

Bits and Bytes im Flow – Netzwerk im Wandel der Zeit

GUUG-Frühjahrsfachgespräch 2016

26. Februar 2016



Michel Rode
Linux/Unix Consultant & Trainer
B1 Systems GmbH
rode@b1-systems.de

Vorstellung B1 Systems

- gegründet 2004
- primär Linux/Open Source-Themen
- national & international tätig
- über 70 Mitarbeiter
- unabhängig von Soft- und Hardware-Herstellern
- Leistungsangebot:
 - Beratung & Consulting
 - Support
 - Entwicklung
 - Training
 - Betrieb
 - Lösungen
- dezentrale Strukturen

Schwerpunkte


- Virtualisierung (XEN, KVM & RHEV)
- Systemmanagement (Spacewalk, Red Hat Satellite, SUSE Manager)
- Konfigurationsmanagement (Puppet & Chef)
- Monitoring (Nagios & Icinga)
- IaaS Cloud (OpenStack & SUSE Cloud & RDO)
- Hochverfügbarkeit (Pacemaker)
- Shared Storage (GPFS, OCFS2, DRBD & CEPH)
- Dateiaustausch (ownCloud)
- Paketierung (Open Build Service)
- Administratoren oder Entwickler zur Unterstützung des Teams vor Ort

Software-defined networking – SDN

Aufteilung


- Teil 1 - Allgemeiner Überblick
- Teil 2 - Umsetzung innerhalb von OpenStack
- Teil 3 - "Ausfallsichere" Kommunikation mit der Außenwelt

Zuerst, warum SDN?



Quelle: <https://www.getdigital-blog.de/>

Zuerst, warum SDN?



Quelle: <https://www.getdigital-blog.de/>

- zu komplex
- zu unflexibel
- ineffizient
- keine Automation

Traditionell vs. modern

- Computing
 - Mainframe

Traditionell vs. modern

- Computing
 - Mainframe
 - Server

Traditionell vs. modern

- Computing
 - Mainframe
 - Server
 - KVM

Traditionell vs. modern

- Provisioning
 - Diskette

Traditionell vs. modern

- Provisioning
 - Diskette
 - CDROM

Traditionell vs. modern

- Provisioning
 - Diskette
 - CDROM
 - qcow Image

Traditionell vs. modern

- Networking
 - Kabel

Traditionell vs. modern

- Networking
 - Kabel
 - ???

Definition SDN

- Centralization
- Separation
- Automation & Orchestration

Definition SDN

- Centralization
- **Separation**
- Automation & Orchestration

Definition SDN

- Centralization
- Separation
- Automation & Orchestration

Ist das neu?

- Neu? Nein mit Perwoll gewaschen!

Ist das neu?


- Wireless Controller
 - Cisco, Ubiquiti, ...

Ist das neu?

- Wireless Controller
 - Cisco, Ubiquiti, ...
- Telco Hardware (zB MPLS)


Ist das neu?

- Marionetten der IT




Warum dann SDN?

