

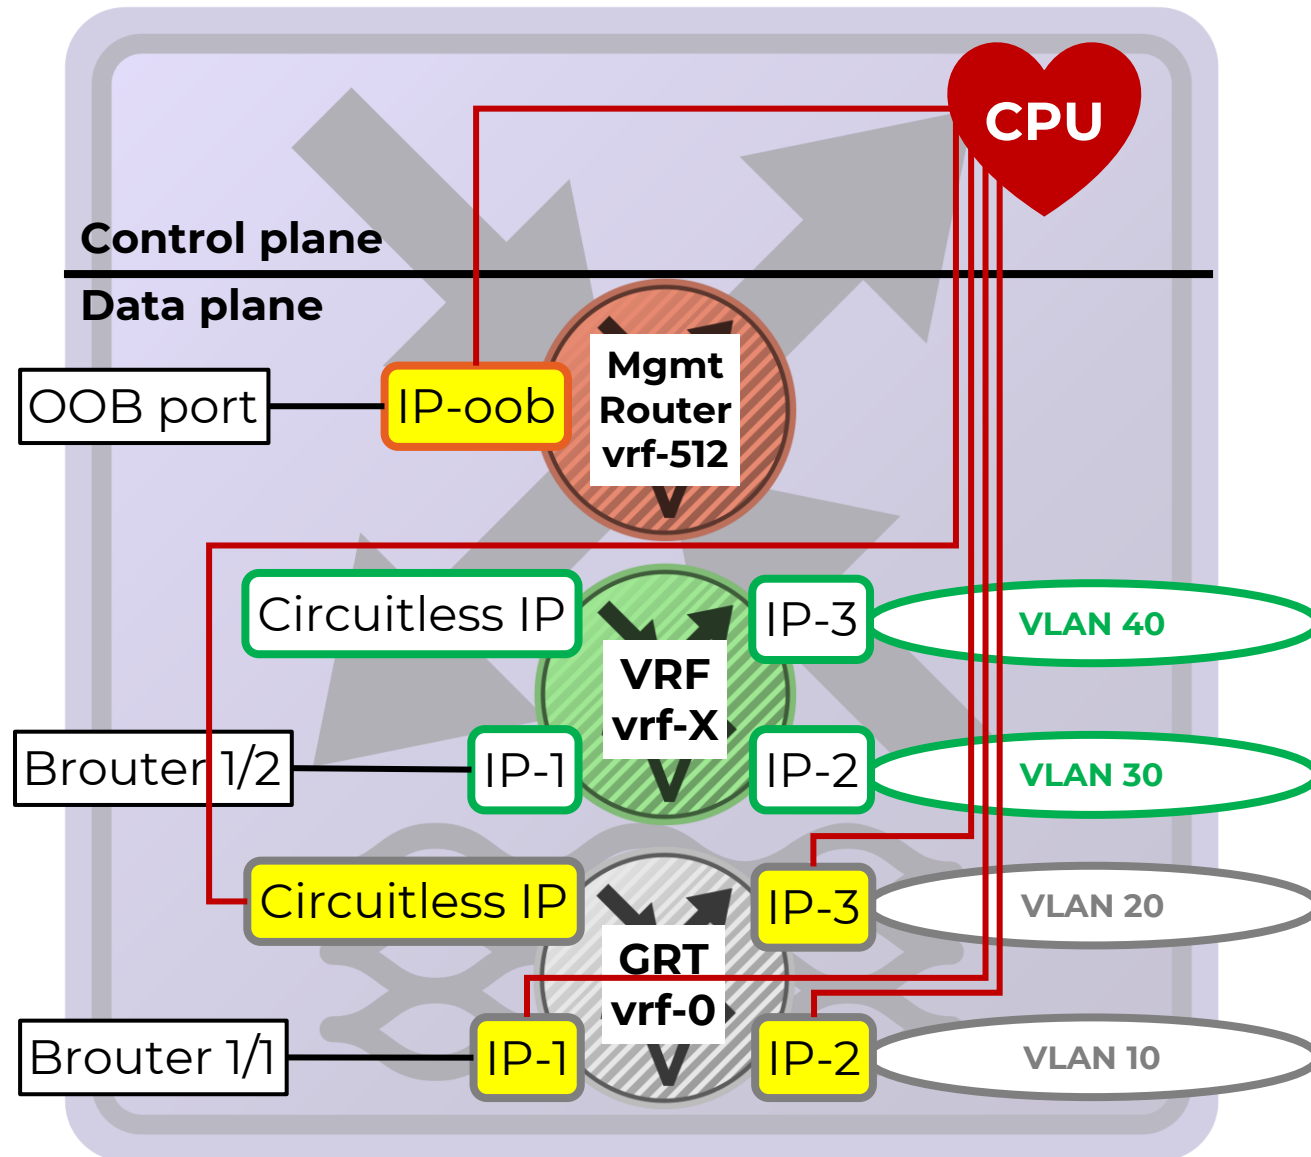


# VOSS 8.2 Segmented Mgmt Stack explained

Ludovico Stevens  
Technical Marketing Engineering  
November 2022


# VOSS Management before 8.2

# VOSS IP mgmt prior to 8.2 (still applies to VSP8600)




- Switch mgmt via
  - Out-of-band: OOB Ethernet port
  - Inband: Any IP address configured on default GRT (vrf-0)
- CPU selects OOB vs. Inband exclusively based on MgmtRouter and GRT routes
  - If OOB and GRT are IP routed together, can result in non-functional asymmetric routing
- Mgmt traffic initiated by switch over inband, selection of source IP ambiguous:
  - GRT IP interface corresponding to next-hop IP for destination non-ISIS route
  - GRT ISIS Source IP for ISIS route
  - Need to configure fixed source IP to use/advertise for some protocols: RADIUS, SNMP, Syslog, LLDP, SONMP, etc..
- NOTE: No OOB port on XA1400, VSP4850, VSP4450
  - VSP4850 support up to VOSS7.1.x only

# VOSS IP mgmt prior to 8.2 (still applies to VSP8600)



# VOSS IP mgmt prior to 8.2 – DVR Leaf

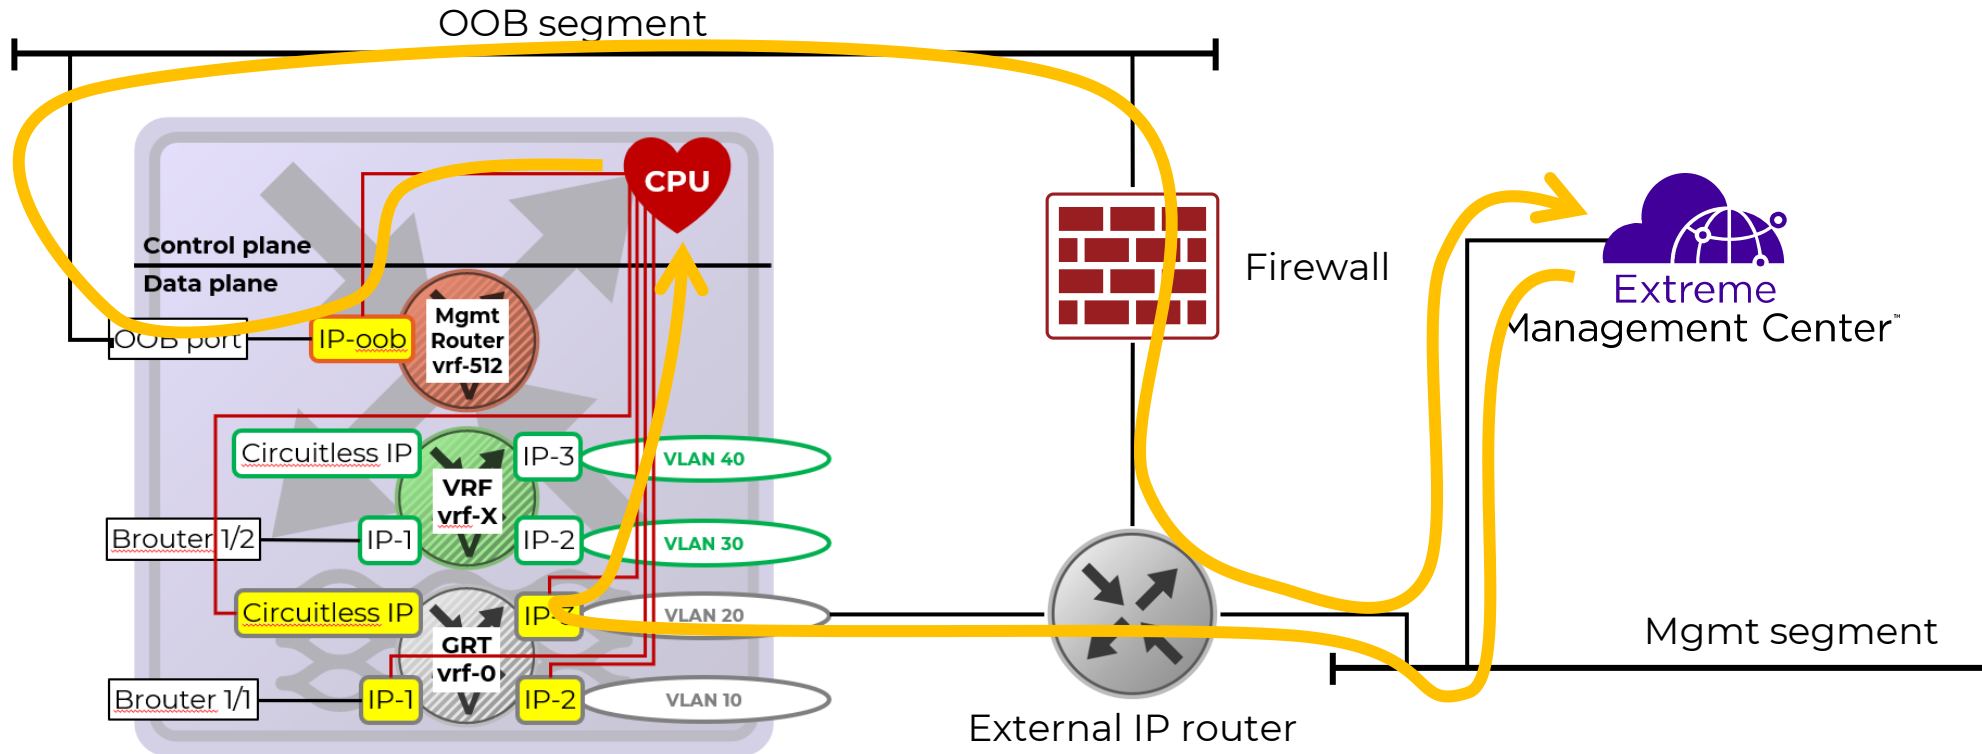


- A DVR Leaf does not actually have a full IP stack for the DVR interfaces
  - The GRT DVR interfaces cannot be used for mgmt

```
DVR Leaf only
router isis
  inband-mgmt-ip <ip>
exit
```

