NSX-v Distributed Load Balancing - GSG

Preview 6.2.3 (not available for production)

Dimitri Desmidt, VMware, Inc

Goal

- Presentation of DLB
- Understand how to enable DLB + configure DLB VIP
- How to demo DLB

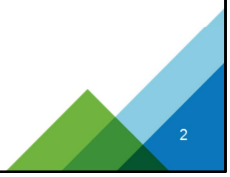

Agenda

1. Presentation of DLB
2. Understand how to enable DLB + configure DLB VIP
3. How to demo DLB
4. Video of a DLB demo
5. Known limitations

3

Agenda

1. **Presentation of DLB**
2. Understand how to enable DLB + configure DLB VIP
3. How to demo DLB
4. Video of a DLB demo
5. Known limitations

Goal of Distributed Load Balancing

- Goal
  - Offer a very scalable and distributed load balancing service
  - Optimized packet flow

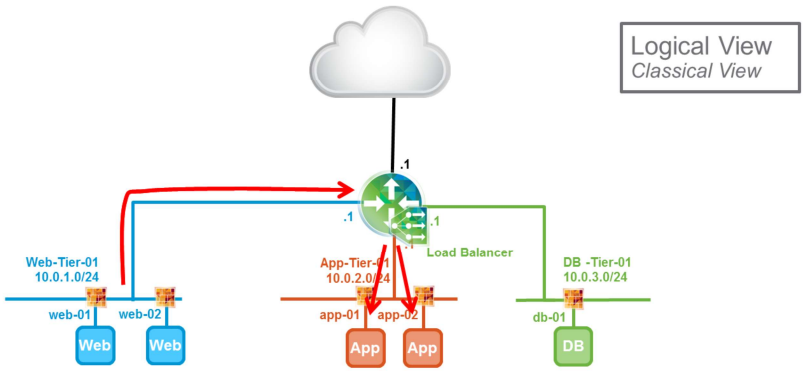

Goal of Distributed Load Balancing

- Goal
  - Offer a very scalable and distributed load balancing service
  - Optimized packet flow

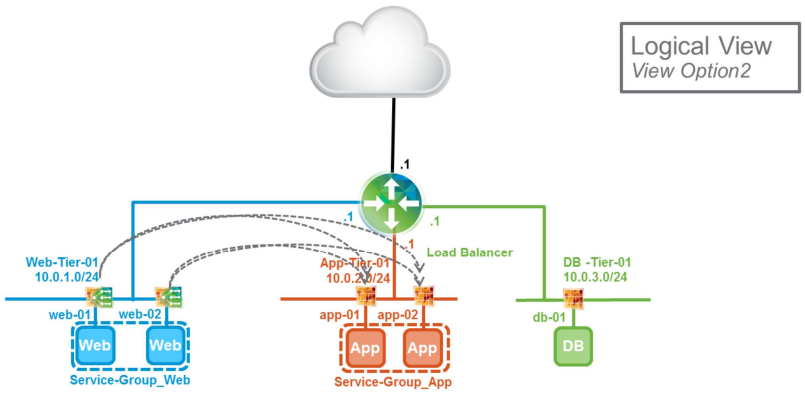

Goal of Distributed Load Balancing

- Goal
  - Offer a very scalable and distributed load balancing service
  - Optimized packet flow

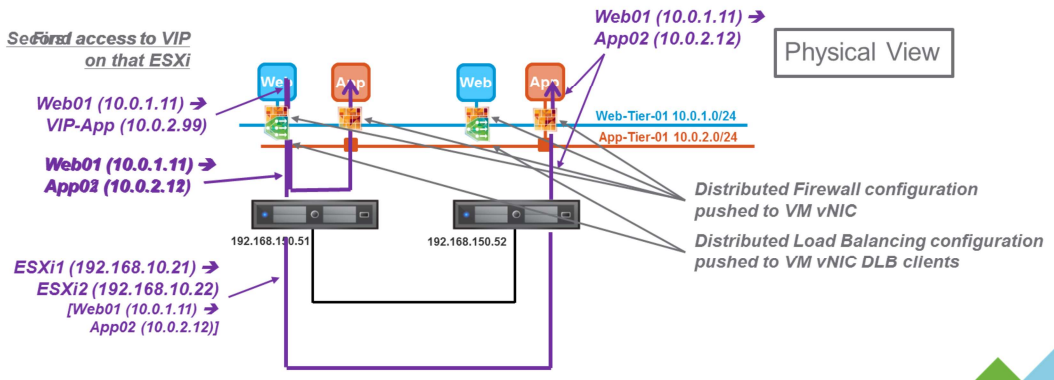

Goal of Distributed Load Balancing

- Goal
  - Offer a very scalable and distributed load balancing service
  - Optimized packet flow