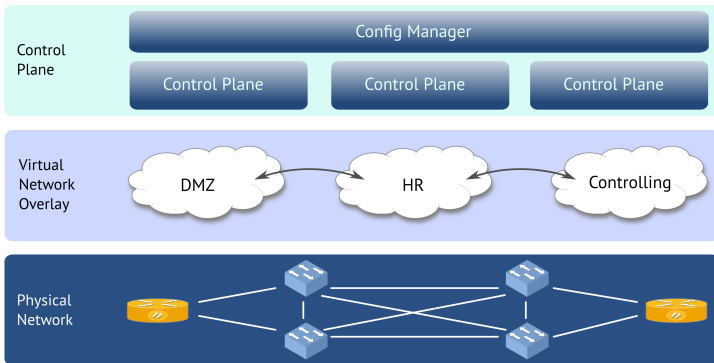
Logischer Aufbau



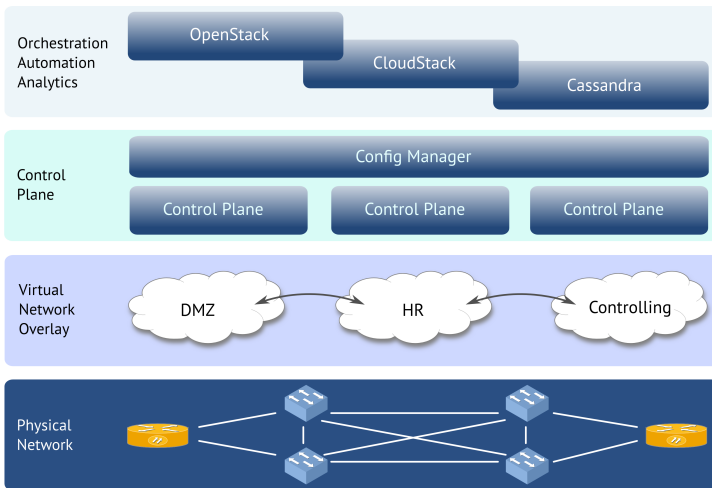
Logischer Aufbau



Logischer Aufbau




Logischer Aufbau



Control Plane & Data Plane

- Trennung von
 - Logik und
 - Datenpfad
- Zentrale Verwaltung
- Dynamische Steuerung

Control Plane & Data Plane





Implementierung

Implementierung

- Midonet
- Contrail/OpenContrail
- PLUMgrid
- OpenDayLight
- Floodlight
- Maestro
- ...

Implementierung

- Midonet
- Contrail/OpenContrail
- PLUMgrid
- OpenDayLight
- Floodlight
- Maestro
- ...

Implementierung

- Midonet
- Contrail/OpenContrail
- PLUMgrid
- OpenDayLight
- Floodlight
- Maestro
- ...

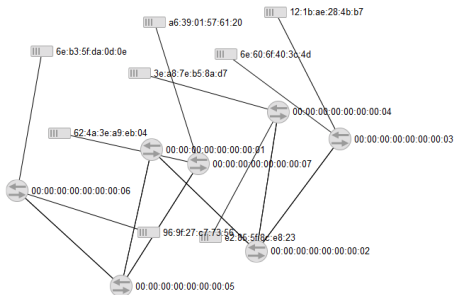
Floodlight



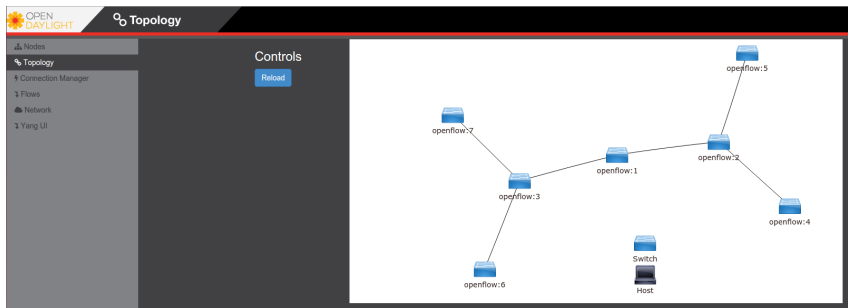
Dashboard [Topology](#) Switches Hosts

10.0.0.1

Network Topology



OpenDayLight



OpenFlow & OpenStack

Was ist (Open)Flow?

- vergleichbar mit einem L2 Switch
- erweitert um Header
- Schnittstelle - "API"
- Kombination aus Data- & Controlplane
- Flow als kleinste Einheit
- Routing-Protokoll für Netzwerkpakete

Wo kommt das zum Einsatz?

- Neutron ML2 w/ openvswitch
 - Controller (Neutron)
 - Compute (Nova)
- Neutron ML2 w/ *

OpenvSwitch?