- Instead, a Circuitless IP was created in GRT, but using a new command as the traditional “interface loopback <n>” config context is not available on a DVR Leaf node

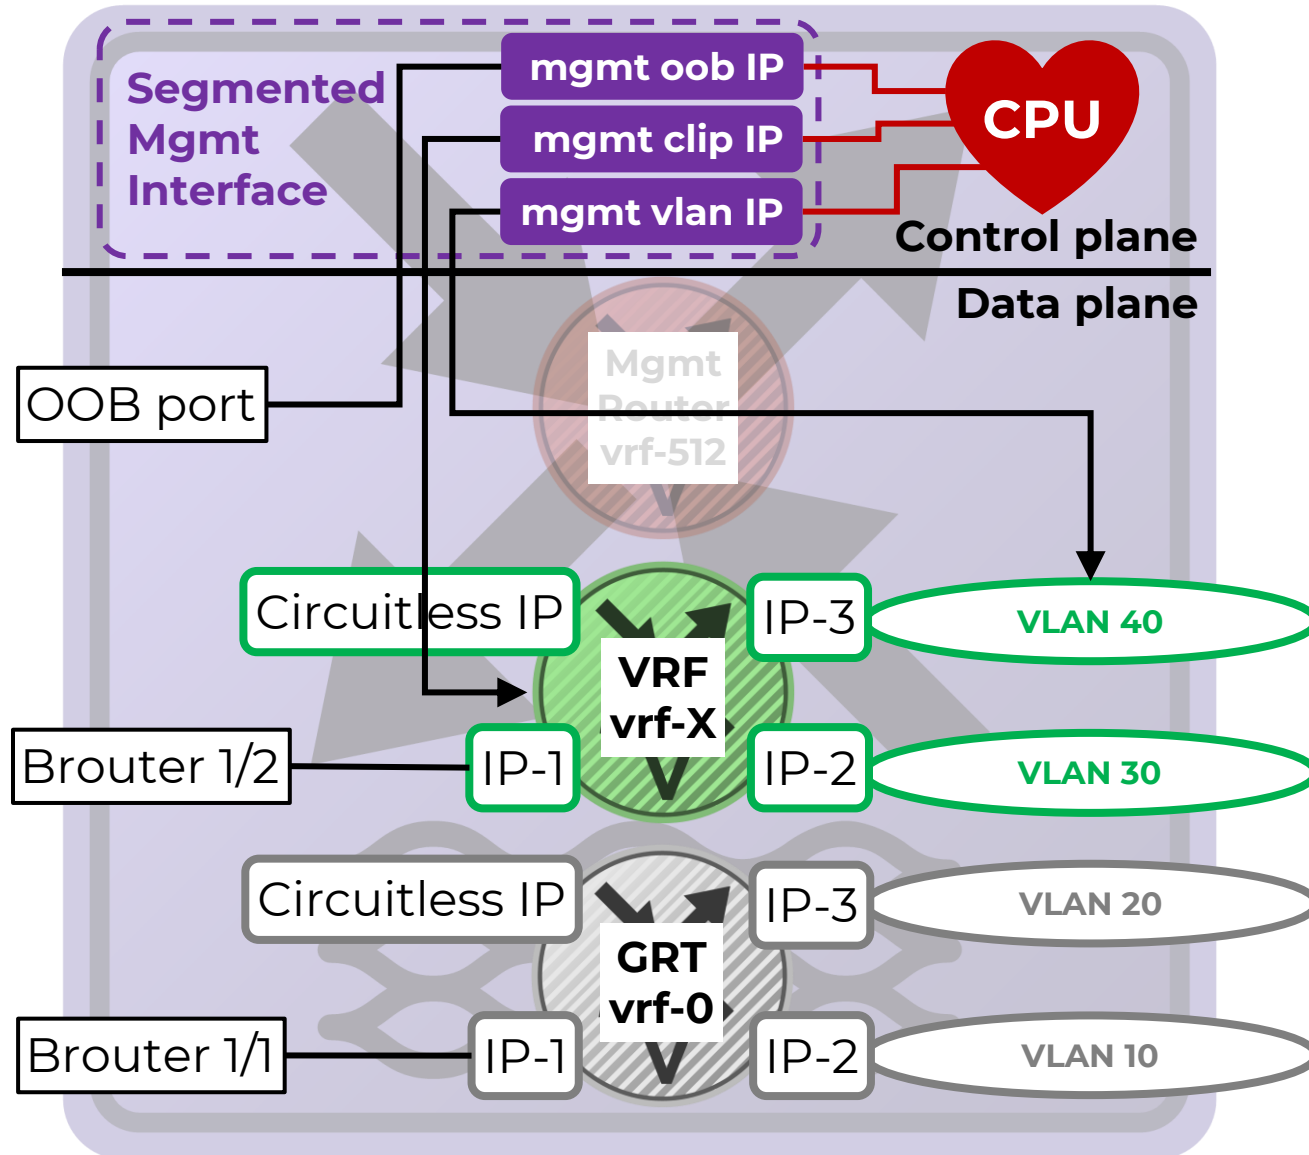
# Pre-8.2 mgmt asymmetrical routing problems



- A mgmt initiated packet (e.g. SNMP Request, or SSH TCP Syn) destined for a VSP inband GRT IP address
- VSP sends response (SNMP Response, or SSH TCP SynAck) via OOB port, if the OOB has a valid IP route
- Communication will fail, for SNMP, SSH, Telnet; but ICMP ping works, so very confusing!
- Recommendation pre-8.2: keep OOB network separate; do not configure a default route in MgmtRouter VRF
- VOSS 8.2 however only solves this problem for TCP based protocols (i.e. not for SNMP, RADIUS, Syslog, etc..)

# VOSS Management from 8.2 onwards


# VOSS IP mgmt 8.2 with Segmented Mgmt Interface



- Switch mgmt via 3 unambiguous IP interfaces:
  - mgmt oob
  - mgmt clip
  - mgmt vlan
- mgmt clip can be assigned to any VRF/GRT
- mgmt vlan can be assigned to any VLAN
- When switch responds to mgmt request, response will now always use same mgmt interface request arrived on
  - No more problems with asymmetrical mgmt routing
- No need to configure source IP for mgmt protocols
- For which mgmt IP LLDP and SONMP should advertise, any of the 3 mgmt interfaces can be selected
- MgmtRouter vrf-512 becomes obsolete
  - CLI show commands & SNMP MIB are maintained and will now show Segmented Mgmt IPs for it
- NOTE: No OOB port on XA1400, VSP4450




# VOSS IP mgmt 8.2 with Segmented Mgmt Interface



- IPv6 also supported (except on XA1400)
- force-topology-ip
  - Determines which mgmt IP used in LLDP advertisements
  - Will advertise both IPv4 and IPv6 if both configured
- Gotchas!
  - if switch booted without a config (ZTF defaults) mgmt vlan will already be created for vlan 4048
  - mgmt IPs must be “enabled”
  - configuring a mgmt IP does not automatically turn off mgmt dhcp-client; remember to turn that off:
    - **no mgmt dhcp-client**

# Segmented Mgmt Interface - quick-config-mgmt




```
VSP:1#% quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the configuration.
Once the basic parameters are configured, additional configuration can
proceed using other management interfaces. Press q to abort at any time.
Management interface types:
  1 - Out of band management port
  3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]:
```

- quick-config-mgmt
  - Integrated interactive script to configure segmented mgmt IP interfaces
  - Useful if starting afresh with 8.2 or later

- IPv4 only is supported
- Can setup only one interface at a time
- Management CLIP is not supported

# Segmented Mgmt Interface – DHCP Client



```
mgmt oob
  enable
exit
mgmt dhcp-client oob
```

```
New Zero-Touch Defaults (8.2)
mgmt oob
  enable
exit
mgmt vlan 4048
  enable
exit
dhcp-client cycle
```

```
mgmt vlan <vid>
  enable
exit
mgmt dhcp-client vlan
```

- New segmented mgmt interface comes with new DHCP Client
  - Only for mgmt vlan and mgmt oob
- Create and enable the mgmt interface type then enable dhcp-client on it
- In practice this will only be used when the VSP boots up in the new 8.2 and 8.3 zero-touch factory defaults, which introduce the concepts of the onboarding Private-VLAN (4048) and ETREE I-SID (15999999) and where all VSP ports are enabled and members of PVLAN 4048
  - This new zero-touch “default” mode applies when the VSP is booted without any config file
  - NOTE: this does not apply to the old “boot config flag factorydefaults” which produces the original default config where all ports are disabled and members of VLAN 1
  - dhcp-client cycle mode will alternatively try and obtain a DHCP IP on either the oob or vlan interfaces

