

CXS2110BCN

vmware® EXPLORE

Accelerate Your Networking and Security Functions with SmartNICs

Nand Kishore Vyas
Technical Lead / Manager (NSX)

Jayaram K G
Staff Network Virtualization Engineer

#vmwareexplore #CXS2110BCN



Presenters



Jayaram K G

Staff Network Virtualization Engineer



Nand Kishore Vyas

Technical Lead / Manager (NSX)

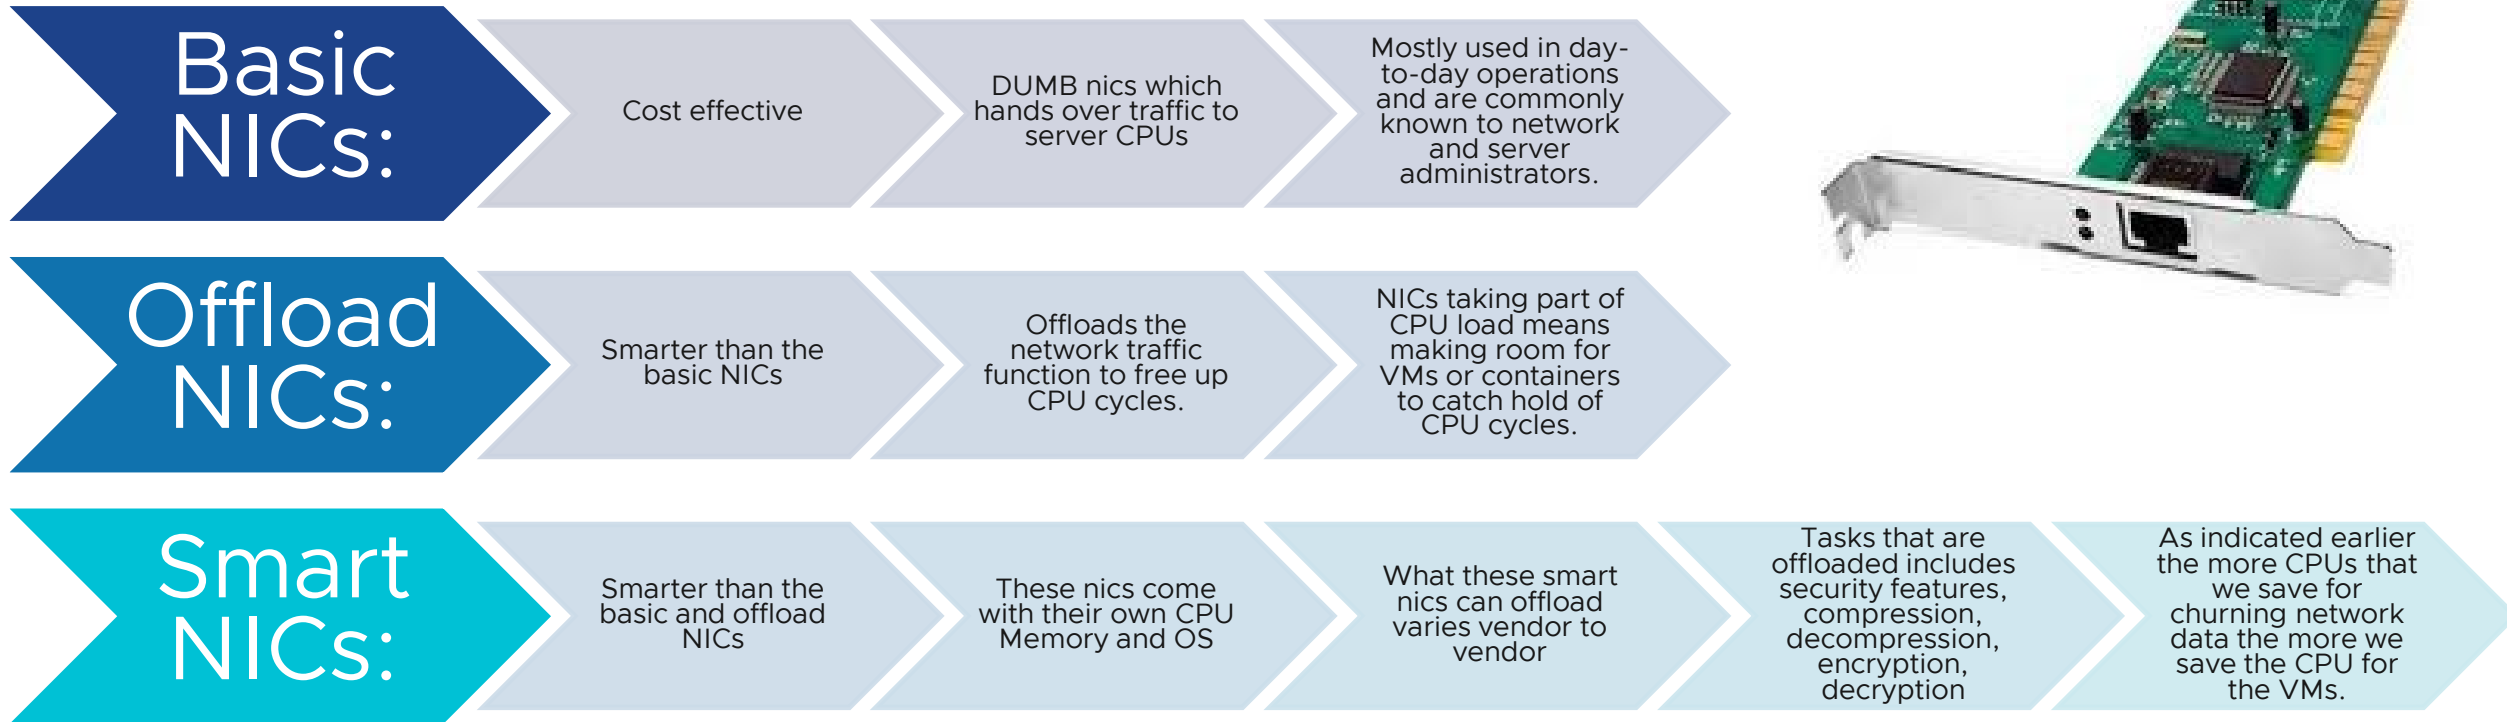


Agenda

1. Evolution of NICs
2. What is DPU or SmartNics ?
3. Unfolding the Datapath
 - Standard Datapath
 - SRIOV & Enhanced Datapath
 - DPUs - Modes
4. Packet walk over DPU
5. Summaries UPTv2 vs Emulated mode