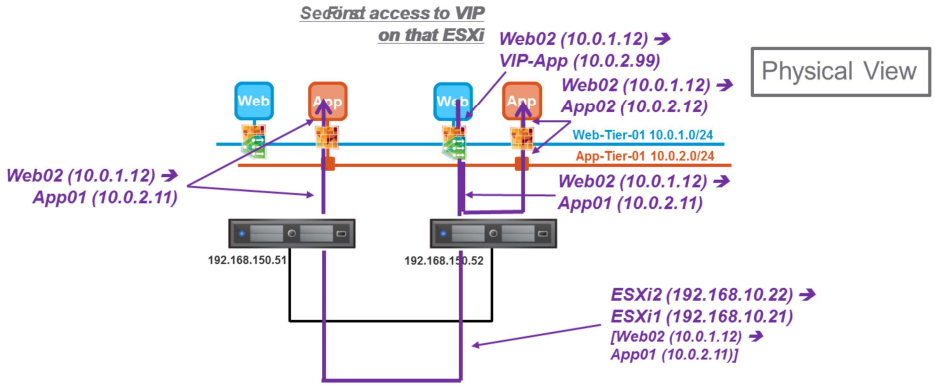

Use Case of Distributed Load Balancing

- Use case
  - East-West load balancing
  - L4 (TCP / UDP) load balancing service
- Currently not a use case
  - North/South load balancing
  - L7 load balancing (SSL off load, URL rewriting, etc)

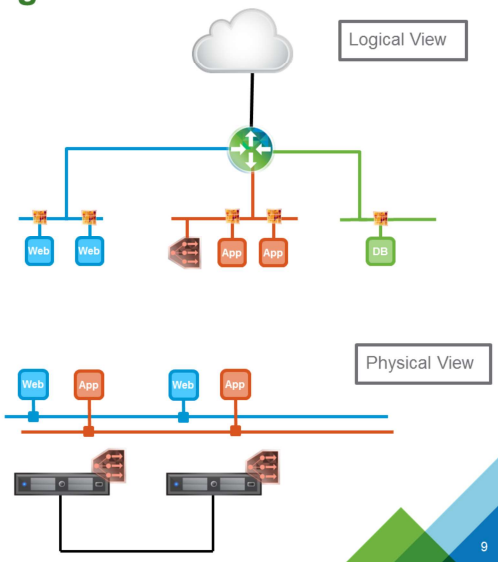

9

Agenda

1. Presentation of DLB
2. **Understand how to enable DLB + configure DLB VIP**
3. How to demo DLB
4. Video of a DLB demo
5. Known limitations

10

Enable DLB (1/4) (Note the steps will simplified for GA)

1. Create a New Service DLB
   - Under "NSX – Service Definitions – Services" – Create a New Service
     a) With "Deployment Mechanism = Host based vNIC"
     b) Service Category: "Load Balancer"

a

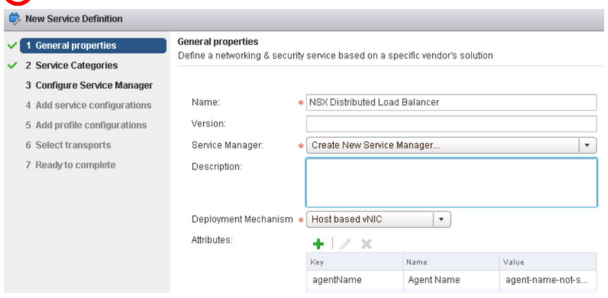

b

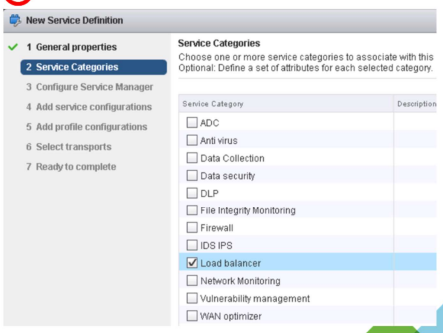

Enable DLB (2/4)

1. Create a New Service DLB
   c) Service Manager: "any name"
   d) Keep other default settings

c

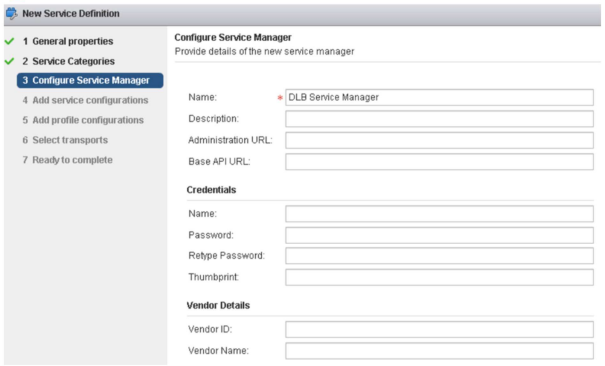

d

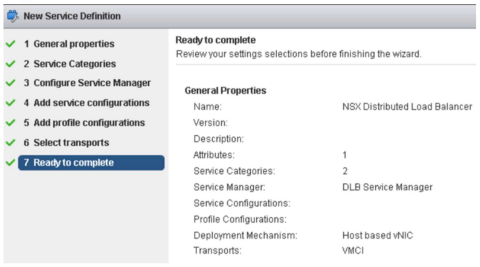

Enable DLB (3/4)

2. Specify what Clusters will have DLB capabilities
   - Under "NSX – Service Definitions – Services" – Edit DLB Service
   - Click on the left "Service Instance" and select Service Instance "NSX Distributed Load Balancer"
   - Select tab "Manage – Deployment"
   - Click "+"

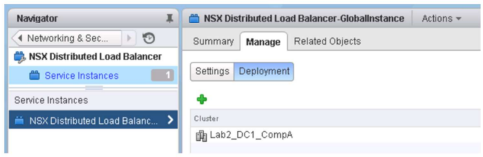

Note: Today only 1 Cluster can be selected.