# Segmented Mgmt Interface – Link-Local IP




Both the OOB Mgmt and inband Mgmt VLAN interfaces can auto-assign an IPv4 link-local address


- 169.254.xx.1/16 on Mgmt OOB and/or 169.254.xx.4/16 on Mgmt VLAN
  - where xx is second last Byte from Base MAC address of the node
- No config required
- Link-local IP address goes away if a manual static IP is configured

Example:

- BaseMacAddr: 00:51:00:f9:20:00  
→ 0x20 = 32
- OOB Mgmt IP: 169.254.32.1/16
- Mgmt VLAN IP: 169.254.32.4/16



# Segmented Mgmt Interface – convert command



- Introduced in VOSS 8.5.0.0
- Allows an existing mgmt IP to be switched to a different IP and/or on a different VLAN-id, I-SID or VRF
- Automatic rollback if user is not able to connect to new IP within configurable rollback time

```
VSP:1(config)#% mgmt vlan
VSP:1(config:vlan)#% convert [vlan <vid>] [ports-tagged <ports>] [ports-untagged <ports>] [i-sid <i-sid>] [ip <addr/mask>] [gateway <ip>] [rollback <secs>]
- Or -
VSP:1(config)#% mgmt clip
VSP:1(config:clip)#% convert [vrf <name>] [ip <addr/mask>] [gateway <ip>] [rollback <secs>]
- Or -
VSP:1(config)#% mgmt oob
VSP:1(config:oob)#% convert [ip <addr/mask>] [gateway <ip>] [rollback <secs>]
```

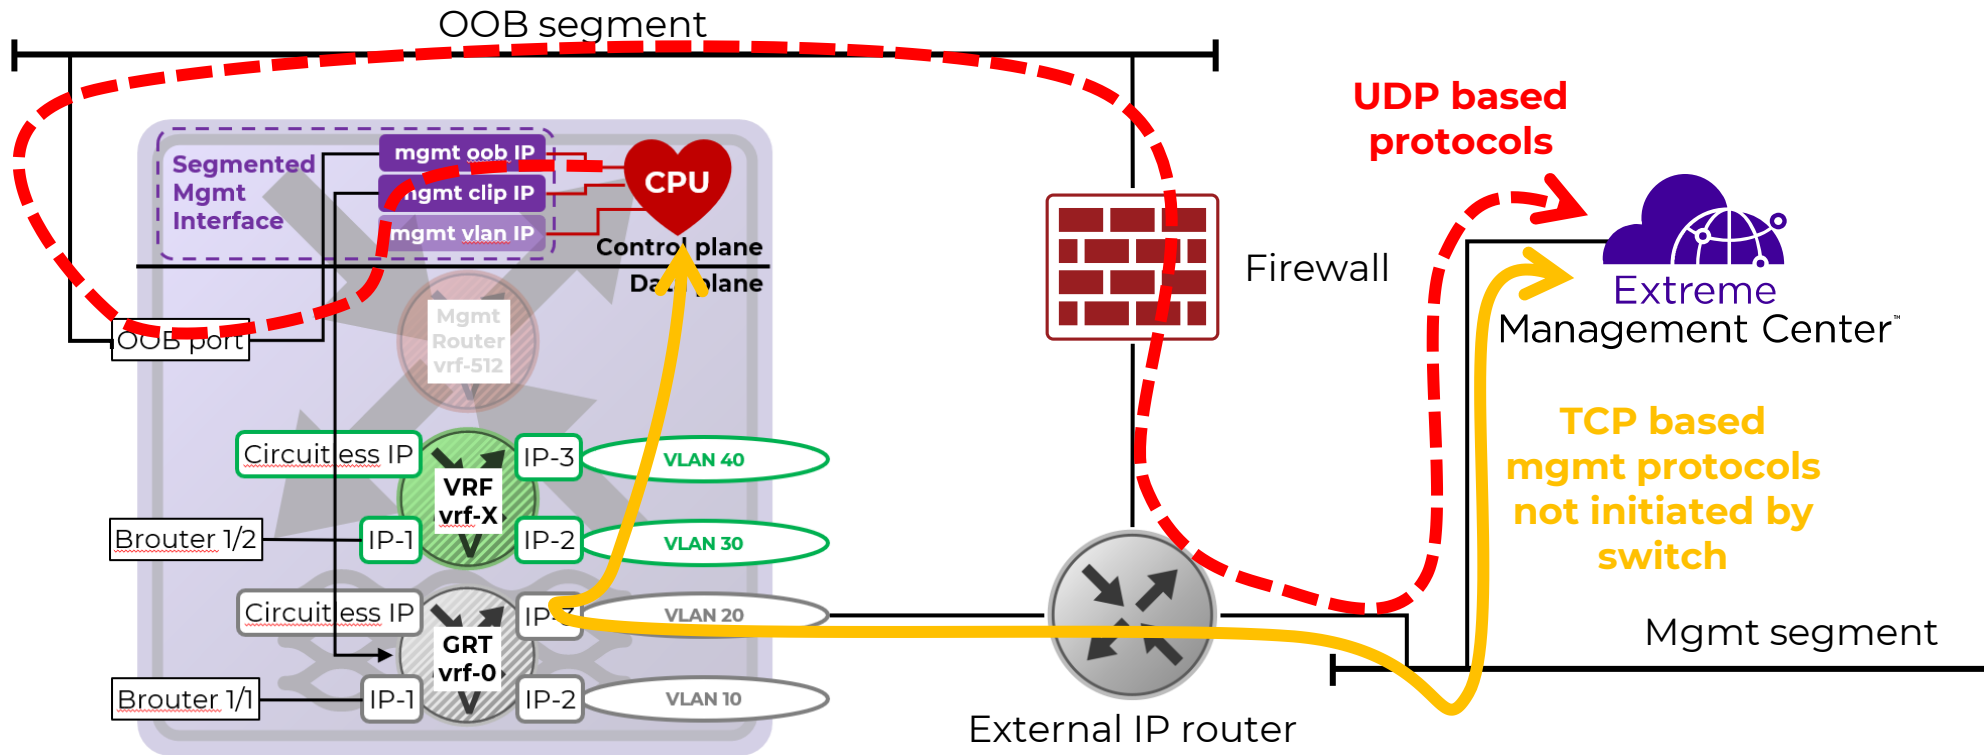
**WARNING: The existing mgmt interface will be deleted and re-created with the given parameters, please reconnect to the switch and issue 'mgmt convert-commit' command before the 120 second rollback timer expires.**  
Continue with this operation (y/n) ?

<SSH/Telnet connection is lost>  
<Re-connect to newly configured IP (including new VLAN/I-SID/VRF if one was set/changed)>

**Mgmt convert: Please issue 'mgmt convert-commit' in the remaining XX seconds before rollback timer expires otherwise mgmt XXXX config change will be reverted**

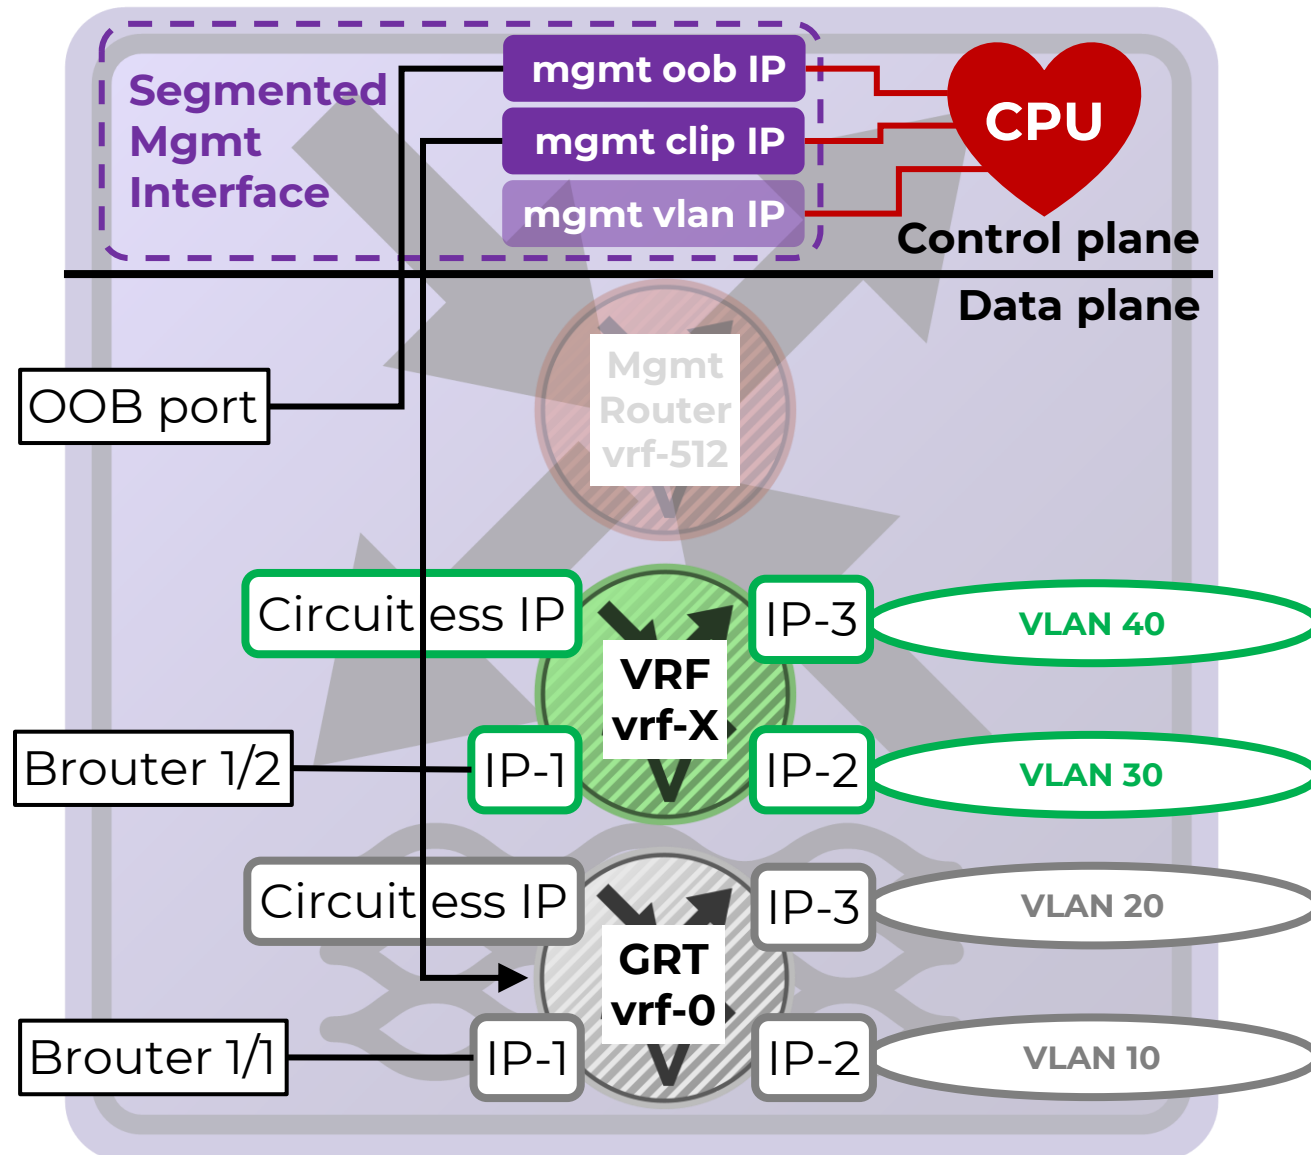
```
VSP8000-1:1(config)#% mgmt convert-commit
```