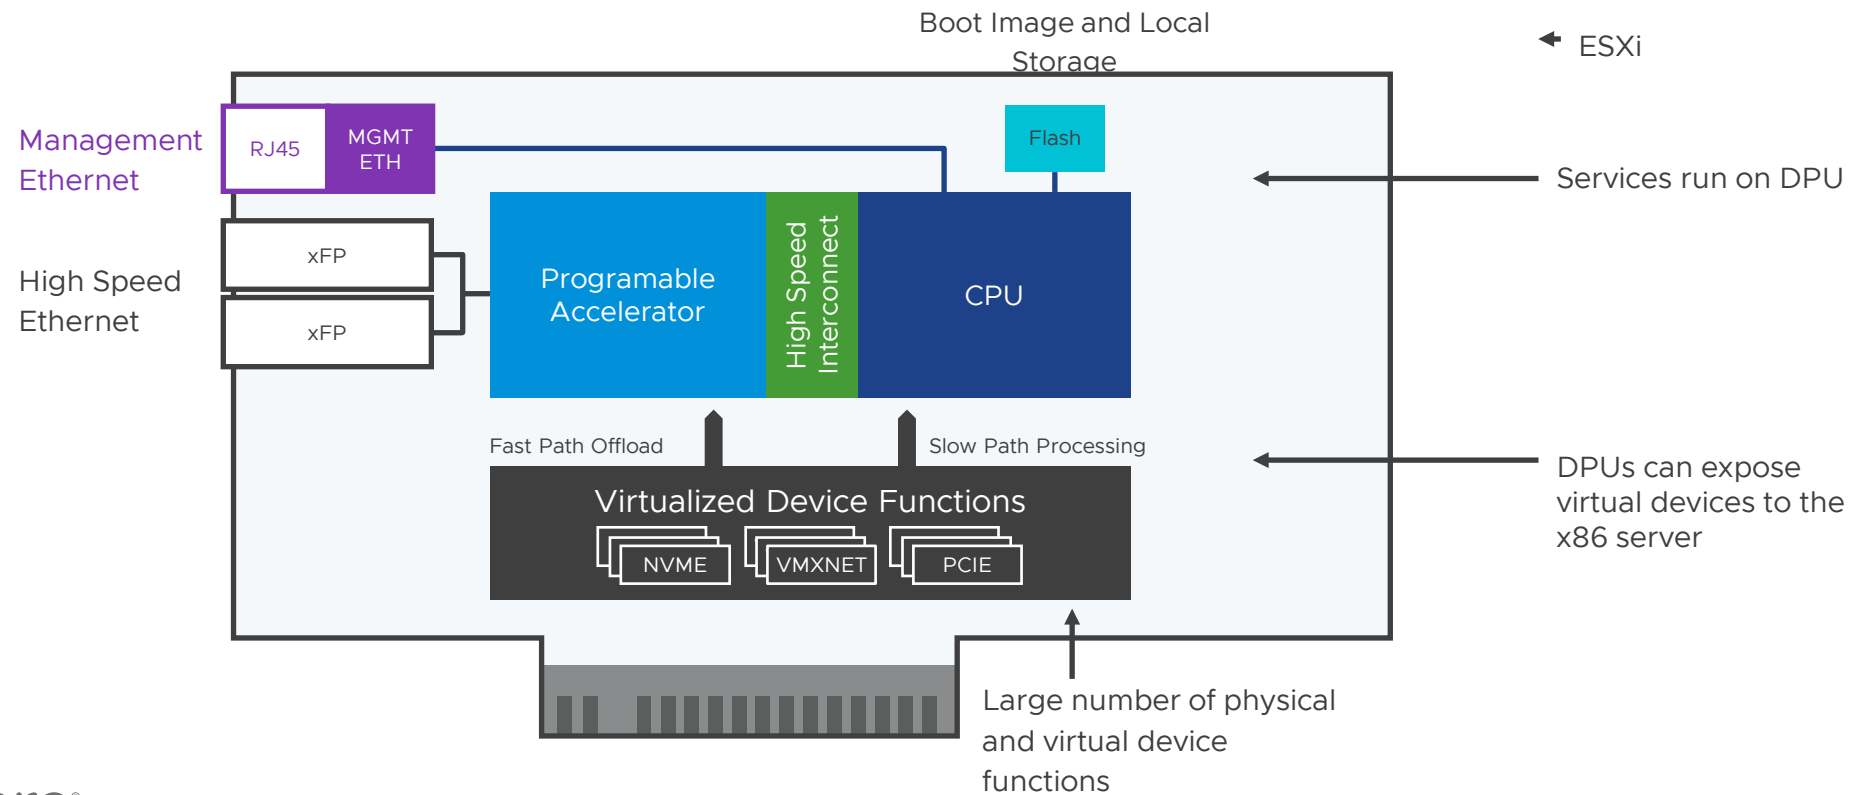
Evolution of NICs

Evolution of NICs



What is DPU or SmartNics ?

- DPU/SmartNics are a new class of programmable processor that combines with multi core ARM CPU, High performance network interface and rich programmable acceleration engines
- DPU / SmartNic allows to install ESXi ARM on the flash drive of the Smart Nic
- Isolated Compute and programmable accelerator to handle task offloaded by x86 ESXi
- Virtualized Device Functions to expose virtual devices to the x86 server.