Enable DLB (4/4)

3. Create a DLB Service
   - Under "NSX – Service Definitions – Services" – Edit DLB Service
   - Click on the left "Service Instance" and select Service Instance "NSX Distributed Load Balancer"
   - Select tab "Related Objects"
   - Click "+"

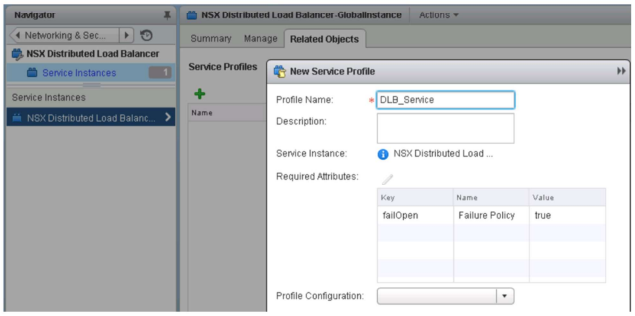

Configure DLB VIP (1/7)

1. Create a Security Group containing the "Clients-VMs" (VMs talking to the DLB VIP App)
   - Under "NSX – Service Composer – Security Group" – Create new group SG_Web containing the Web VMs

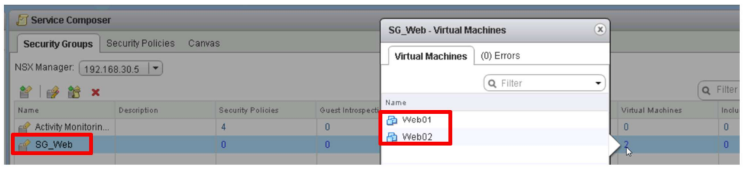

2. Create a Security Group containing the "Servers-VMs" (VMs in the DLB VIP Pool)

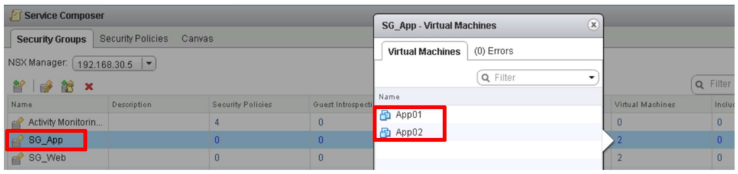

Configure DLB VIP (2/7)

3. Add in DLB_Service the Security Group SG_Web
   - Under "NSX – Service Definitions – Services" – Edit DLB Service
   - Click on the left "Service Instance" and select Service Instance "NSX Distributed Load Balancer"
   - Select tab "Related Objects"
   - Edit DLB_Service, and under "Manage – Applied Object", add the Security Group where are the clients VMs

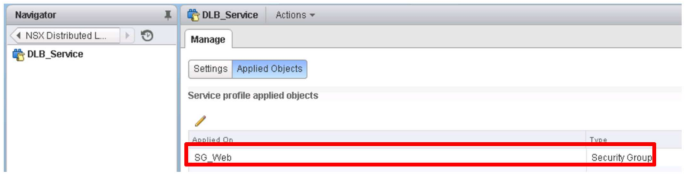

Configure DLB VIP (3/7)

4. Publish DLB Filter to the VMs NIC in the SG_Web
   - Under "NSX – Service Definitions – Services" – Edit DLB Service
   - Click on the left "Service Instance" and select Service Instance "NSX Distributed Load Balancer"
   - Select tab "Manage - Settings", click Publish

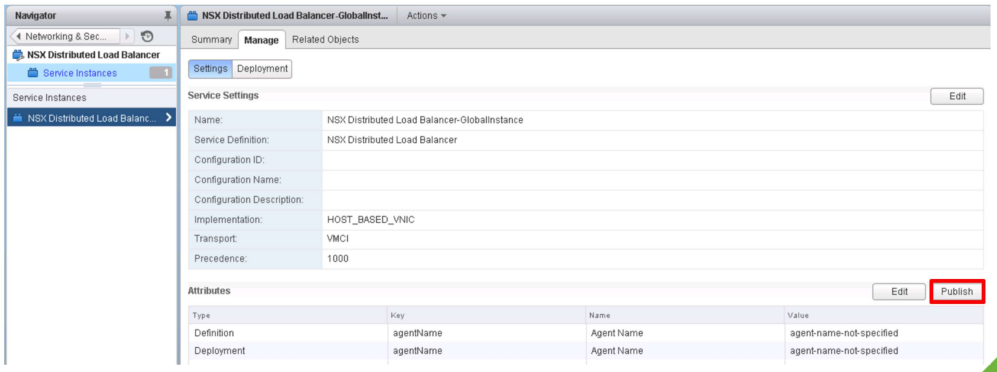

Configure DLB VIP (4/7)