- Simple: Software-Switch
- Detail:
 - ovs-vs,of,dpctl
 - vswitchd & ovssdb
 - Verteile Installation mögl.
 - Unterstützt OpenFlow!!111

```

controller# ovs-ofctl dump-flows br-tunnel
NXST_FLOW reply (xid=0x4):
[...] table=0, n_bytes=0, idle_age=1395, priority=0 actions=drop
[...] table=0, n_bytes=1310, idle_age=1390, priority=1,in_port=1 actions=resubmit(,1)
[...] table=1, n_bytes=0, idle_age=1392, priority=2,dl_vlan=1,dl_dst=fa:16:3e:c9:10:fd actions=drop
[...] table=1, n_bytes=1268, idle_age=1390, priority=0 actions=resubmit(,2)
[...] table=1, n_bytes=0, idle_age=1392, priority=1,dl_vlan=1,dl_src=fa:16:3e:c9:10:fd \
[...] actions=mod_dl_src:fa:16:3f:1e:ba:3e,resubmit(,2)
[...] table=1, n_bytes=0, idle_age=1392, priority=3,arp,dl_vlan=1,arp_tpa=192.168.123.1 actions=drop
[...] table=2, n_bytes=378, idle_age=1390, priority=1,arp,dl_dst=ff:ff:ff:ff:ff:ff \
[...] actions=resubmit(,21)
[...] table=2, n_bytes=1160, idle_age=1392, priority=0,dl_dst=01:00:00:00:00:00/01:00:00:00:00:00 \
[...] actions=resubmit(,22)
[...] table=2, n_bytes=0, idle_age=1395, priority=0,dl_dst=00:00:00:00:00:00/01:00:00:00:00:00 \
[...] actions=resubmit(,20)
[...] table=3, n_bytes=0, idle_age=1395, priority=0 actions=drop
[...] table=3, n_bytes=0, idle_age=1392, priority=1,tun_id=0x1 actions=mod_vlan_vid:1,resubmit(,9)
[...] table=4, n_bytes=0, idle_age=1395, priority=0 actions=drop
[...] table=6, n_bytes=0, idle_age=1395, priority=0 actions=drop
[...] table=9, n_bytes=0, idle_age=1395, priority=0 actions=resubmit(,10)
[...] table=9, n_bytes=0, idle_age=1394, priority=1,dl_src=fa:16:3f:11:d2:65 actions=output:1
[...] table=10, n_bytes=0, idle_age=1395, priority=1 actions=learn(table=20,hard_timeout=300,\
[...] priority=1,cookie=0x99074aa5abe39b41,NXM_OF_VLAN_TCI[0..11],\
[...] NXM_OF_ETH_DST[]=NXM_OF_ETH_SRC[], \
[...] load:0->NXM_OF_VLAN_TCI[],load:NXM_NX_TUN_ID[]->NXM_NX_TUN_ID[],\
[...] output:NXM_OF_IN_PORT[]),output:1
[...] table=20, n_bytes=0, idle_age=1395, priority=0 actions=resubmit(,22)
[...] table=21, n_bytes=378, idle_age=1390, priority=0 actions=resubmit(,22)
[...] table=22, n_bytes=1448, idle_age=1390, priority=0 actions=drop

```



```
compute# ovs-ofctl dump-flows br-public
NXST_FLOW reply (xid=0x4):
[...] table=0, n_packets=698, n_bytes=59520, idle_age=1, priority=1 actions=resubmit(,3)
[...] table=0, n_packets=45, n_bytes=3868, idle_age=864, priority=0 actions=NORMAL
[...] table=0, n_packets=0, n_bytes=0, idle_age=863, priority=2,in_port=2 actions=resubmit(,1)
[...] table=1, n_packets=0, n_bytes=0, idle_age=863, priority=0 actions=resubmit(,2)
[...] table=2, n_packets=0, n_bytes=0, idle_age=863, priority=2,in_port=2 actions=drop
[...] table=3, n_packets=698, n_bytes=59520, idle_age=1, priority=1 actions=NORMAL
[...] table=3, n_packets=0, n_bytes=0, idle_age=863, priority=2,dl_src=fa:16:3f:1e:ba:3e \
[...] actions=output:2
```

```
controller# ovs-ofctl dump-flows br-public
NXST_FLOW reply (xid=0x4):
[...] table=0, n_packets=659, n_bytes=55965, idle_age=0, priority=1 actions=resubmit(,3)
[...] table=0, n_packets=49, n_bytes=4076, idle_age=824, priority=0 actions=NORMAL
[...] table=0, n_packets=0, n_bytes=0, idle_age=823, priority=2,in_port=2 actions=resubmit(,1)
[...] table=1, n_packets=0, n_bytes=0, idle_age=823, priority=0 actions=resubmit(,2)
[...] table=2, n_packets=0, n_bytes=0, idle_age=822, priority=2,in_port=2 actions=drop
[...] table=3, n_packets=659, n_bytes=55965, idle_age=0, priority=1 actions=NORMAL
[...] table=3, n_packets=0, n_bytes=0, idle_age=821, priority=2,dl_src=fa:16:3f:11:d2:65 \
[...] actions=output:2
```