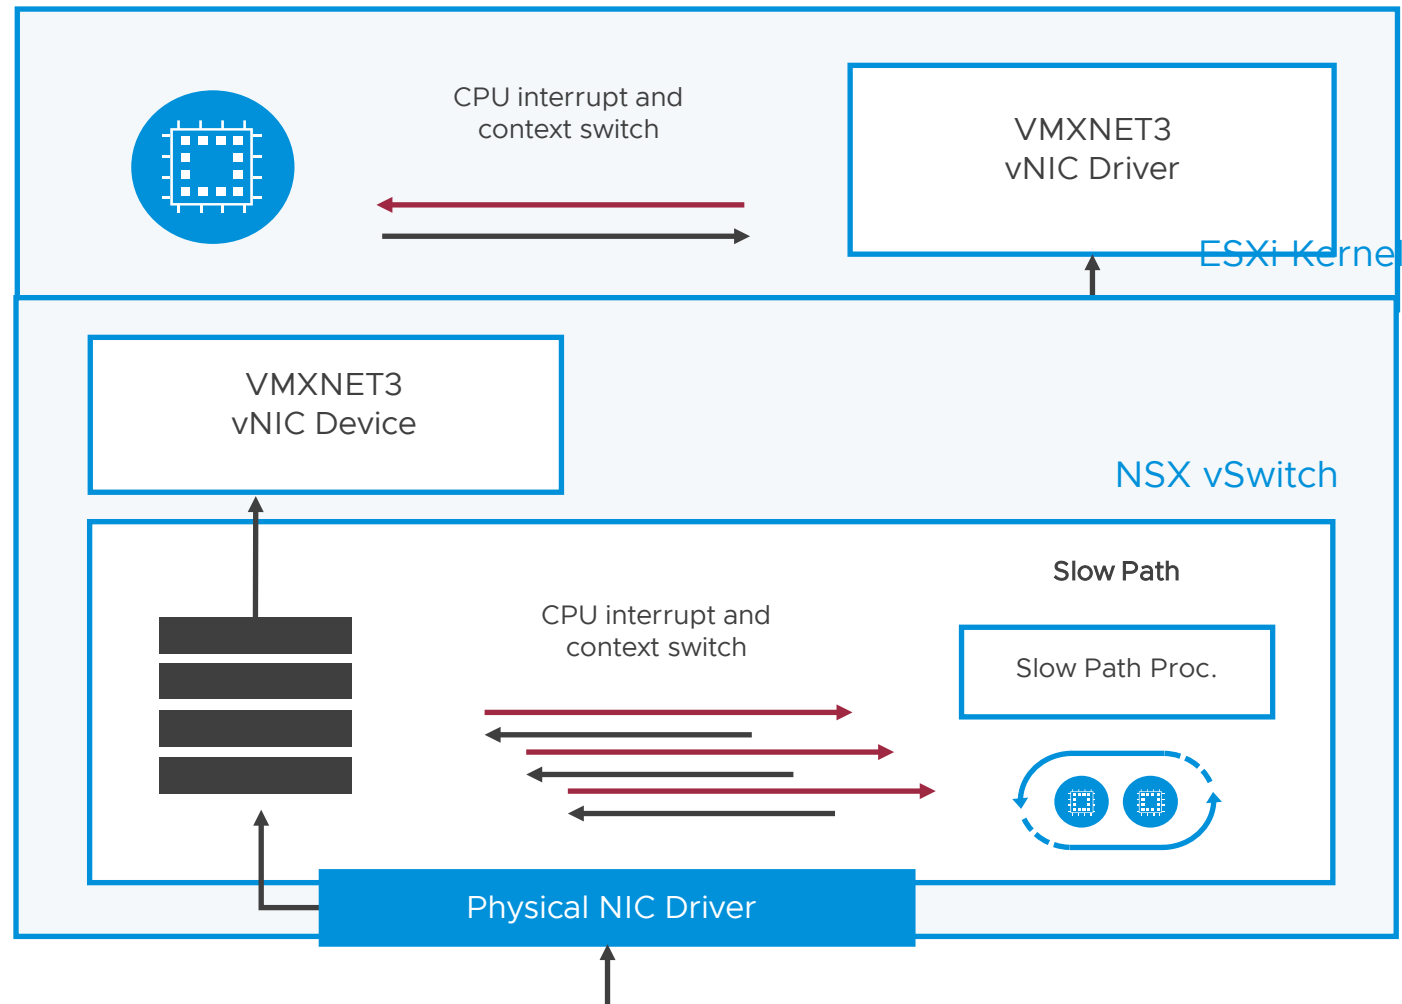


Unfolding the Datapath

Standard Datapath

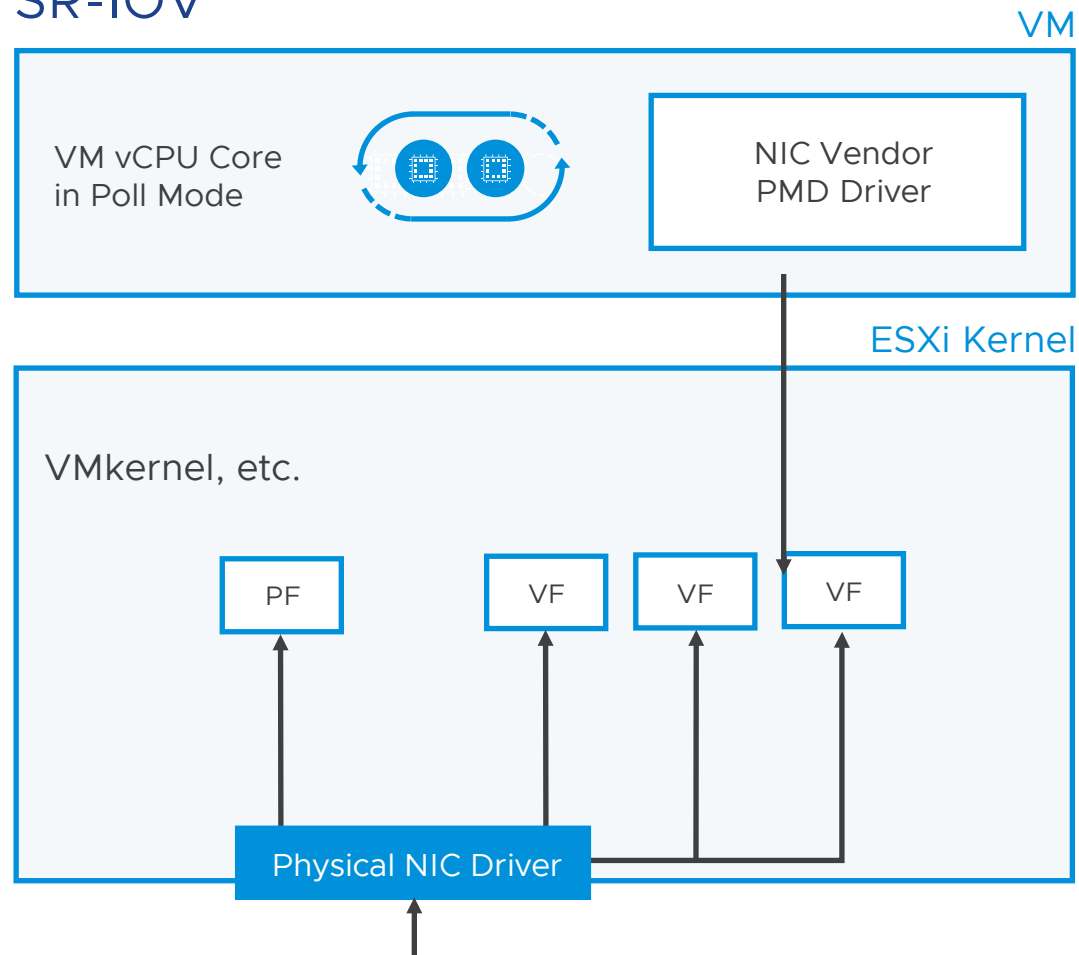
Standard VDS

VM

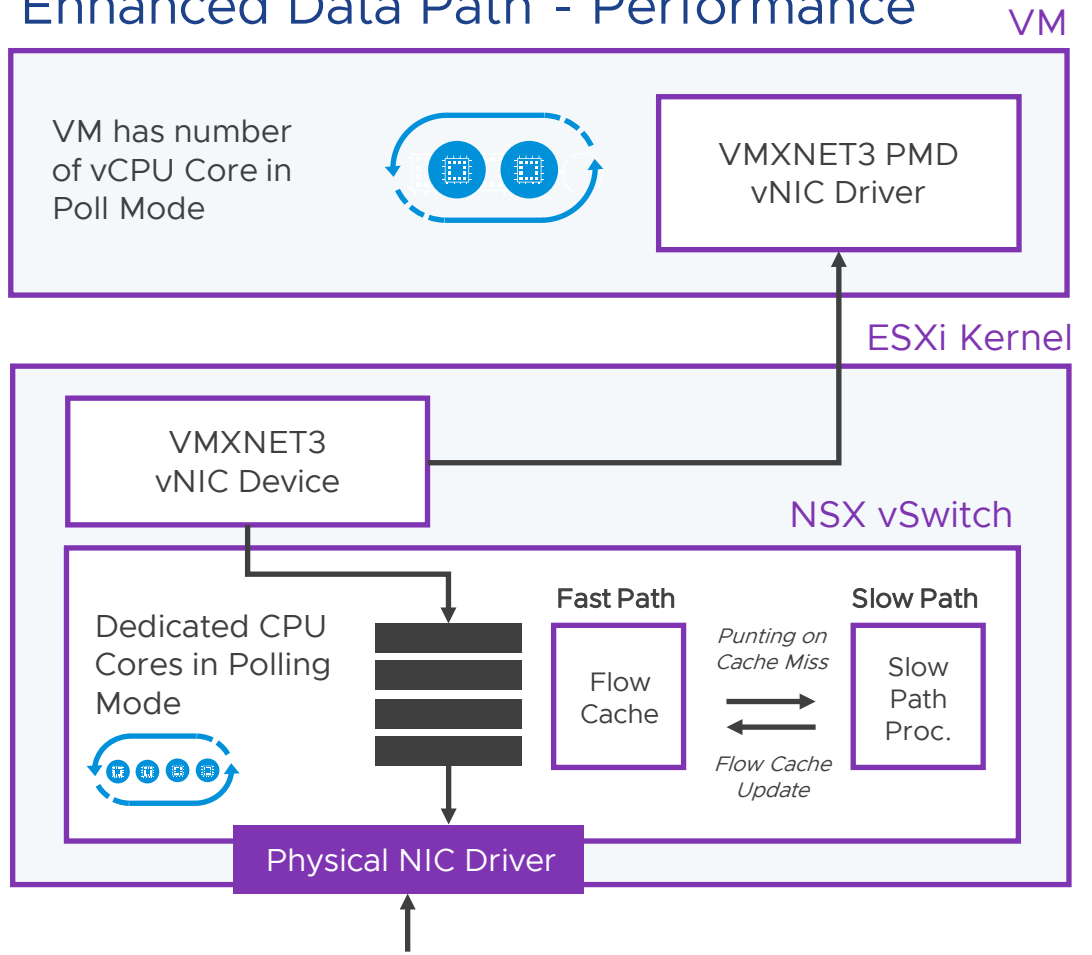


SR-IOV and Enhanced Datapath

SR-IOV



Enhanced Data Path - Performance



EDP-Polling vs. SRIOV

Why do we need more than what we support today?

SR-IOV / Passthrough

- vSwitch is completely bypassed
- No additional CPU load and forwarding latency because no Hypervisor vSwitch is used
- Introduces NIC Driver dependency in VM / Application
- VMs are “blocked” from moving, all vSphere high-availability and Hypervisor operational features are lost



