible to have it out of tree
  - Our extensions have a similar role to Neutron Segments
- Gradual enroll, cell by cell
- Vanilla upstream packages for Neutron, much smaller patch on Nova

Future Datacentre network design

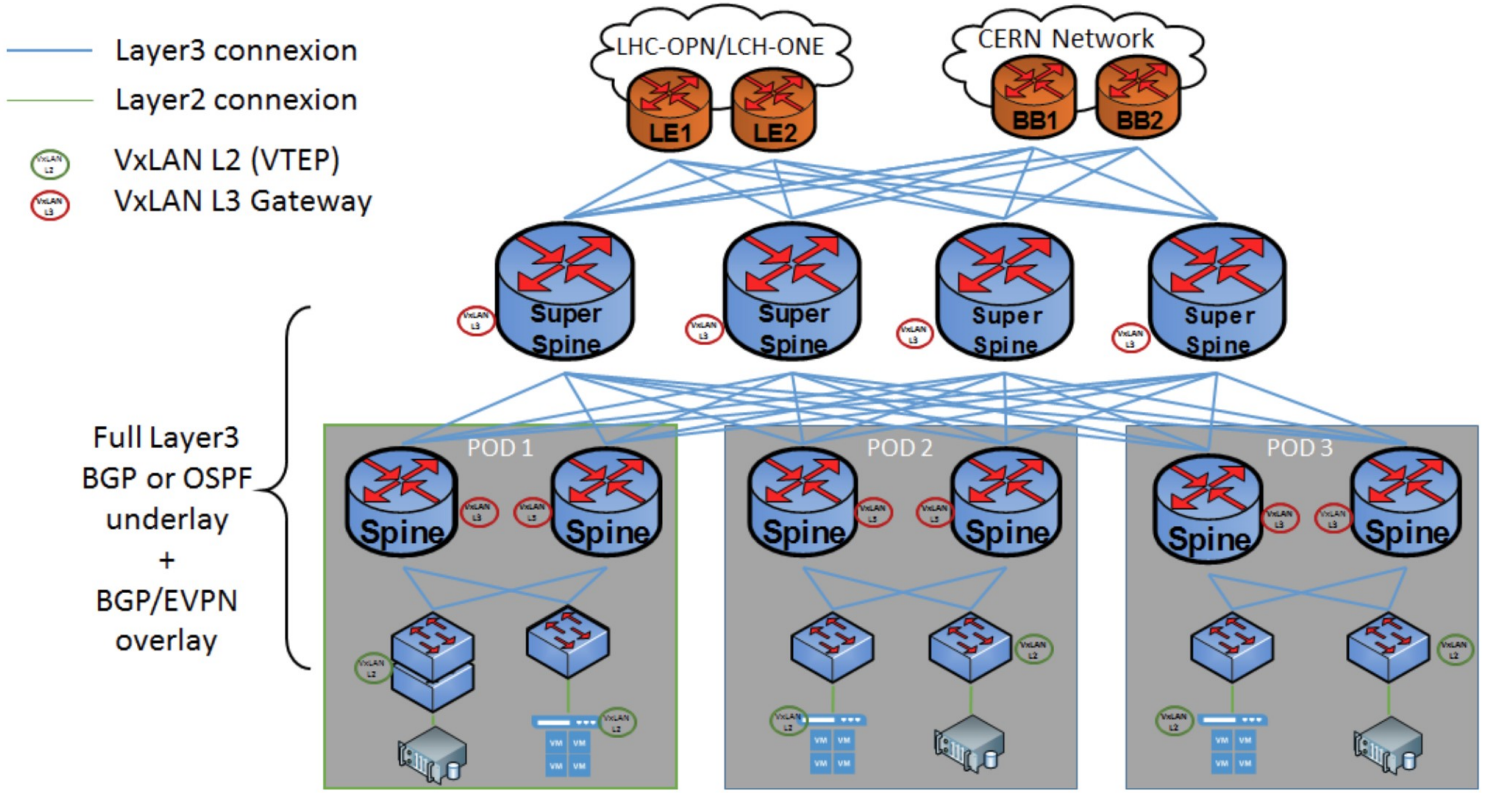

SDN Datacenter schema

- Orchestrator
  - Compute (VMs, containers, BM)
  - Storage
  - Network
- SDN Controller
- DC Network virtual and physical
  - switches
  - routers / gateways.