# VOSS IP mgmt 8.2 – reduced asymmetrical routing (TCP only)




- Segmented mgmt interfaces use Linux VR contexts
  - If a mgmt request is received on mgmt clip, the switch response will always use the same mgmt interface
  - However, this only works for TCP based protocols, not initiated by the switch, like SSH, Telnet, HTTP, HTTPS, etc..
- For switch-initiated messages (TCP or UDP) and all UDP based protocols (SNMP, Syslog, RADIUS, etc..), the same issues of asymmetrical routing persists and per mgmt interface routes are inspected and the best route with the lowest metric will determine the outgoing segmented mgmt interface
  - Default metric weights: clip = 100, vlan = 200, oob = 300
  - Static routes can only be configured for mgmt vlan & mgmt oob (and different weight can be configured)
  - For mgmt clip, the IP routes of the associated VRF/GRT apply (always with weighting 100)

# Segmented Mgmt Interface: L3 BEB / L3 Router



- If the VSP is a L3 BEB (or a non-Fabric IP router), inband management should use mgmt clip
  - The mgmt vlan interface “should” not be used
- The mgmt clip interface can be associated with the GRT (as before) but can now also be easily associated with any VRF
  - If IP Shortcuts or L3VSN is enabled on the GRT/VRF, the mgmt clip will automatically be redistributed even if redistribution of directs is not enabled
- Note that management via a GRT Circuitless IP was already best practice pre-8.2 for L3 BEBs
- The mgmt oob interface can also be used

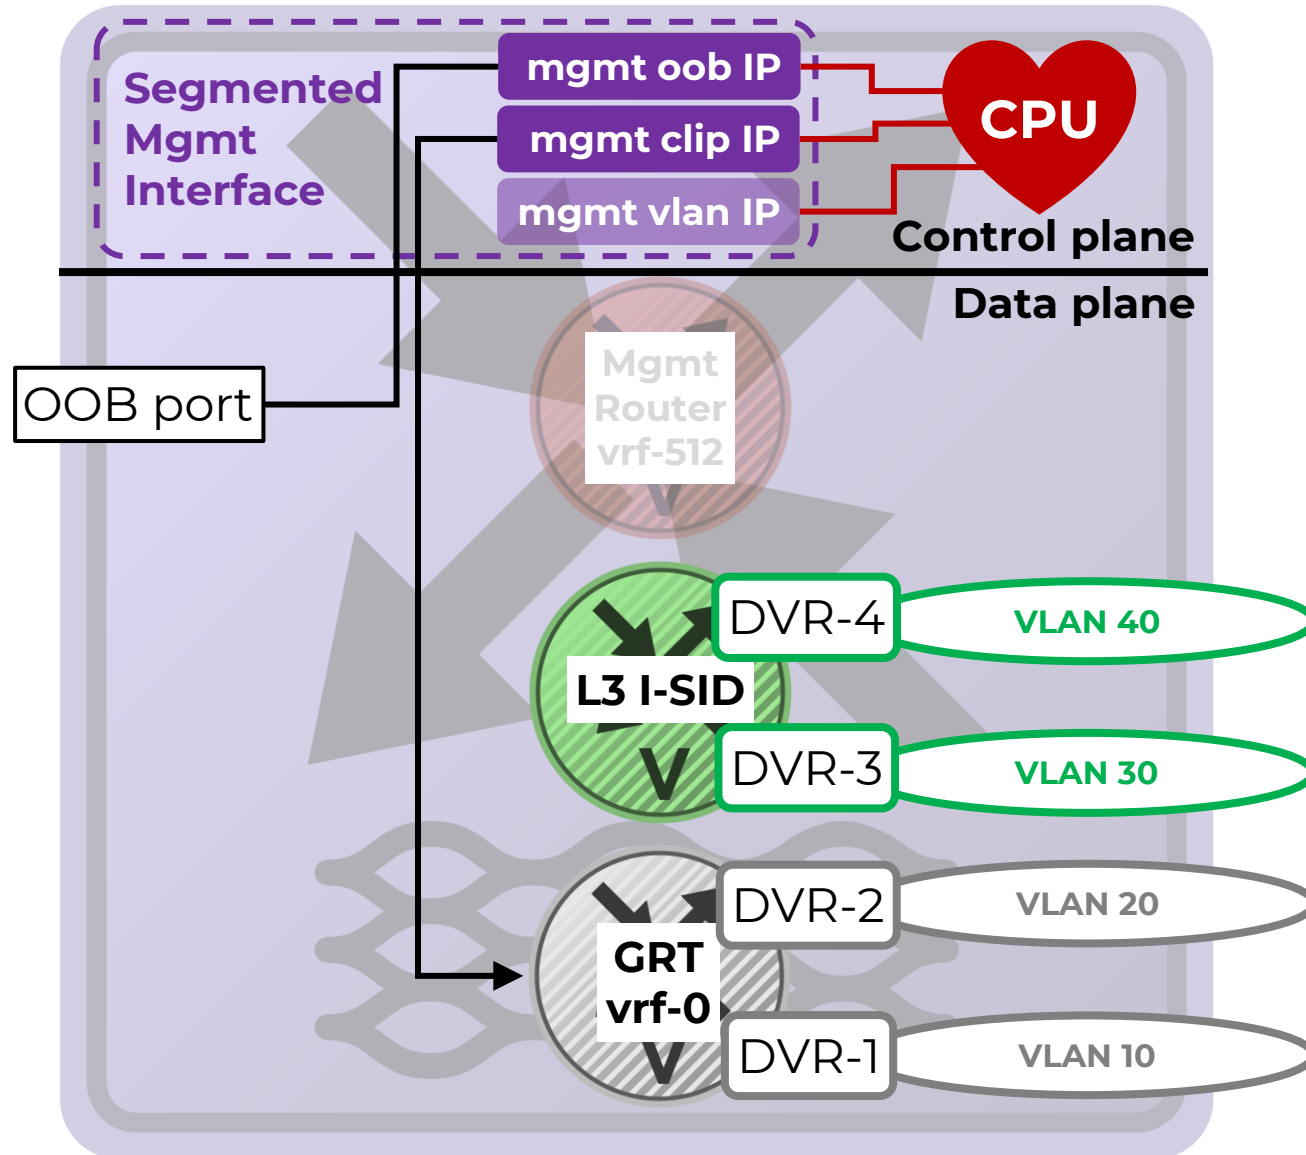
# Segmented Mgmt Interface: L2 BEB / L2 Switch



- If the VSP is a L2 BEB (or non-Fabric L2 switch), inband management should use mgmt vlan
  - The mgmt clip can however still be used on a L2BEB, on the GRT, but it will require IP enabling SPBM
  - On a non-Fabric L2 switch, the mgmt clip cannot really be used as there are no IP interfaces to route to/from that clip
    - It would require turning the VSP switch into a L3 switch
- The mgmt vlan interface can be associated with any platform VLAN already created on the switch
  - The VLAN can of course be made into a fabric wide L2VSN by assigning an I-SID to it
- The mgmt oob interface can also be used