OpenFlow – Flows I

```
cookie=0x0, duration=2427.523s, table=2, n_packets=18, n_bytes=1719,  
idle_age=434, priority=2,in_port=2 actions=drop
```

- duration
- table
- packets/bytes
- idle_age

OpenFlow – Flows II

```
[..] idle_age=433, priority=2,dl_vlan=1,dl_dst=fa:16:3e:c1:64:05  
actions=strip_vlan,set_tunnel:0x1,output:2
```

- filter
 - priority
 - dl_vlan
 - dl_dst
 - nw_src
 - nw_proto
 - tp_dst
- actions
 - strip_vlan
 - set_tunnel
 - output
 - resubmit

Wer steuert das Ganze?

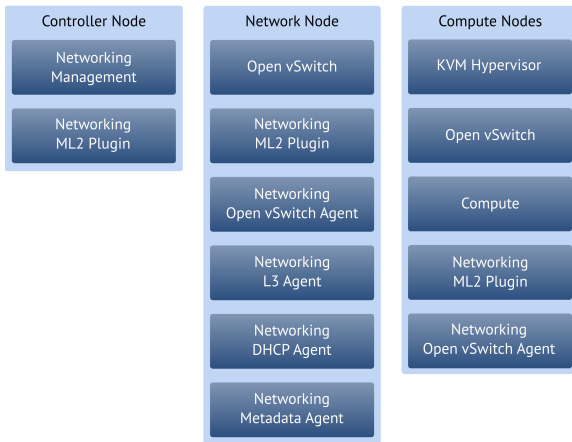
- Neutron Server
- L2-Agent/OpenvSwitch-Agent
- L3-Agent

Wer steuert das Ganze?

- SDN-Controller
 - 'set-controller'
 - 'set-manager'
 - Plugins

Hier entstehen Probleme...

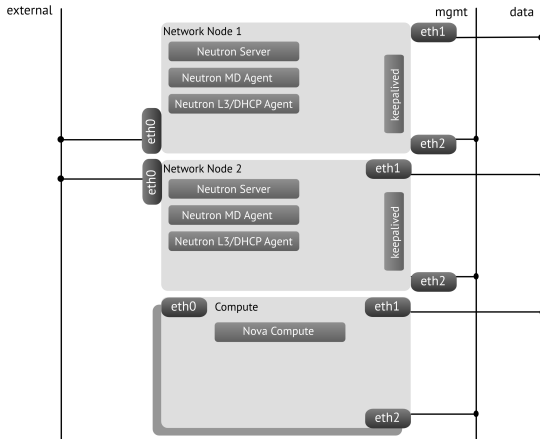
Service Layout



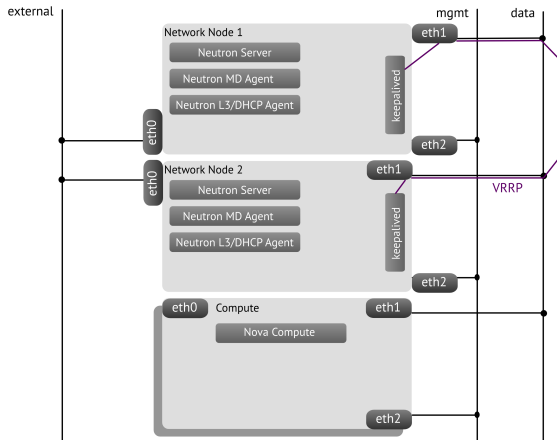
Was tun?