5. Configure the VIP
   - Under "NSX – Firewall", tab "Configuration – Partner Security Services", create a new rule

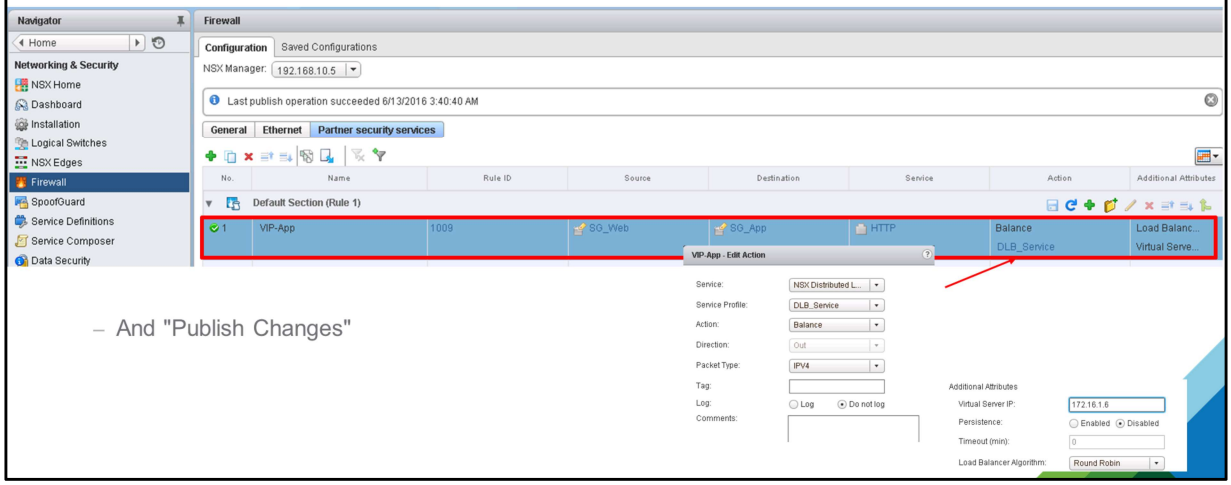

   - And "Publish Changes"

Configure DLB VIP (5/7)

6. Configure Central Healthcheck
   - From any deployed Edge VM with load balancing enabled

     Note: The Edge-VM must have an IP@ with connectivity to the Servers-VM to do their healthchecks

     - Validate Load Balancing is enabled

       Under "NSX – Edge – Manage – Load Balancer – Global Configuration"

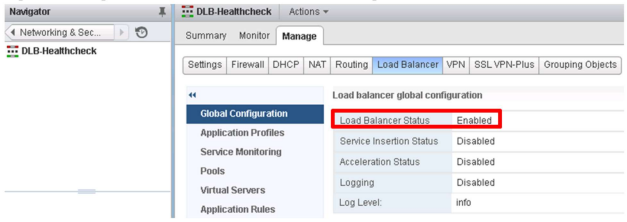

     - Validate it has connectivity to the Servers-VM

       In Edge (console or SSH)

```
NSX-edge-91-0> ping 10.1.2.11
PING 10.1.2.11 (10.1.2.11) 56(84) bytes of data.
64 bytes from 10.1.2.11: icmp_seq=1 ttl=63 time=1.05 ms
64 bytes from 10.1.2.11: icmp_seq=2 ttl=63 time=1.17 ms
```

Configure DLB VIP (6/7)

6. Configure Central Healthcheck – cont.
   - From any deployed Edge VM with load balancing enabled
     - Configure DLB healthchecks

       Under "NSX – Edge – Manage – Load Balancer – Pools"

       Note: The DLB healthchecks must start with name "_DLB_"

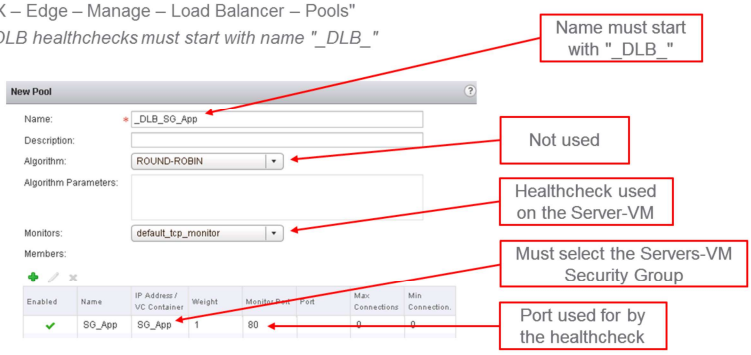

Configure DLB VIP (7/7)

6. Configure Central Healthcheck – cont.
   - Validate healthcheck
     - Via NSX-Mgr UI

       Under "NSX – Edge – Manage – Grouping Objects – IP Sets"