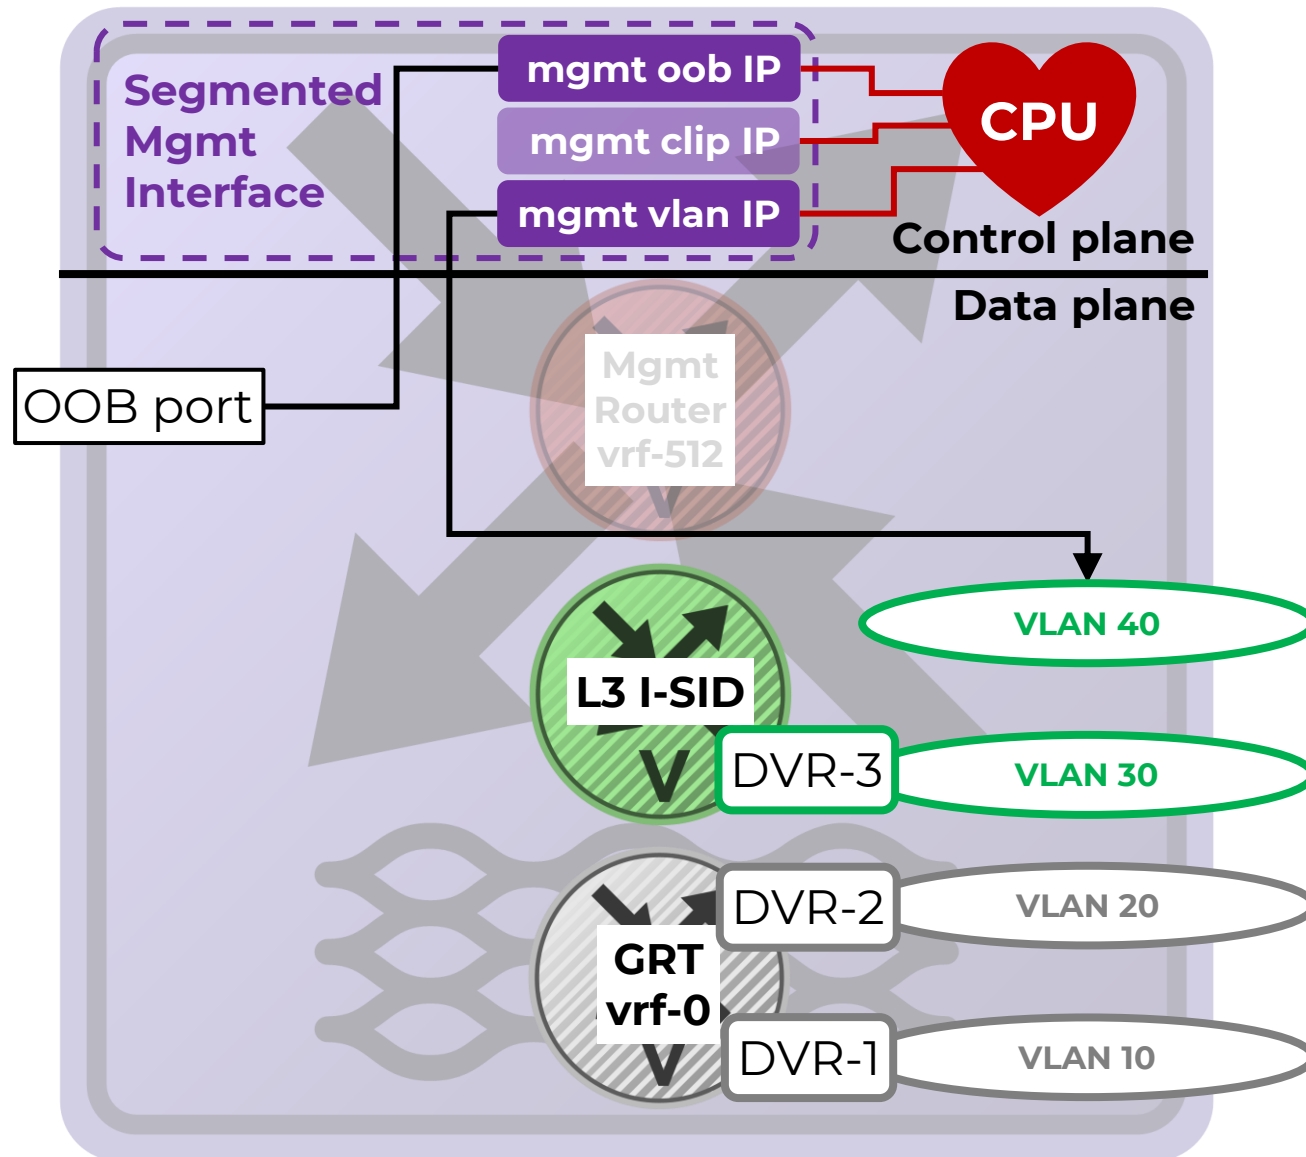


# Segmented Mgmt Interface: DVR Leaf, CLIP mgmt



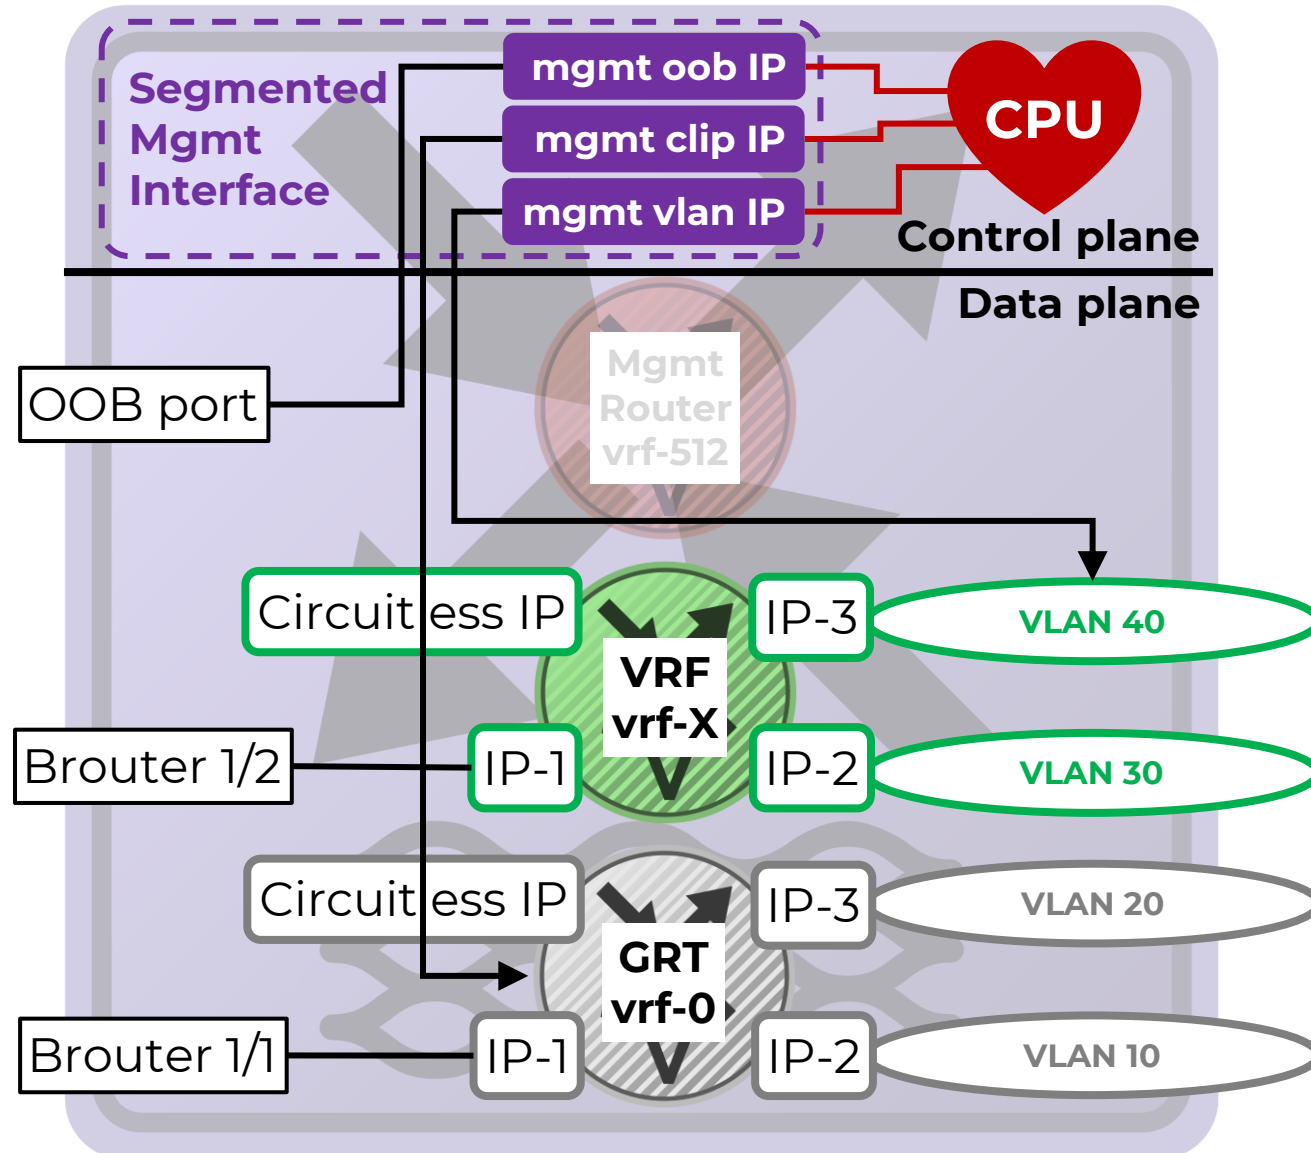
- A DVR Leaf is a special case as it is a L3 BEB in the data plane but a L2 BEB from a configuration management perspective
- If mgmt will be done over the GRT then mgmt clip can be used
  - This will be equivalent to the pre-8.2 inband-mgmt-ip
- However, on a DVR Leaf, the mgmt clip can only be associated with GRT
  - As a DVR Leaf does not have any locally configured VRFs
- A DVR Leaf already has SPBM IP enabled, so no extra config is required if a mgmt CLIP is created
- The mgmt oob interface can also be used