- VRRP
- DVR
- neutron-ha-tool

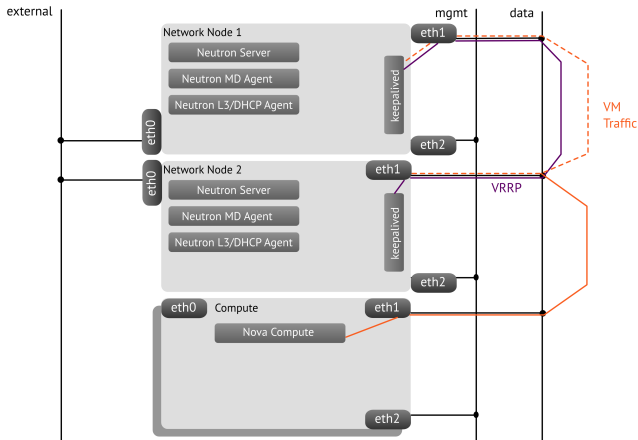
VRRP



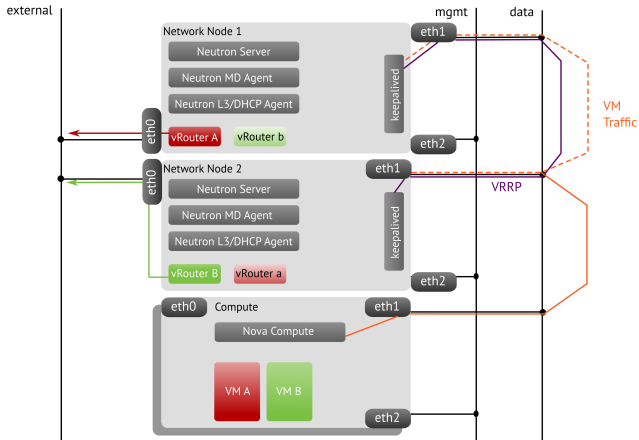
VRRP



VRRP



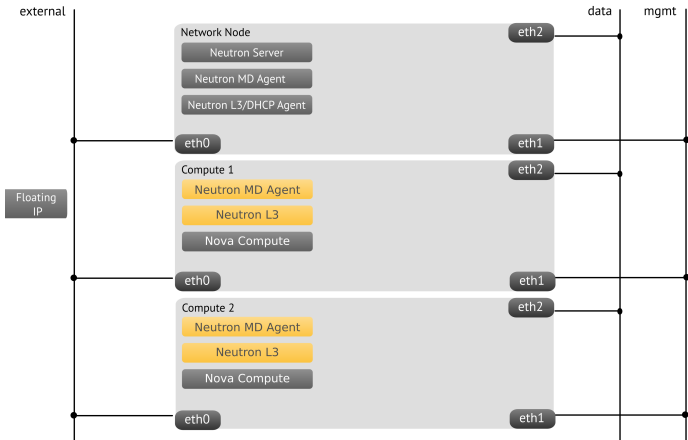
VRRP



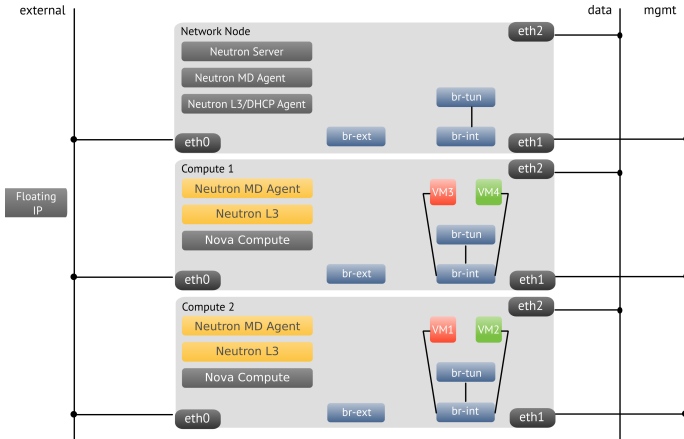
VRRP - Eigenschaften

- keepalived
- contrackd
- ähnlich wie legacy
- Gedacht für "Ausfälle"