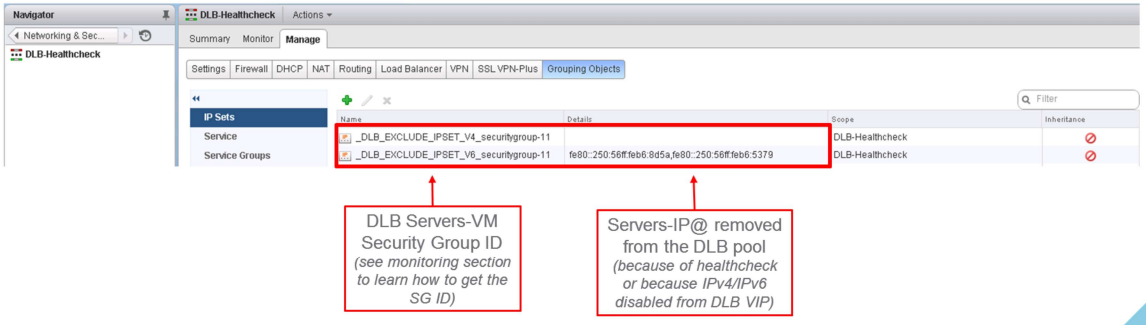

DLB Servers-VM Security Group ID *(see monitoring section to learn how to get the SG ID)*

Servers-IP@ removed from the DLB pool *(because of healthcheck or because IPv4/IPv6 disabled from DLB VIP)*

Monitoring – DLB VIP stats (1/4)

- DLB VIP Statistics
  - Find information to build NSX-Mgr Central CLI DLB stats request
    1. Find ESXi "host-id" (ESXi where the Client-VM is hosted)

From NSX-Mgr:

```
nsxmgr.lab1.vsphere.local> show cluster all
No.  Cluster Name      Cluster Id              Datacenter Name   Firewall Status
1    Cluster-MgtEdge   domain-c9               Lab1              Enabled
2    Cluster-CompA     domain-c7               Lab1              Enabled

nsxmgr.lab1.vsphere.local> show cluster domain-c7
Datacenter: Lab1
Cluster: Cluster-CompA
No.  Host Name       Host Id                 Installation Status
1    192.168.10.21   host-12                 Enabled
2    192.168.10.22   host-15                 Enabled
```

Monitoring – DLB VIP stats (2/4)

- DLB VIP Statistics
  - Find information to build NSX-Mgr Central CLI DLB stats request
    2. Find DLB "VIP filter"

From ESXi (ESXi where the Client-VM is hosted):

```
[root@localhost:~] summarize-dvfilter | grep Web01 -A 20
world 36806 vmm0:Web01 vcUuid:'50 36 6c 22 3b 8f 31 68-9c dc dc f1 0c 30 13 6b'
 port 67108878 Web01.eth0
  vNic slot 2
   name: nic-36806-eth0-vmware-sfw.2
   agentName: vmware-sfw
   state: IOChain Attached
   vmState: Detached
   failurePolicy: failClosed
   slowPathID: none
   filter source: Dynamic Filter Creation
  vNic slot 4
   name: nic-36806-eth0-serviceinstance-1.4
   agentName: serviceinstance-1
   state: IOChain Attached
   vmState: Detached
   failurePolicy: failOpen
   slowPathID: none
   filter source: Dynamic Filter Creation
  dvPort slot 0
   name: 54-sw0a 06 36 50 d7 83 cf 29-b7 4e 21 d1 a2 f3 05 35.dvfilter-generic-vmware.0
   agentName: dvfilter-generic-vmware
   state: IOChain Attached
```

Name of the Client-VM

Look for the filter name of vNic slot 4

Monitoring – DLB VIP stats (3/4)

- DLB VIP Statistics
  - Find information to build NSX-Mgr Central CLI DLB stats request
    3. Find DLB "Pool Security Group"

From ESXi (ESXi where the Client-VM is hosted):

```
[root@localhost:~] vsipioctl getaddrsets -f nic-36806-eth0-serviceinstance-1.4
addrset ip-securitygroup-10 {
ip 10.1.1.11,
ip 10.1.1.12,
}
addrset ip-securitygroup-11 {
ip 10.1.2.11,
ip 10.1.2.12,
```

Name of the filter vNic slot 4

Name of filter Security Group for Servers-VMs

Monitoring – DLB VIP stats (4/4)

- DLB VIP Statistics
  - Get DLB stats centrally per ESXi (ESXi where the Client-VM is hosted)

Under NSX-Mgr

```
nsxmgr.lab1.vsphere.local> show dlb host host-12 filter nic-36806-eth0-serviceinstance-1.4 addrsets ip-securitygroup-11
                           stats all
ip-securitygroup-11:
    in+pass    :                0                0
    out+pass   :                0                0
    in+block   :                0                0
    out+block  :                0                0
    match      :                0
    nomatch    :                0
    uptime     :            11649

                             addr :    pkt[in][pass]    pkt[in][drop]   pkt[out][pass]   pkt[out][drop] tcp curr  tcp hwm
                                     bytes[in][pass]  bytes[in][drop] bytes[out][pass] bytes[out][drop] uptime
addrset ip-securitygroup-11 {
                       ip 10.1.2.11 :              68                0              102                0        0       13
                                                 7973                0             6834                0 218
                       ip 10.1.2.12 :             744                0             1116                0        0       17
                                                87234                0            74772                0 11649
```

Note: The stats are also available directly on the ESXi

```
[root@localhost:~] vsipioctl getaddrsets -f nic-36806-eth0-serviceinstance-1.4 -a ip-securitygroup-11 -s all
```

Monitoring – DLB Pool status (1/1)