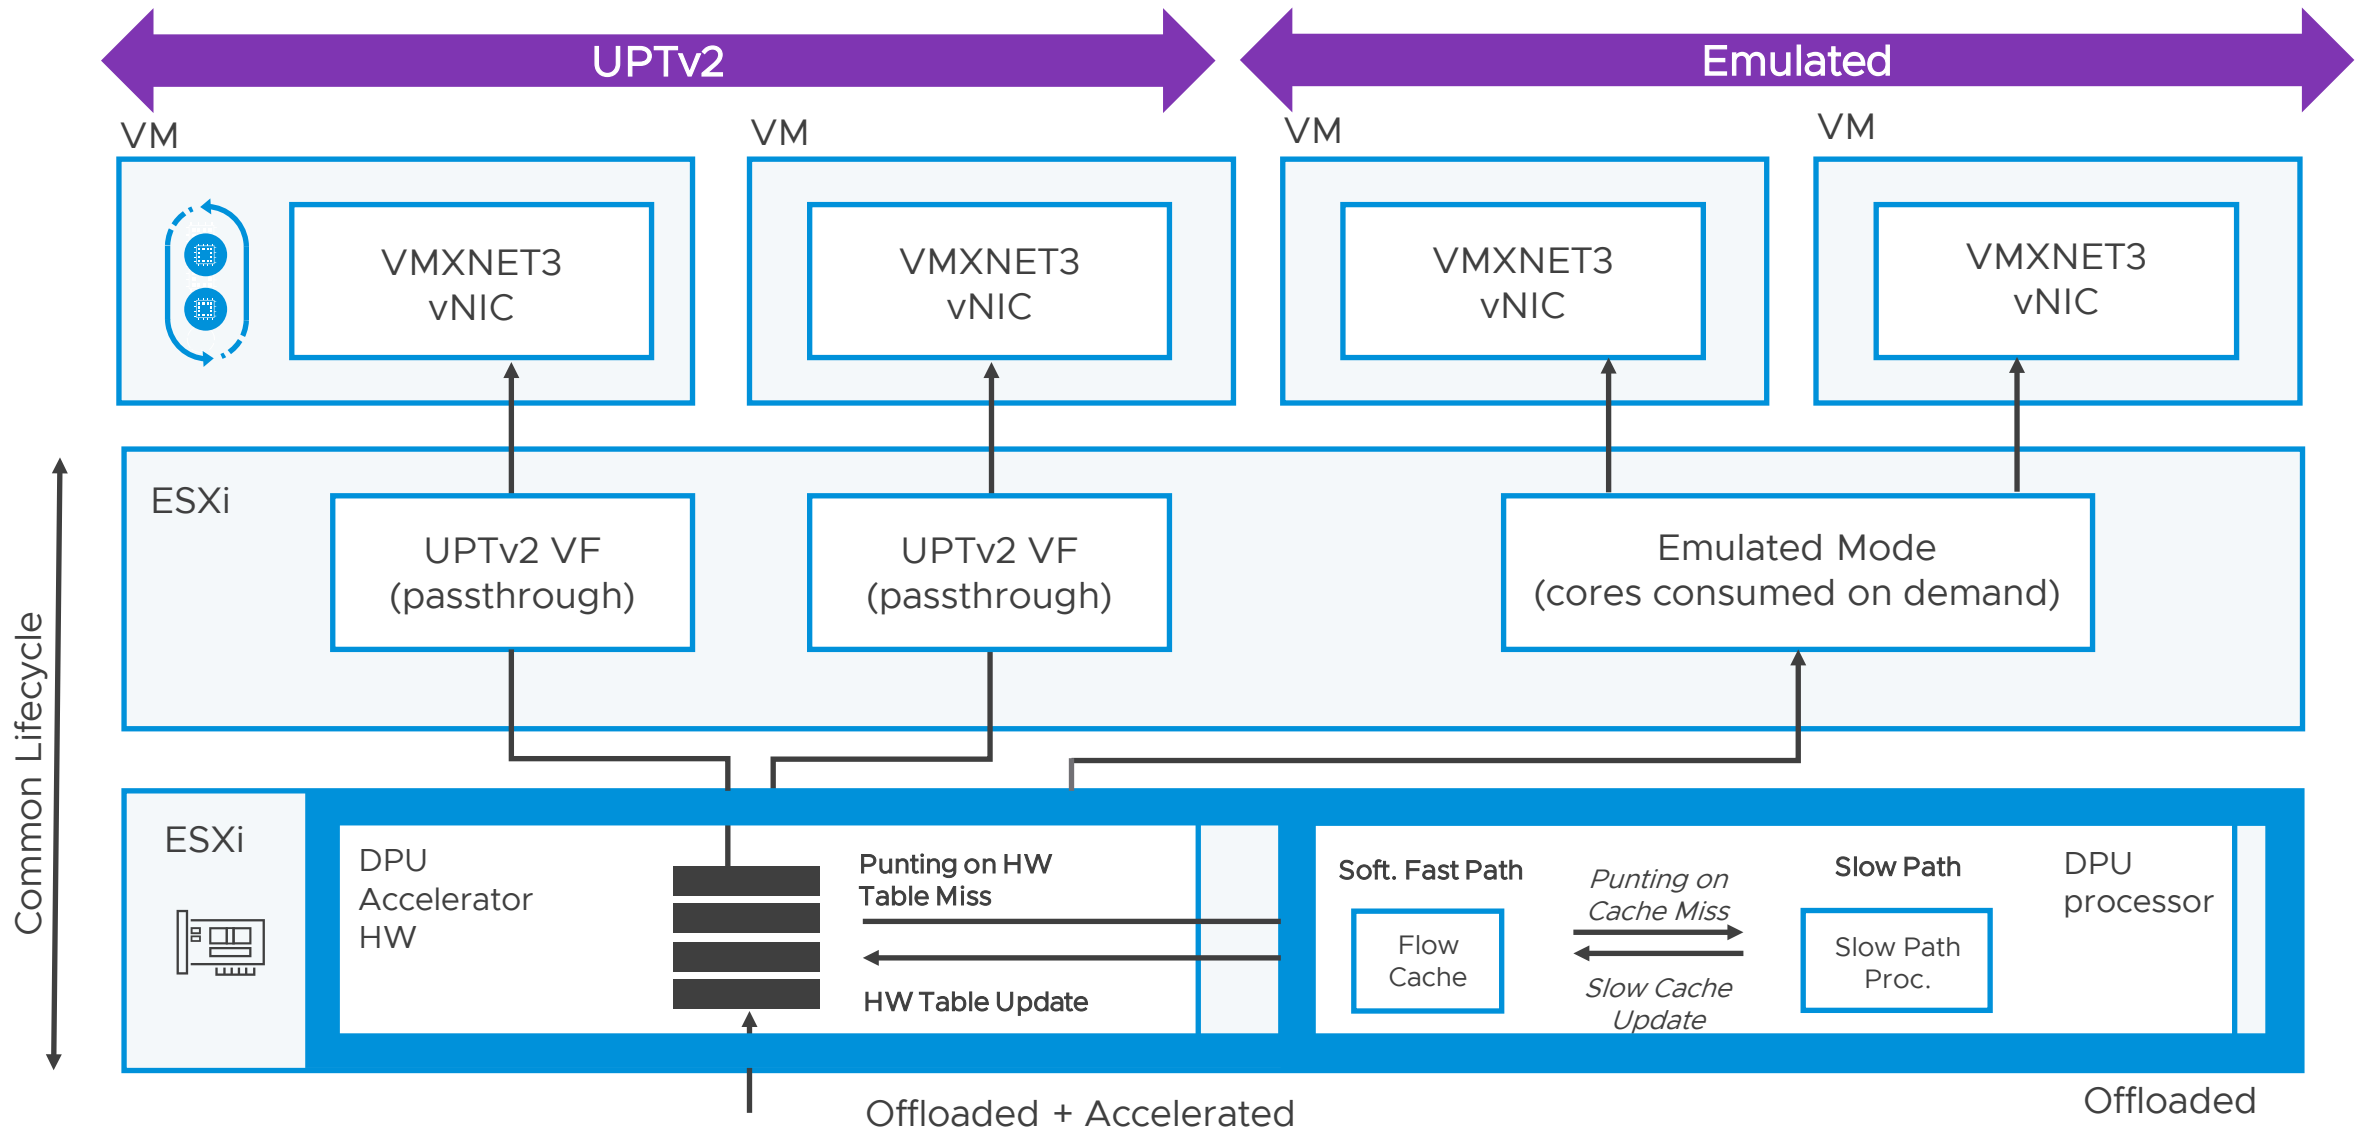
Enhanced Data Path (EDP)

- Single VM Driver (VMXNET3) abstracts away individual NIC Vendors from VM Driver
- Single driver model reduces complex Multi-party issue resolution
- All vSphere Hypervisor high-availability and operational features are usable (vMotion, DRS, HA, Teaming, etc.)
- Decoupling VMs from hardware enables virtualization benefits