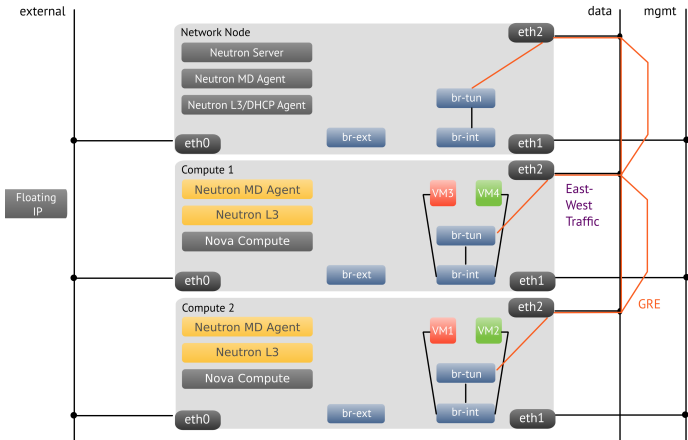
DVR



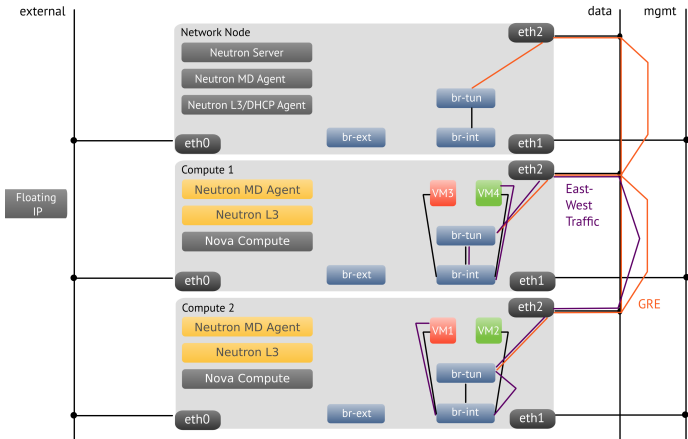
DVR



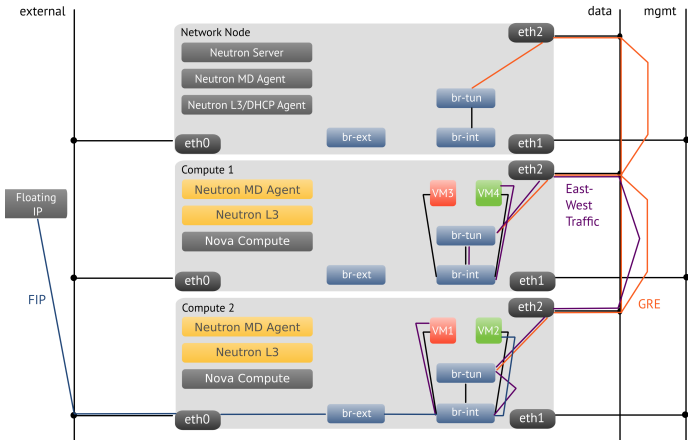
DVR



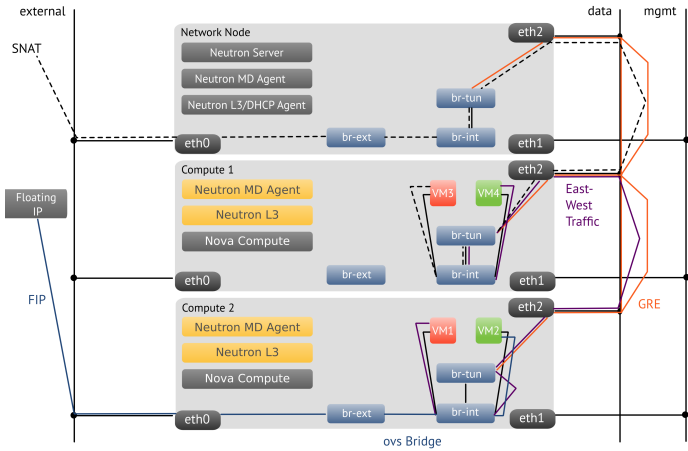
DVR



DVR



DVR



DVR - Eigenschaften

- Computes brauchen "Public"-Netzwerk
- zusätzliche Namespaces
- Netzwerk-Performance
- Seit Juno

neutron-ha-tool

- Pacemaker Installationen
- sichert L3 Agents ab
- geringe Downtime

Vielen Dank für Ihre Aufmerksamkeit!

Bei weiteren Fragen wenden Sie sich bitte an info@b1-systems.de
oder +49 (0)8457 - 931096