- DLB Pool status
  - Get DLB pool status centrally per ESXi (ESXi where the Client-VM is hosted)
    Under NSX-Mgr

```
nsxmgr.lab1.vsphere.local> show dlb host host-12 filter nic-36806-eth0-serviceinstance-1.4 addrsets ip-securitygroup-11
                           validity show
addrset ip-securitygroup-11 {
ip 10.1.2.11    dlb_valid_cnt = 0x05,
ip 10.1.2.12    dlb_valid_cnt = 0x05,
}
```

Note:DLB Pool status is also available via NSX-Mgr UI (see Configure DLB VIP - Configure Central Healthcheck)

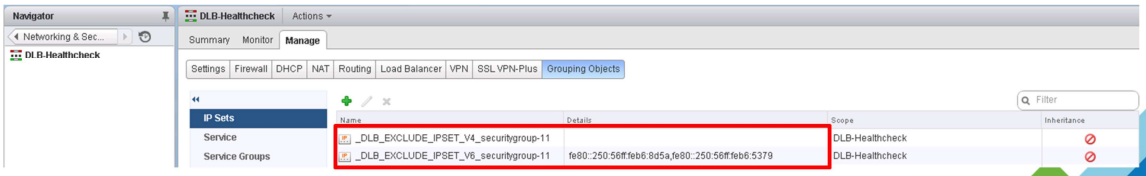

DLB Log Format (1/2)

- Enable DLB Logging
  - Under "NSX – Firewall", tab "Configuration – Partner Security Services", enable logging

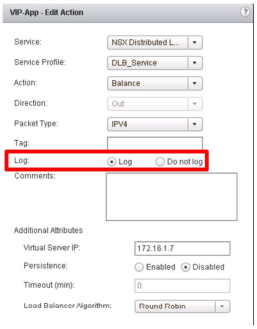

- Validate ESXi syslog configuration
  - Under "vCenter – Hosts and Clusters" – select "ESXi", tab "Manage – Settings – System – Advanced System Settings", and filter with "syslog", configure syslog:

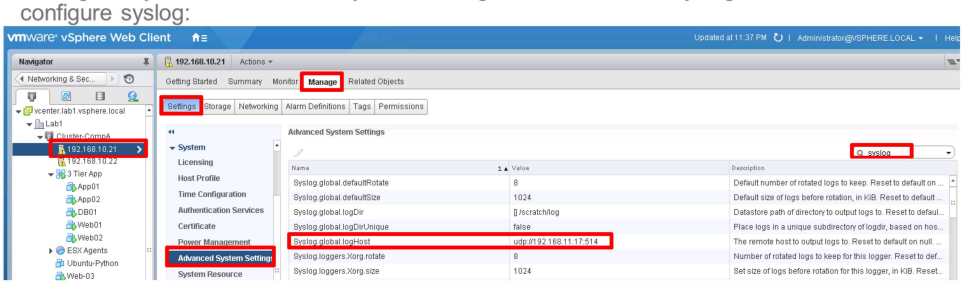

DLB Log Format (2/2)

- DLB Log Format

```
Msg: 2016-07-01T13:22:43.490Z localhost.localdomain dfwpktlogs: 165 INET match RDR
837/15476 OUT 60 TCP 10.1.1.11/39437->10.1.2.12/80 S\0x0a
```

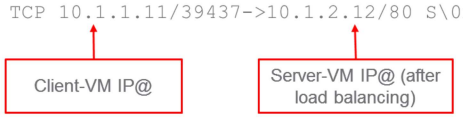

Deep Technical Validation (1/2)

- Validate the Web VMs have the DLB filter (slot 4)
  - SSH to the ESXi hosting the Web VM and run the command "summarize-dvfilter"

```
[root@localhost:~] summarize-dvfilter
<snip>
world 148217 vmm0:Web01 vcUuid:'50 09 f1 b7 88 0e c4 f3-a0 b7 7d 28 2a 09 e9 76'
 port 50331658 Web01.eth0
  vNic slot 2
   name: nic-148217-eth0-vmware-sfw.2
   agentName: vmware-sfw
   state: IOChain Attached
   vmState: Detached
   failurePolicy: failClosed
   slowPathID: none
   filter source: Dynamic Filter Creation
  vNic slot 1
   name: nic-148217-eth0-dvfilter-generic-vmware-swsec.1
   agentName: dvfilter-generic-vmware-swsec
   state: IOChain Attached
   vmState: Detached
   failurePolicy: failClosed
   slowPathID: none
   filter source: Alternate Opaque Channel
  vNic slot 4
   name: nic-148217-eth0-serviceinstance-1.4
   agentName: serviceinstance-1
   state: IOChain Attached
   vmState: Detached
   failurePolicy: failOpen
   slowPathID: none
   filter source: Dynamic Filter Creation
```