DPU-based acceleration combines the benefits of SR-IOV and EDP

DPU - Emulated and UPTv2 Mode



Packet walk Over DPU/SmartNic Emulated Mode

Enable Emulated Mode or Disable UPTv2 Mode

Edit Settings | gs-ubuntu-01

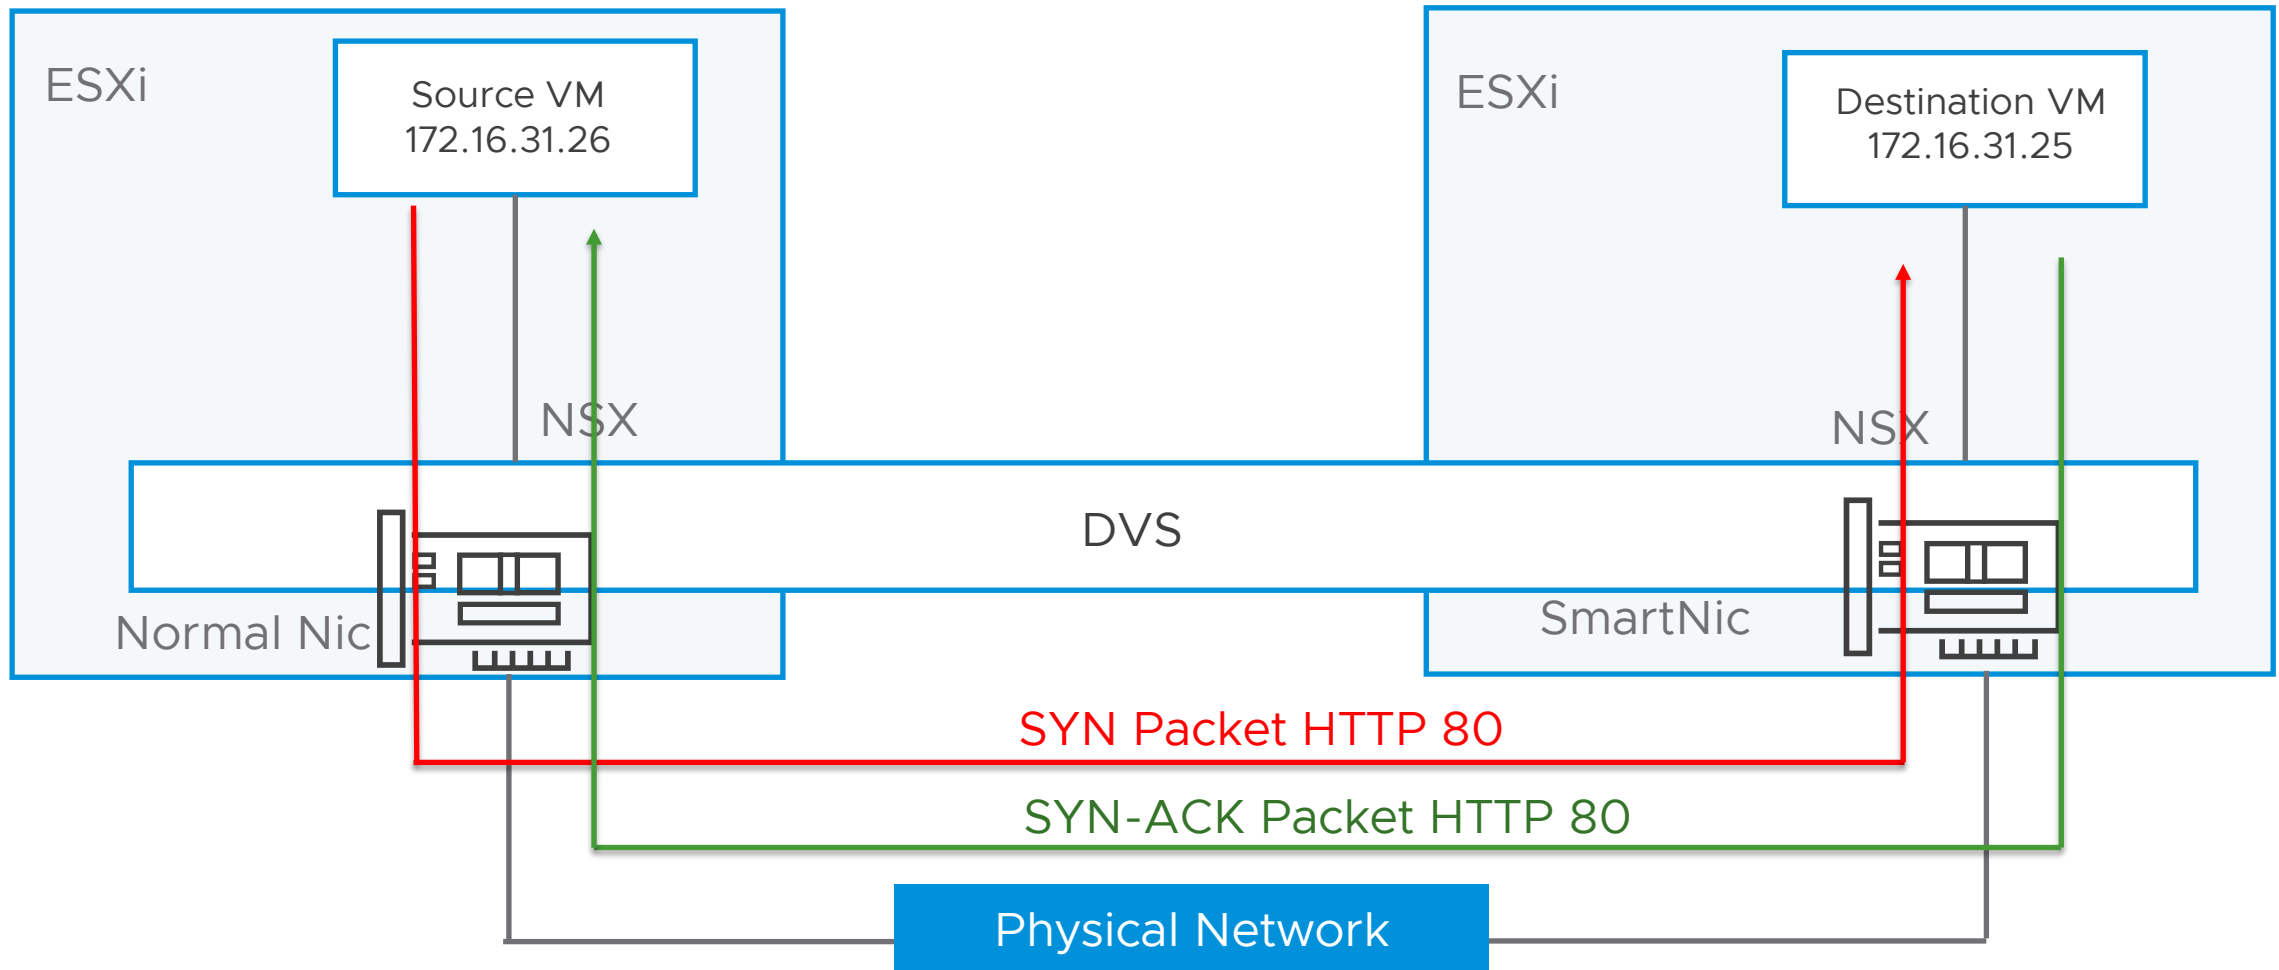
Virtual Hardware | VM Options | Advanced Parameters

ADD NEW DEVICE ▾

> CPU	2	ⓘ
> Memory *	2	GB ▾
> Hard disk 1	16	GB ▾
> SCSI controller 0	LSI Logic Parallel	⋮
▾ Network adapter 1 *	cluster1-segment1	Connected
Status	<input checked="" type="checkbox"/> Connect At Power On	
Port ID	c380b93d-ee7f-4eb5-92	
Adapter Type	VMXNET 3	▾
UPT Support	<input type="checkbox"/> Use UPT Support	
MAC Address	00:50:56:a1:99:9e	Automatic ▾
> CD/DVD drive 1	Client Device	<input type="checkbox"/> Connect At Power On
> Video card	Specify custom settings	▾
> SATA controller 0	AHCI	⋮
> Security Devices	Not Configured	
> Other	Additional Hardware	