# Segmented Mgmt Interface: DVR Leaf, VLAN mgmt




- A DVR Leaf is a special case as it is a L3 BEB in the data plane but a L2 BEB from a configuration management perspective
- Configuration of mgmt VLAN on DVR Leaf changed in VOSS 8.5
  - No need to create any VLANs, which is not allowed on a DVR Leaf anyway
  - Simply create the mgmt VLAN against a L2 I-SID
    - `mgmt vlan i-sid <isid>`
- For network wide VLAN management the above is sufficient
- If the network management is IP routed on a VRF (or the GRT) then the mgmt VLAN I-SID can be attached to an IP VLAN routed interface on the DVR Controllers
  - The DVR Controllers should have an IP VRRP interface for this same mgmt I-SID associated with the VRF (or GRT) used for management
    - Do not configure DVR on this VLAN !
  - Local DVR interfaces on the same mgmt VRF (or GRT) on the DVR Leaf will not be IP routed directly to the local mgmt i-sid but will be able to reach it via the DVR Controller
- The mgmt oob interface can also be used

# Segmented Mgmt Interface: L3 BEB special cases



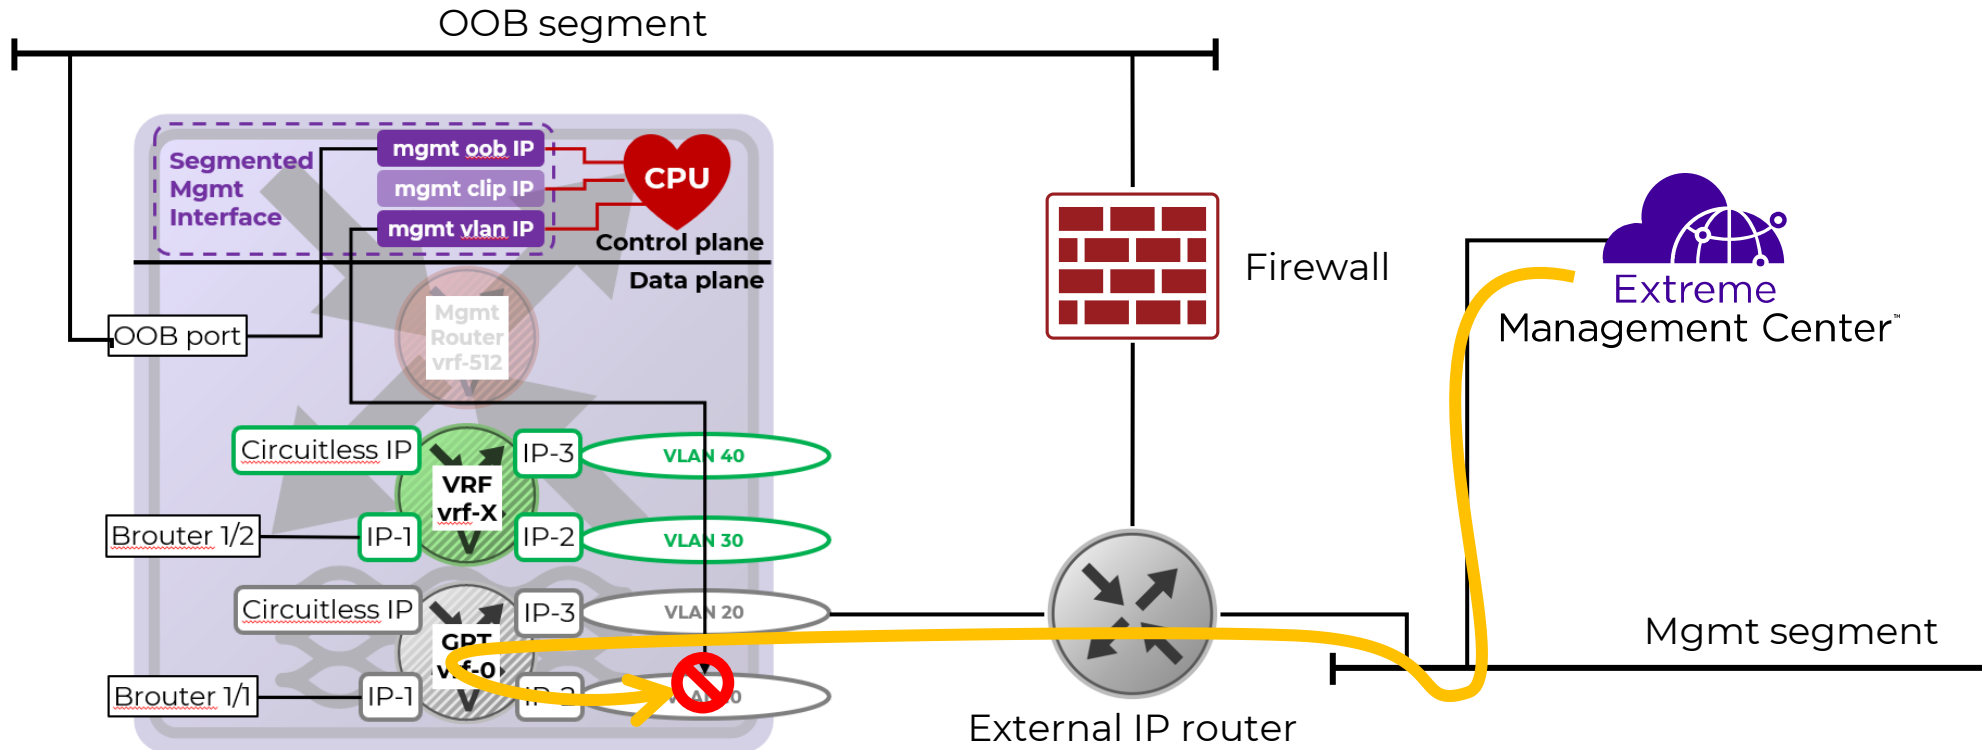
- In some cases, it might be necessary to configure mgmt vlan even on a L3 BEB:
  - XA1400 or VSP running Fabric Extend over a dedicated VRF and it is desired to reach the switch on that VRF from the Internet (e.g., Cloud-IQ) or from WAN underlay
  - VSP7400 or VSP4900 with FIGW VM and it is desired to SSH/FTP the VM from the VSP host switch
  - In both the above cases a mgmt clip also exists for normal inband mgmt
  - If a mgmt vlan is created on a VLAN which already has an IP address in the GRT/VRF, then the mgmt vlan IP must be made the same as that IP address
- All three mgmt interfaces can be used in this example

# Segmented Mgmt Interface: L3 BEB mistake to avoid!



- For a L3 VSP (BEB or non-Fabric), management via a GRT Circuitless IP was already best practice pre-8.2 for L3 BEBs
- However, some customers may not have followed that best practice, and used a GRT VLAN IP for managing all of their L3 BEBs and L2 BEBs alike
  - This did work pre-8.2
- However, this may NOT work properly on a L3 BEB with the new Segmented Mgmt interface
  - The mgmt vlan IP can only be reached if traffic destined to it enters the VSP switch on the same VLAN
  - If the traffic destined to it enters the switch on a different IP interface of the same GRT/VRF, then it will not get IP routed to the mgmt vlan IP destination
  - Of course, if an external Firewall IP routes onto the mgmt vlan segment then it will work fine


# Segmented Mgmt Interface: L3 BEB mistake to avoid!



- In this example, the VSP mgmt vlan IP cannot be reached because the mgmt packet entered the switch on a different IP interface
  - This is true even if a routing VLAN IP is already also configured on the underlying platform VLAN and IP routing is possible between both IP interfaces
- This is a mistake. As the VSP is clearly a L3 router and would have to route traffic to the mgmt vlan subnet, mgmt clip must be used


# Migration to 8.2

# Migration of L3 BEB / L3 Router




```
interface loopback <id>
  migrate-to-mgmt
exit
```

- “migrate-to-mgmt” command is available since VOSS 7.1.3, 8.0.1 and 8.1.0
- save config and upgrade



- NOTE, after the upgrade the GRT CLIP will have gone
- If an ISIS Source IP was in use, re-create a new GRT CLIP (using a different IP address) and assign that as the new ISIS Source IP
  - This operation can also be done before the upgrade by creating a second CLIP on GRT and moving the ISIS Source IP to that second CLIP, while the first CLIP is set to migrate-to-mgmt and will disappear after the upgrade
- As of 8.2 an ISIS Source IP is not mandatory but is still recommended if using IP Shortcuts and will be required again by DVR-One-IP

# Migration of L2 BEB / L2 Switch




```
interface vlan <vid>  
  migrate-to-mgmt  
exit
```

- “migrate-to-mgmt” command is available since VOSS 7.1.3, 8.0.1 and 8.1.0
- save config and upgrade

- NOTE, after the upgrade the GRT VLAN IP will have gone
- If the VSP has more than 1 IP address on more than 1 VLAN before the upgrade, then think twice; the VSP is probably a L3 BEB and should be managed via a CLIP instead!
- If Application Telemetry / sFlow is in use, this will not work with mgmt vlan; in this case consider using mgmt clip or mgmt oob