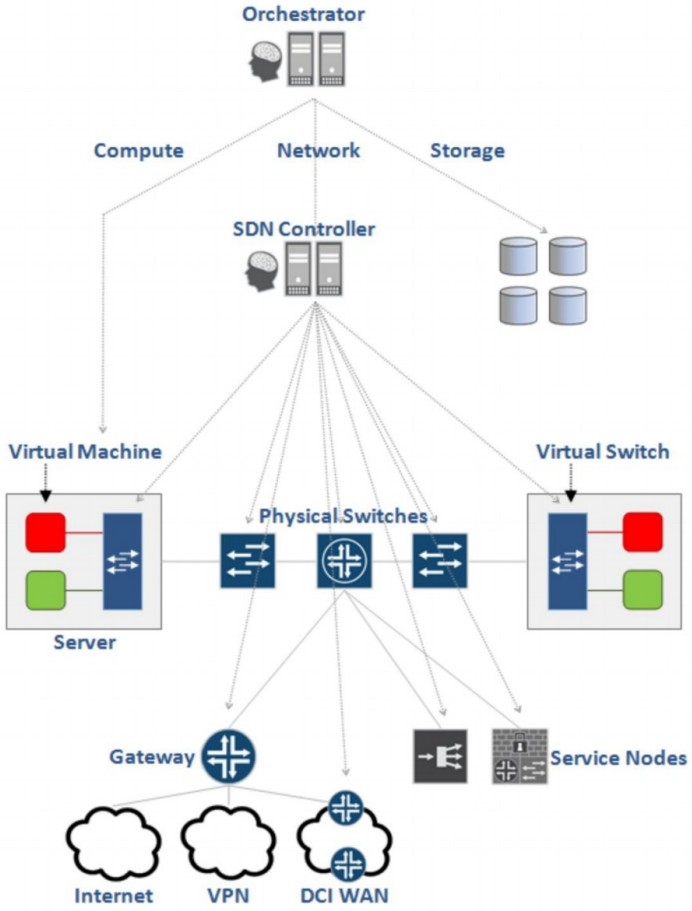

SDN Region @CERN

- Tungsten Fabric
  - Deployed using Contrail Helm
  - Full Cluster
  - Docker images from Tungsten
  - Hypervisors configured with Puppet
    - Vrouter module and agent