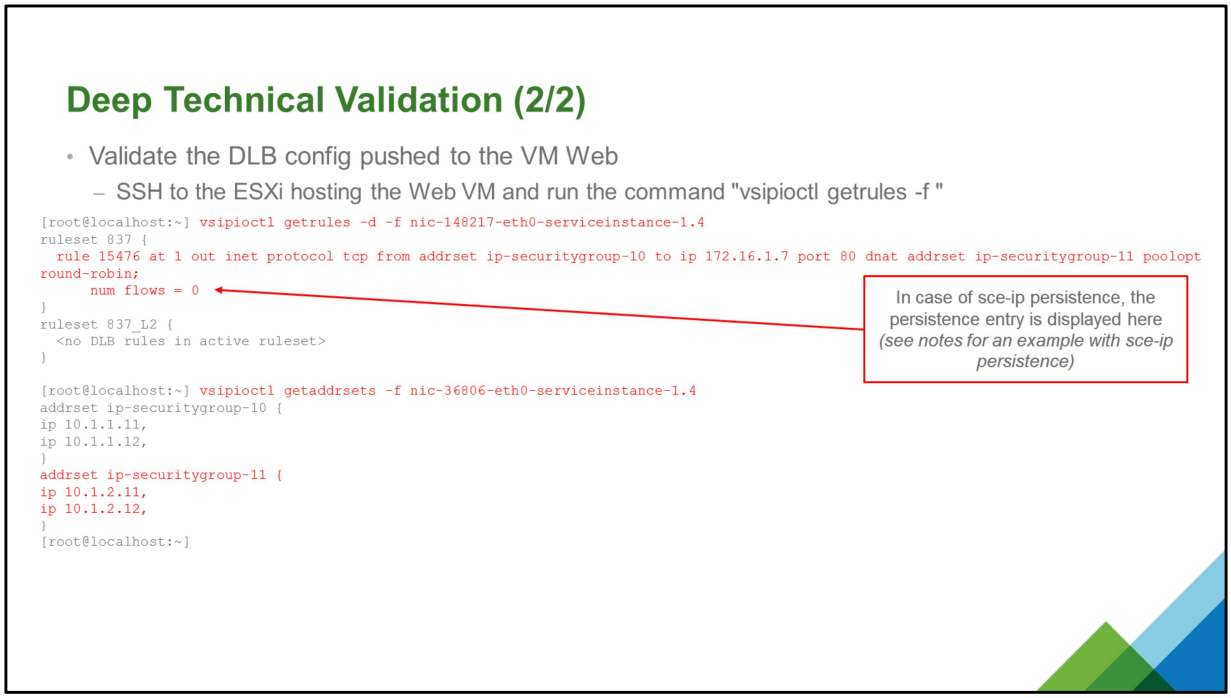


```
[root@localhost:~] vsipioctl getrules -d -f nic-36806-eth0-serviceinstance-1.4
ruleset 837 {
  rule 15476 at 1 out inet protocol tcp from addrset ip-securitygroup-10 to ip 172.16.1.7 port 80 dnat addrset ip-securitygroup-11 poolopt round-robin with log;
      num flows = 0, persist timeout = 60m, timeout remaining = 59:45s, persist addr = 10.1.2.11
}

ruleset 837_L2 {
  <no DLB rules in active ruleset>
}
```

Agenda

1. Presentation of DLB
2. Understand how to enable DLB + configure DLB VIP
3. **How to demo DLB**
4. Video of a DLB demo
5. Known limitations

31

How to demo DLB

1. Deploy a 3-Tier App

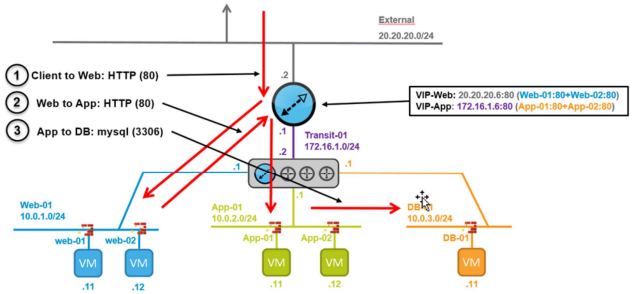

– VMware employees can find an example of 3-Tier App on Vault (https://vault.vmware.com/group/nsx/nsx-poc-resources)

3 Tier App for PoC

To quickly demonstrate at the customer PoC site a 3 Tier App (Web / App / DB). Details on how to install here.

Note: The OVA download link is in the install ppt deck (or directly here).

How to demo DLB

2. Define the **DLB VIP** for VIP App-Tier

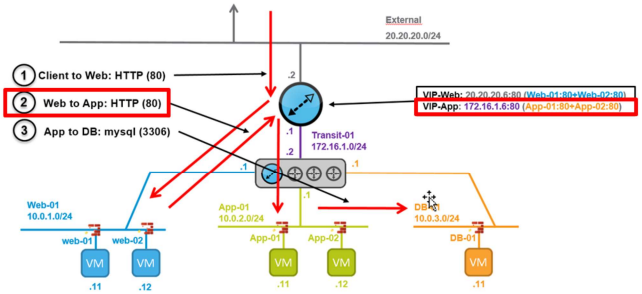

Agenda

1. Presentation of DLB
2. Understand how to enable DLB + configure DLB VIP
3. How to demo DLB
4. **Video of a DLB demo**
5. Known limitations