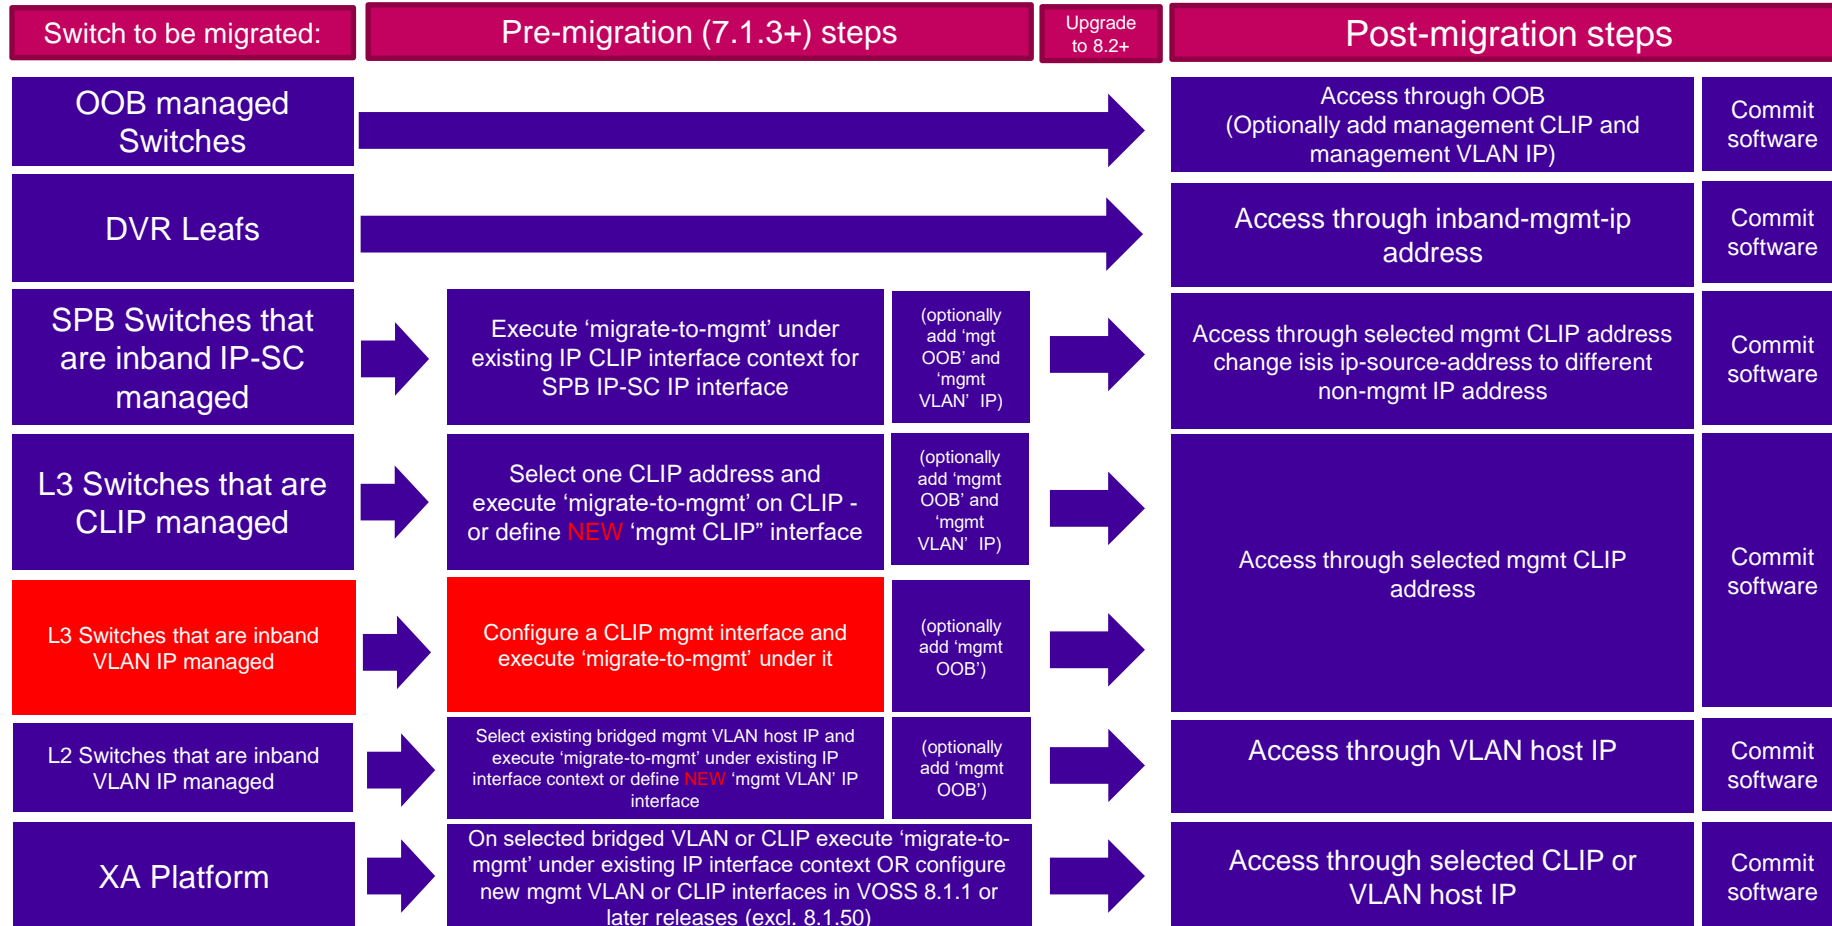
# Migration of DVR Leaf



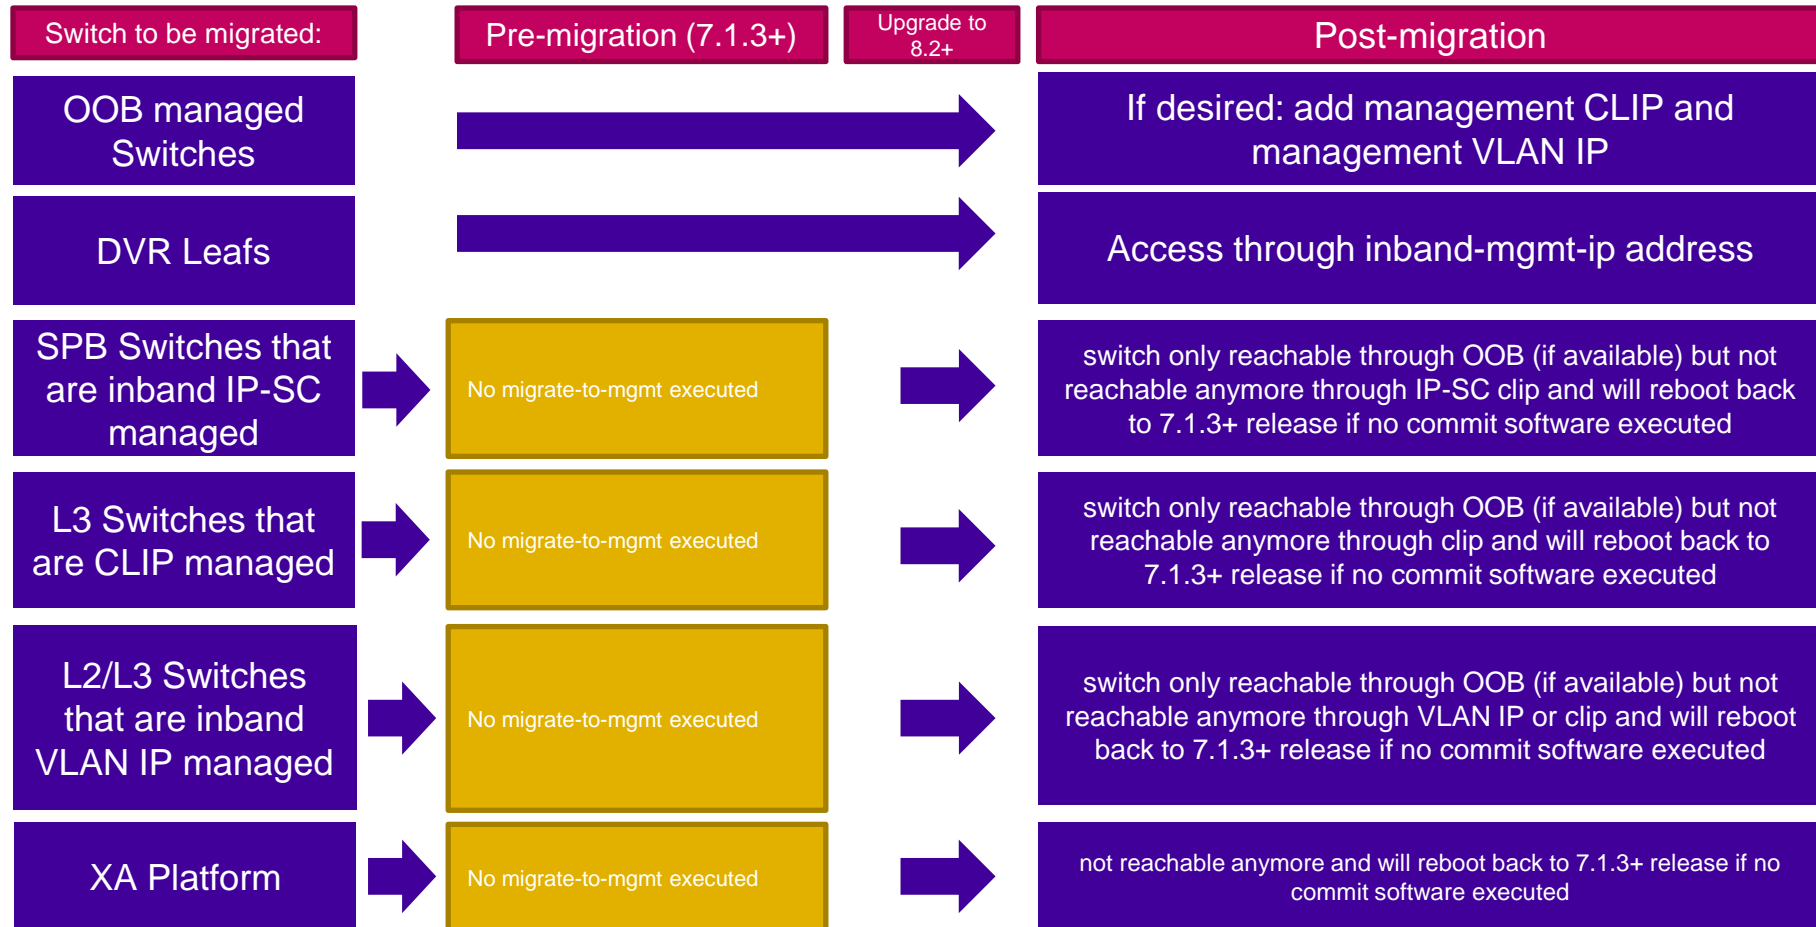
- simply upgrade

- The DVR inband-mgmt-ip CLIP automatically becomes the new segmented mgmt clip
- The ISIS inband-mgmt-ip command becomes obsolete in 8.2