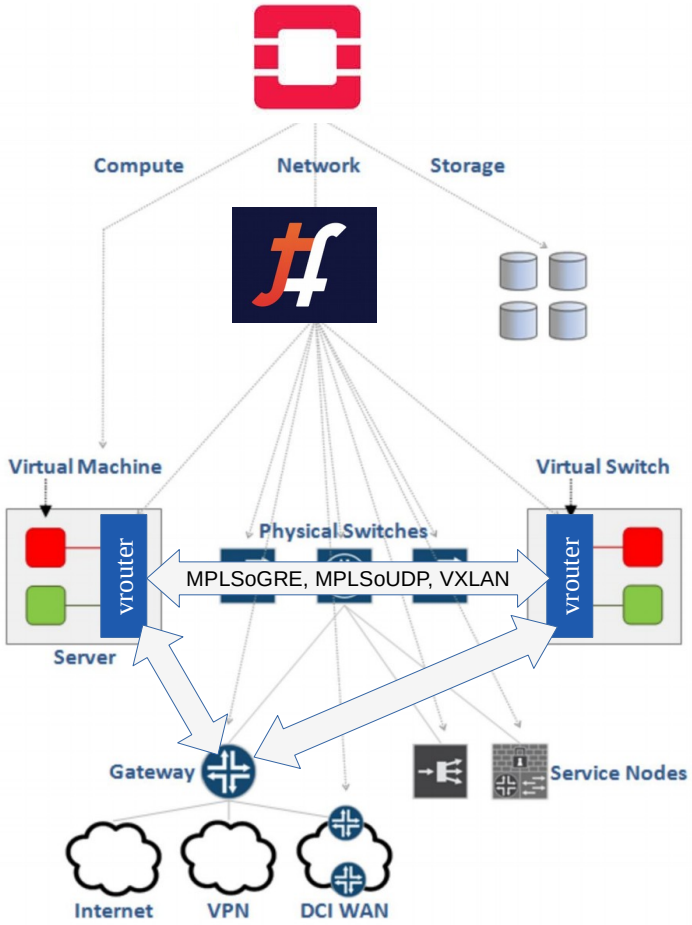

Control plane deployment

- Installed using contrail-helm-deployer
- Using Docker Hub latest images

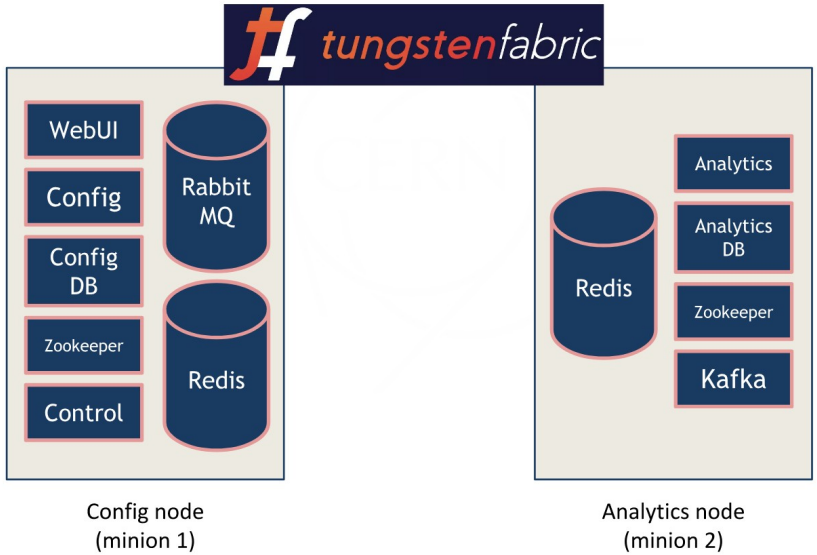

https://github.com/juniper/contrail-helm-deployer

VRouter / Hypervisor setup

- Hypervisor managed with puppet as other regions
- Docker module using multiple containers
  - kernel-vrouter-init
  - contrail-vrouter-agent
  - contrail-vrouter-nodemgr
  - Nova-compute-init
- Region fully deployed using OpenStack Helm and LOCI images

Work in progress

- Software Defined Network region
  - Works perfectly on Virtual machines
  - Evaluating integration with Physical nodes

- Contribute to upstream community
  - And looking forward to contribute even more :D

Thank you

gitlab.cern.ch/cloud-infrastructure

openstack-in-production.blogspot.ch

jose.castro.leon@cern.ch

@josecastroleon

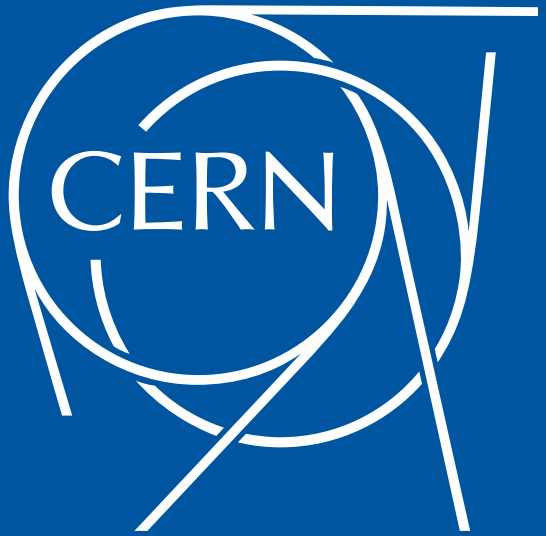
CERN

BACKUP SLIDES

Future Datacentre network

- Based on Juniper QFX devices (re-use some of existing Brocade ToR switches)
- Full fabric routed up to the ToR
- Router redundancy for all ToR switches
- 2 levels of spines (Spines and SuperSpines) ⇒ QFX10000 chassis
- ToR switches ⇒ QFX5xx0 switches
- Support of VxLAN:
  - Use of BGP/EVPN on the overlay
  - VxLAN Layer3 gateway position still to be defined
  - Integration with OpenContrail/Tungsten on study
  - VTEP on Hypervisor or on ToR switches