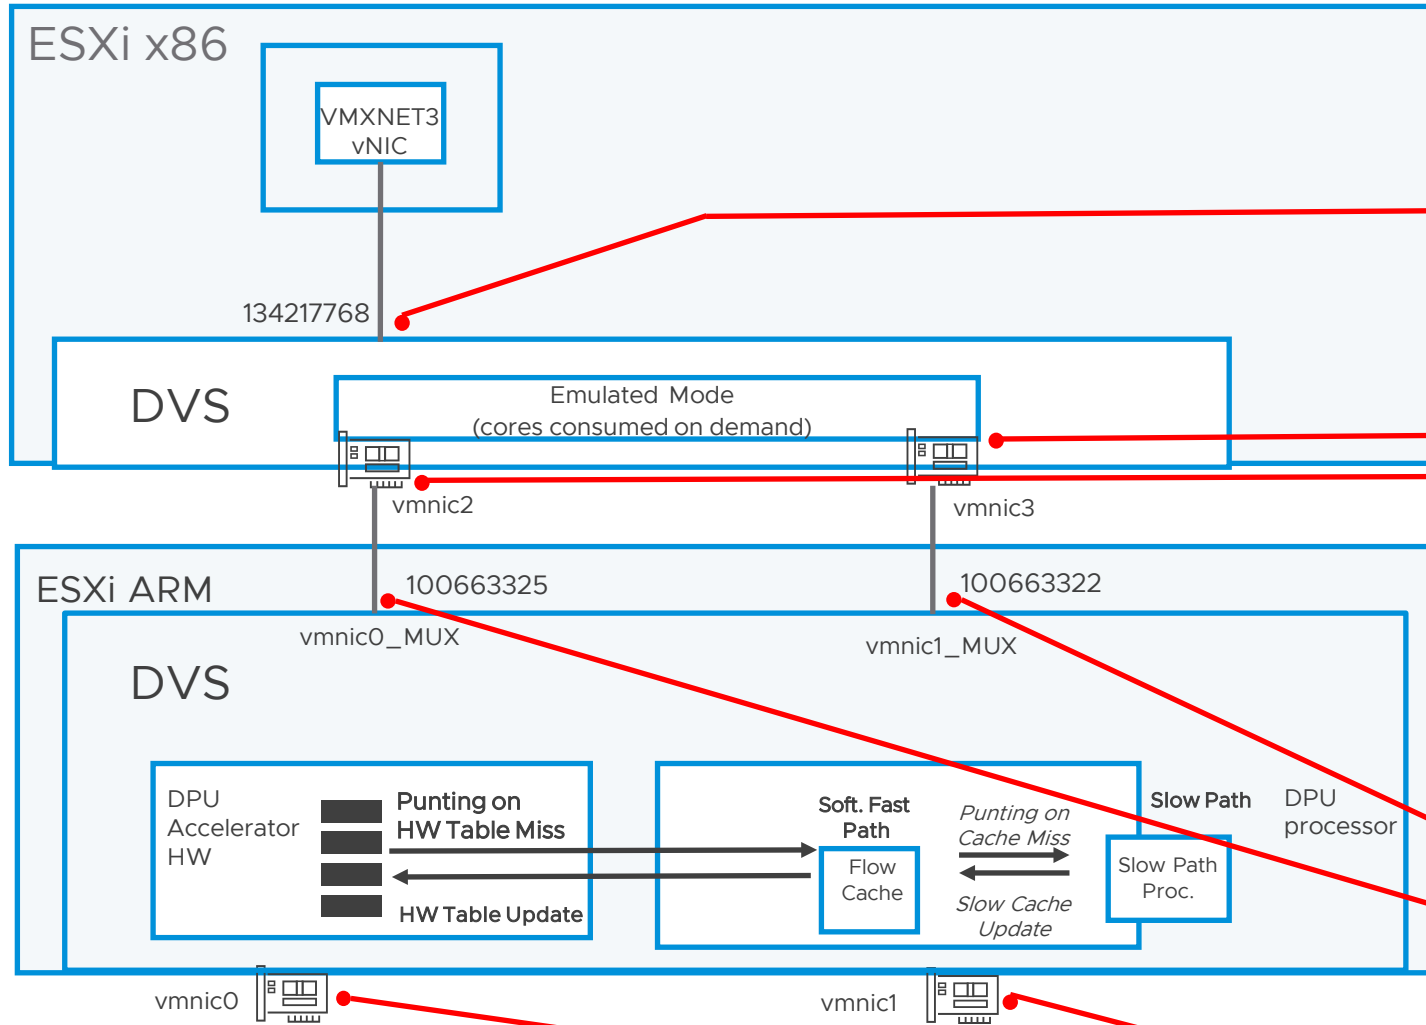
CANCEL OK

Datapath Emulated Mode



Mapping of ports in Emulated Mode

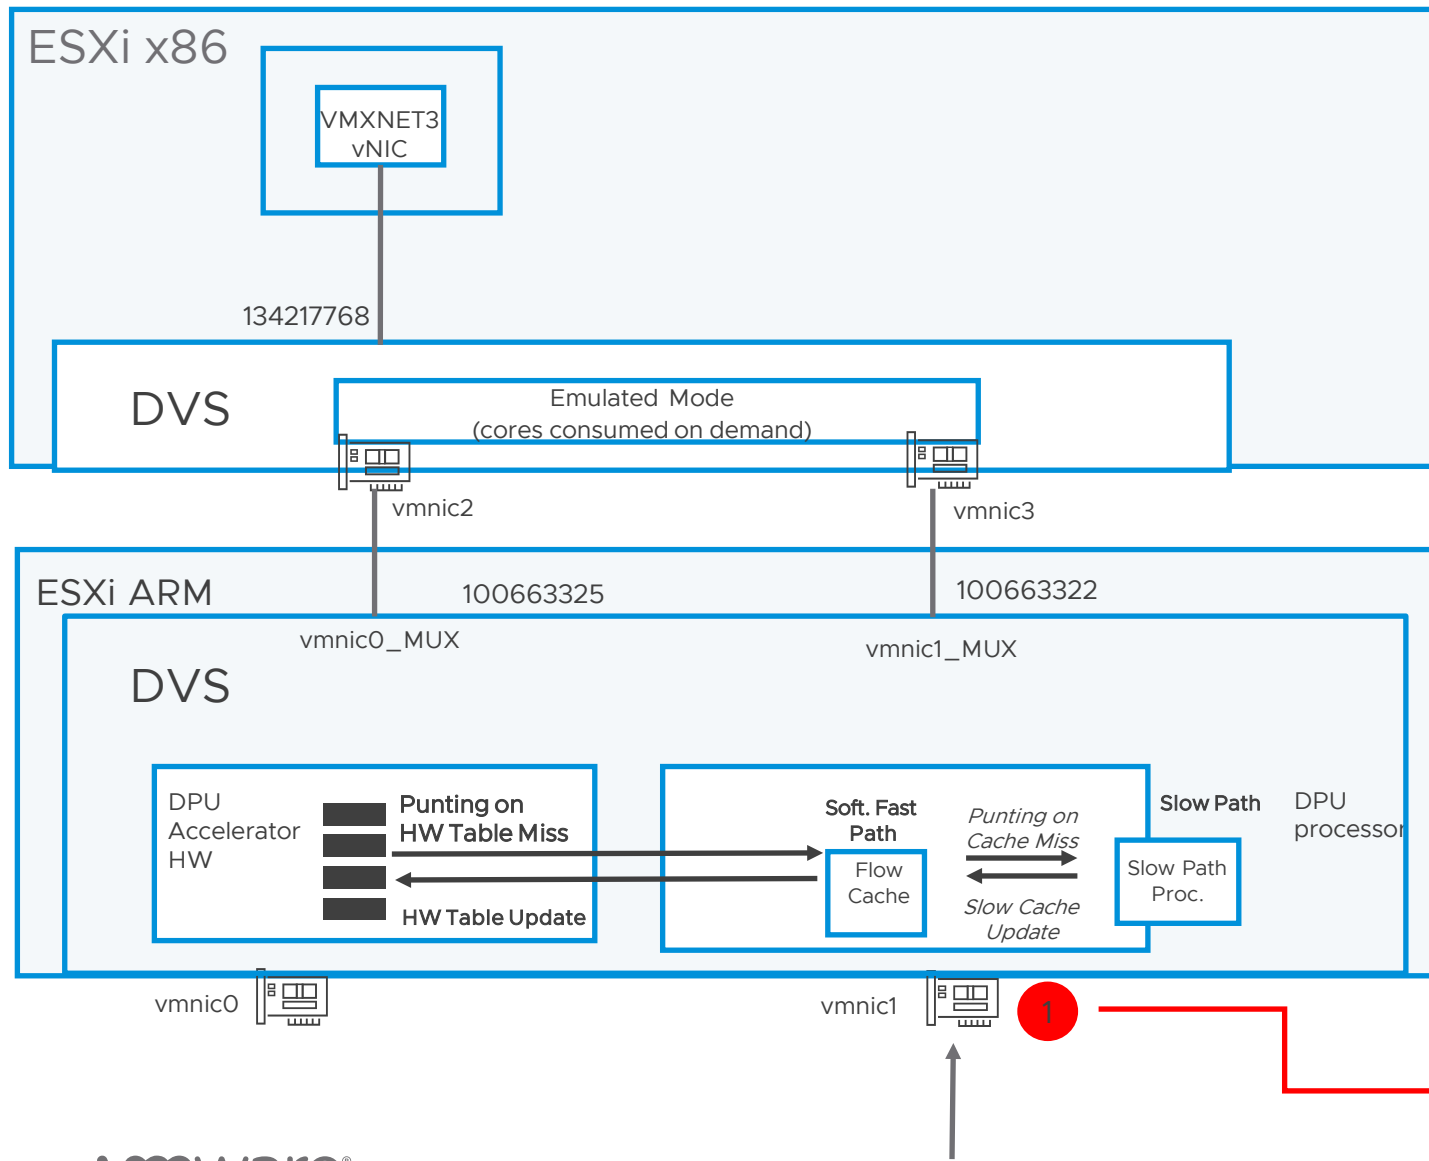
Follow the RED Dot



ESXi	PORT-ID USED-BY	TEAM-PNIC DNAME
x86	67108872 Management	n/a esxioVswitch0
	100663307 Management	n/a DvsPortset-0
	100663317 Shadow of vmnic0	n/a DvsPortset-0
	100663319 Shadow of vmnic1	n/a DvsPortset-0
	100663320 vmk0	vmnic0 DvsPortset-0
	100663322 vmk2	vmnic1 DvsPortset-0
	100663323 vmk3	vmnic1 DvsPortset-0
	134217740 Management	n/a DvsPortset-1
	134217745 Shadow of vmnic3	n/a DvsPortset-1
	134217747 Shadow of vmnic2	n/a DvsPortset-1
	134217768 4241327:gs-ubuntu-01.eth0	vmnic2 DvsPortset-1
	167772189 Management	n/a vSwitchBMC
	167772191 Shadow of vusb0	n/a vSwitchBMC
	167772192 vmk1	vusb0 vSwitchBMC
	2214592521 vmnic128	- esxioVswitch0
	2248146964 vmnic0	- DvsPortset-0
	2248146966 vmnic1	- DvsPortset-0
2281701392 vmnic3	- DvsPortset-1	
2281701394 vmnic2	- DvsPortset-1	
ARM	67108871 Management	n/a
	67108873 vmk7379	vmnic129 esxioesxioVswitch0Vswitch0