# Upgrade paths to VOSS 8.2+

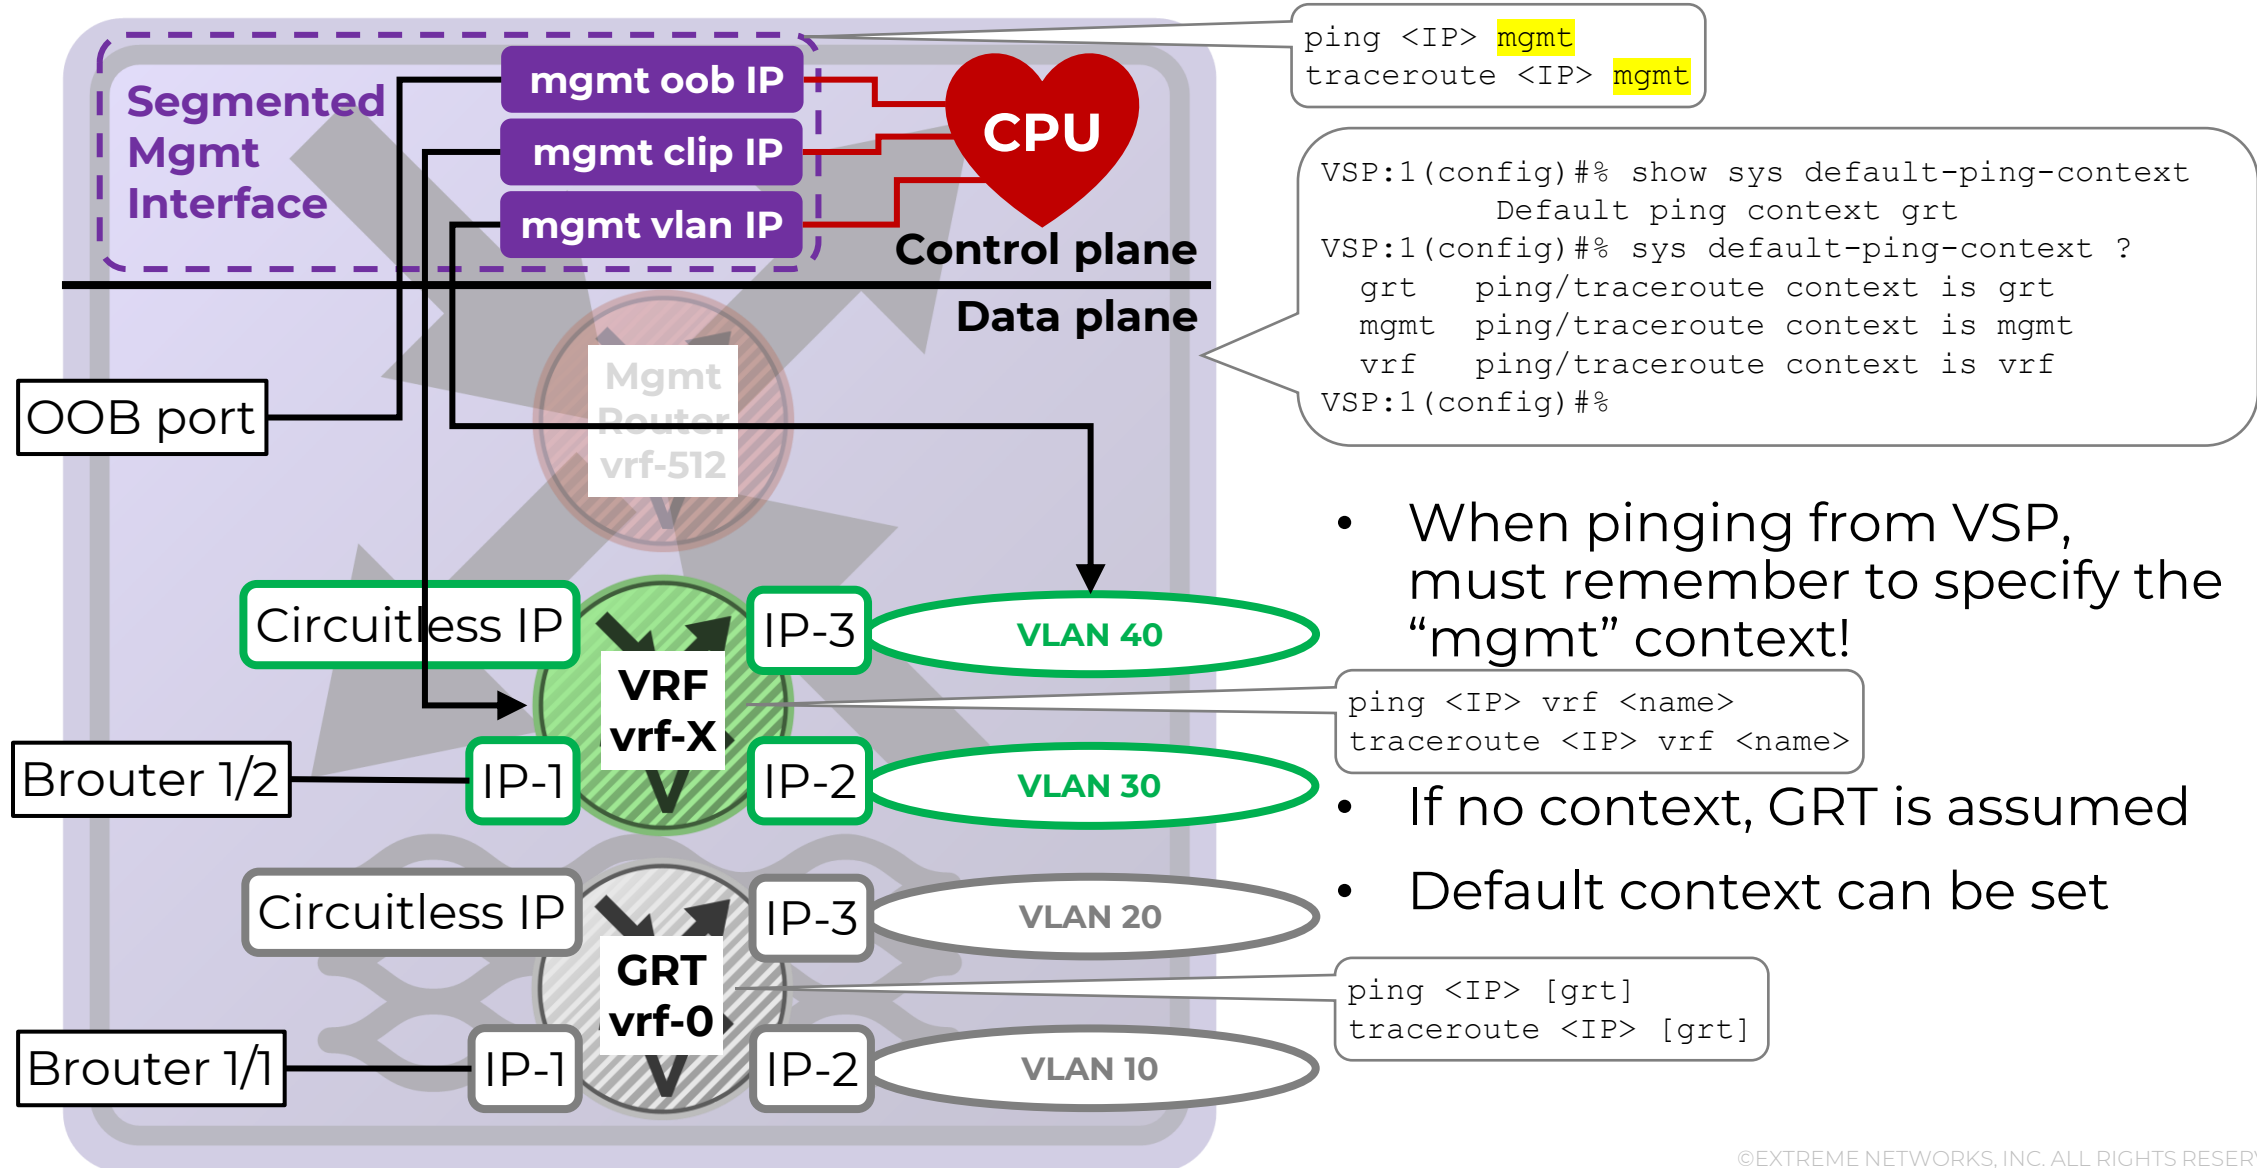


# VOSS 8.2+ upgrade – what if?



# Ping/Traceroute changes with 8.2

# VOSS IP mgmt 8.2 with Segmented Mgmt Interface



```
ping <IP> mgmt
traceroute <IP> mgmt
```

```
VSP:1(config)#% show sys default-ping-context
Default ping context grt
VSP:1(config)#% sys default-ping-context ?
grt ping/traceroute context is grt
mgmt ping/traceroute context is mgmt
vrf ping/traceroute context is vrf
VSP:1(config)#%
```

- When pinging from VSP, must remember to specify the "mgmt" context!

```
ping <IP> vrf <name>
traceroute <IP> vrf <name>
```

- If no context, GRT is assumed
- Default context can be set

```
ping <IP> [grt]
traceroute <IP> [grt]
```