34

Video of a DLB demo – xxx redo

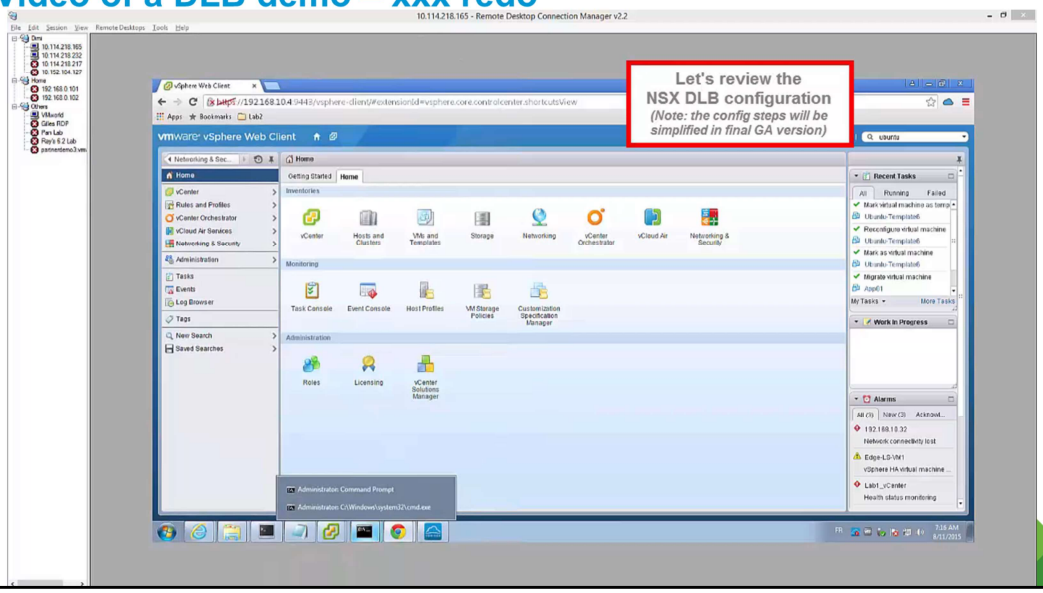

Agenda

1. Presentation of DLB
2. Understand how to enable DLB + configure DLB VIP
3. How to demo DLB
4. Video of a DLB demo
5. **Known limitations**

Known limitations

- VIP in the same subnet as the Client-VM / Client and Server in same subnet => fails (bug 1479932)
- Can add only one cluster in DLB Global Instance (bug 1497514)
- DLB logging format enhanced (bug 1686586)

Backup

Packet Flow (1/2)

- Case with the DGW = DLB:

Client-VM-vNIC (out)	DLB (out)	DLR (out)	ESXi Client VTEP	ESXi Server VTEP	Server-VM-vNIC (in)
Client-IP@ (Client-mac@) => VIP-IP@ (dlr-mac@)	Client-IP@ (Client-mac@) => Server1-IP@ (dlr-mac@) *dest-IP@ is changed but not dest-mac@*	Client-IP@ (dlr-mac@) => Server1-IP@ (Server1-mac@) *sce-mac@ and dest-mac@ are changed*	Encapsulate to ESXi hosting Server1.	De-encapsulate (usual)	Client-IP@ (dlr-mac@) => Server1-IP@ (Server1-mac@)

Server-VM-vNIC (out)	DLR (out)	ESXi Server VTEP	ESXi Client VTEP	DLR (out)	Client-VM-vNIC (in)
Server-IP@ (Server-mac@) => Client-IP@ (dlr-mac@)	Server-IP@ (dlr-mac@) => Client-IP@ (client-mac@) *sce-mac@ and dest-mac@ are changed*	Encapsulate to ESXi hosting Client.	De-encapsulate (usual)	VIP-IP@ (dlr-mac@) => Client-IP@ (Client-mac@) *sce-IP@ is changed but not sce-mac@*	VIP-IP@ (dlr-mac@) => Client-IP@ (Client-mac@)

Packet Flow (2/2)

- Case with the DGW = Edge:

Client-VM-vNIC (out)	DLB (out)	ESXi Client VTEP	ESXi Edge VTEP	Edge (out)	ESXi Edge VTEP	ESXi Server VTEP	Server-VM-vNIC (in)
Client-IP@ (Client-mac@) => VIP-IP@ (edge-mac@)	Client-IP@ (Client-mac@) => Server1-IP@ (edge-mac@) *dest-IP@ is changed but not dest-mac@*	Encapsulate to ESXi hosting Edge.	De-encapsulate (usual)	Client-IP@ (edge-mac@) => Server1-IP@ (Server1-mac@) *sce-mac@ and dest-mac@ are changed*	Encapsulate to ESXi hosting Server1.	De-encapsulate (usual)	Client-IP@ edge-mac@) => Server1-IP@ (Server1-mac@)

Server-VM-vNIC (out)	ESXi Server VTEP	ESXi Edge VTEP	Edge (out)	ESXi Edge VTEP	ESXi Client VTEP	DLR (out)	Client-VM-vNIC (in)
Server-IP@ (Server-mac@) => Client-IP@ (edge-mac@)	Encapsulate to ESXi hosting Edge.	De-encapsulate (usual)	Server-IP@ (edge-mac@) => Client-IP@ (client-mac@) *sce-mac@ and dest-mac@ are*	Encapsulate to ESXi hosting Client.	De-encapsulate (usual)	VIP-IP@ (edge-mac@) => Client-IP@ (Client-mac@) *sce-IP@ is changed but*	VIP-IP@ (edge-mac@) => Client-IP@ (Client-mac@)