	100663306 Management	n/a DvsPortset-1
	100663310 vmk10	vmnic1 DvsPortset-1
	100663311 vmk11	vmnic0 DvsPortset-1
	100663313 vmk50	void DvsPortset-1
	100663320 Shadow of vmnic1	n/a DvsPortset-1
	100663321 vdr-vdrPort	vmnic1 DvsPortset-1
	100663322 vmnic1_MUX_PF20483	vmnic1 DvsPortset-1
	100663324 Shadow of vmnic0	n/a DvsPortset-1
	100663325 vmnic0_MUX_PF16387	vmnic0 DvsPortset-1
	100663333 gs-ubuntu-01.eth0.7c520dd4	vmnic1 DvsPortset-1
	134217747 Management	n/a vSwitchESXioBMC
	134217749 Shadow of vmnic2	n/a vSwitchESXioBMC
	134217750 vmk255	vmnic2 vSwitchESXioBMC
	2214592520 vmnic129	- esxioVswitch0
	2248146967 vmnic1	- DvsPortset-1
2248146971 vmnic0	- DvsPortset-1	
2281701396 vmnic2	- vSwitchESXioBMC	

Packet Walk in Emulated Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.25 HTTP

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
2248146967   4    0 DvsPortset-1
08:c0:eb:5d:ae:0f vmnic1
```

Command:

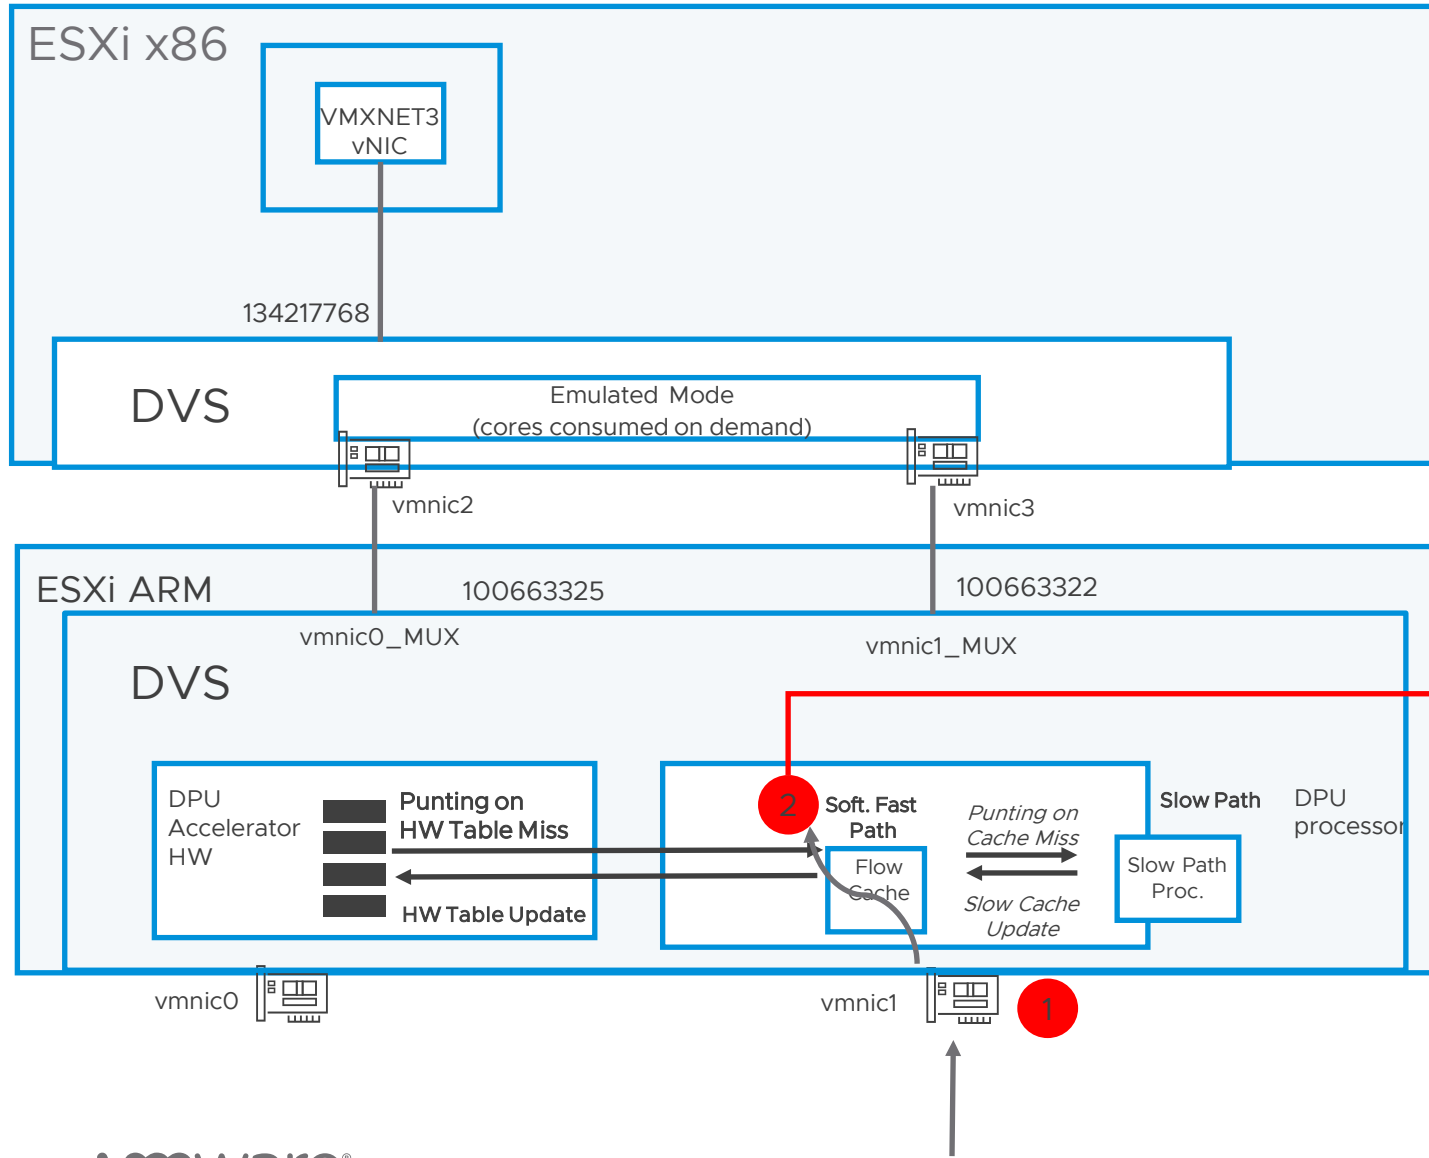
```
pktcap-uw --uplink vmnic1 --dir 2 -o - | tcpdump-uw -enr -
```

Packet:

```
00:50:56:63:04:5c > 00:50:56:63:04:5c, ethertype IPv4
(0x0800), length 132: 192.168.4.57.49190 >
192.168.4.57.8080: Geneve, Flags [none], vni 0x12400, proto
TEB (0x6558), options [8 bytes]: 00:50:56:a1:fc:f3 >
00:50:56:63:04:5c > 00:50:56:63:04:5c, ethertype IPv4 (0x0800), length 74:
172.16.31.26.57448 > 172.16.31.25: Flags [S], seq
2198858042, win 64240, options [mss 1460,sackOK,TS val
3020130186 ecr 0,nop,wscale 7], length 0
```


Packet Walk in Emulated Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.25 HTTP

Command:

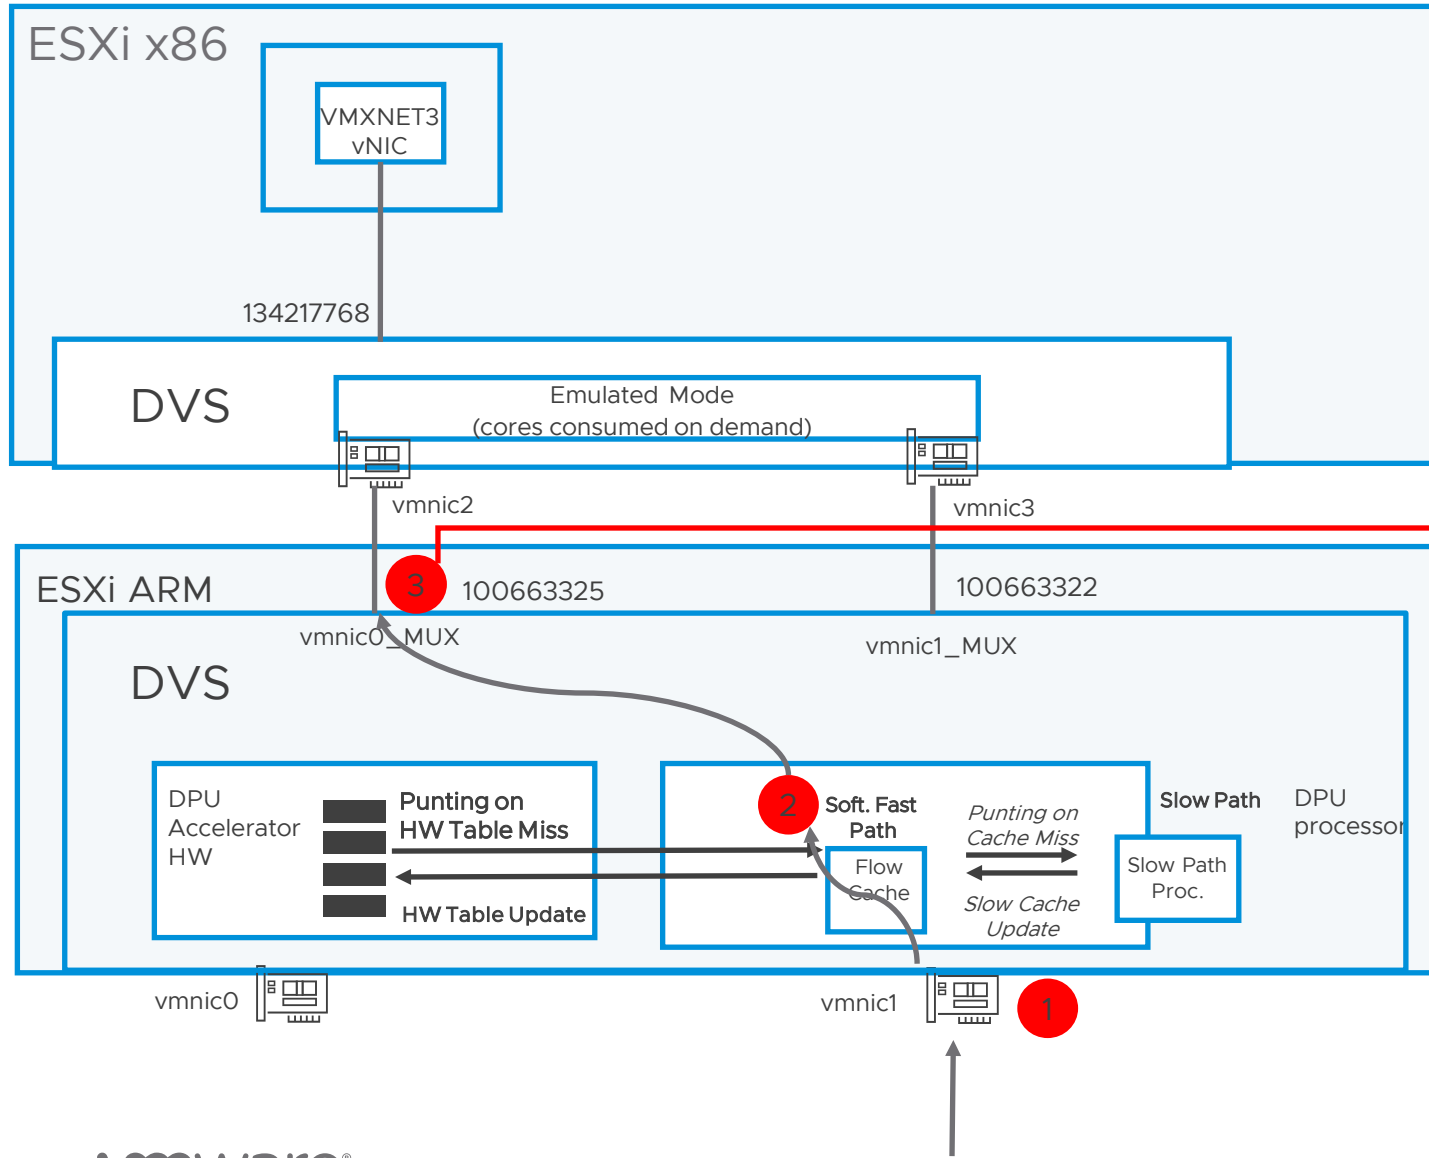
```
nsxcli ens flow-table dump
```

Flow Cache:

FT	dstMAC	srcMAC	VLAN	srcPort	dstIP	proto	VNI	srcPort/type	dstPort/code	Actions	hwHits
L4	00:50:56:a1:fc:f3	00:50:56:a1:fc:f3	0	3	172.16.31.26	80	74752	57448	0	bmap:0x80800080	6
inval(s):150 cg:2123 dp:0x7 len:704; GENEVE DECAP; DFW on dstPort; 0 0											
L4	00:50:56:a1:99:9e	00:50:56:a1:99:9e	0	7	172.16.31.25	80	80	0	0	bmap:0x20000c0	0
cg:2123 dp:0x3 len:814; DFW on srcPort; VNI: 74752; GENEVE ENCAP VNI: 74752; 0 0											

Packet Walk in Emulated Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.25 HTTP

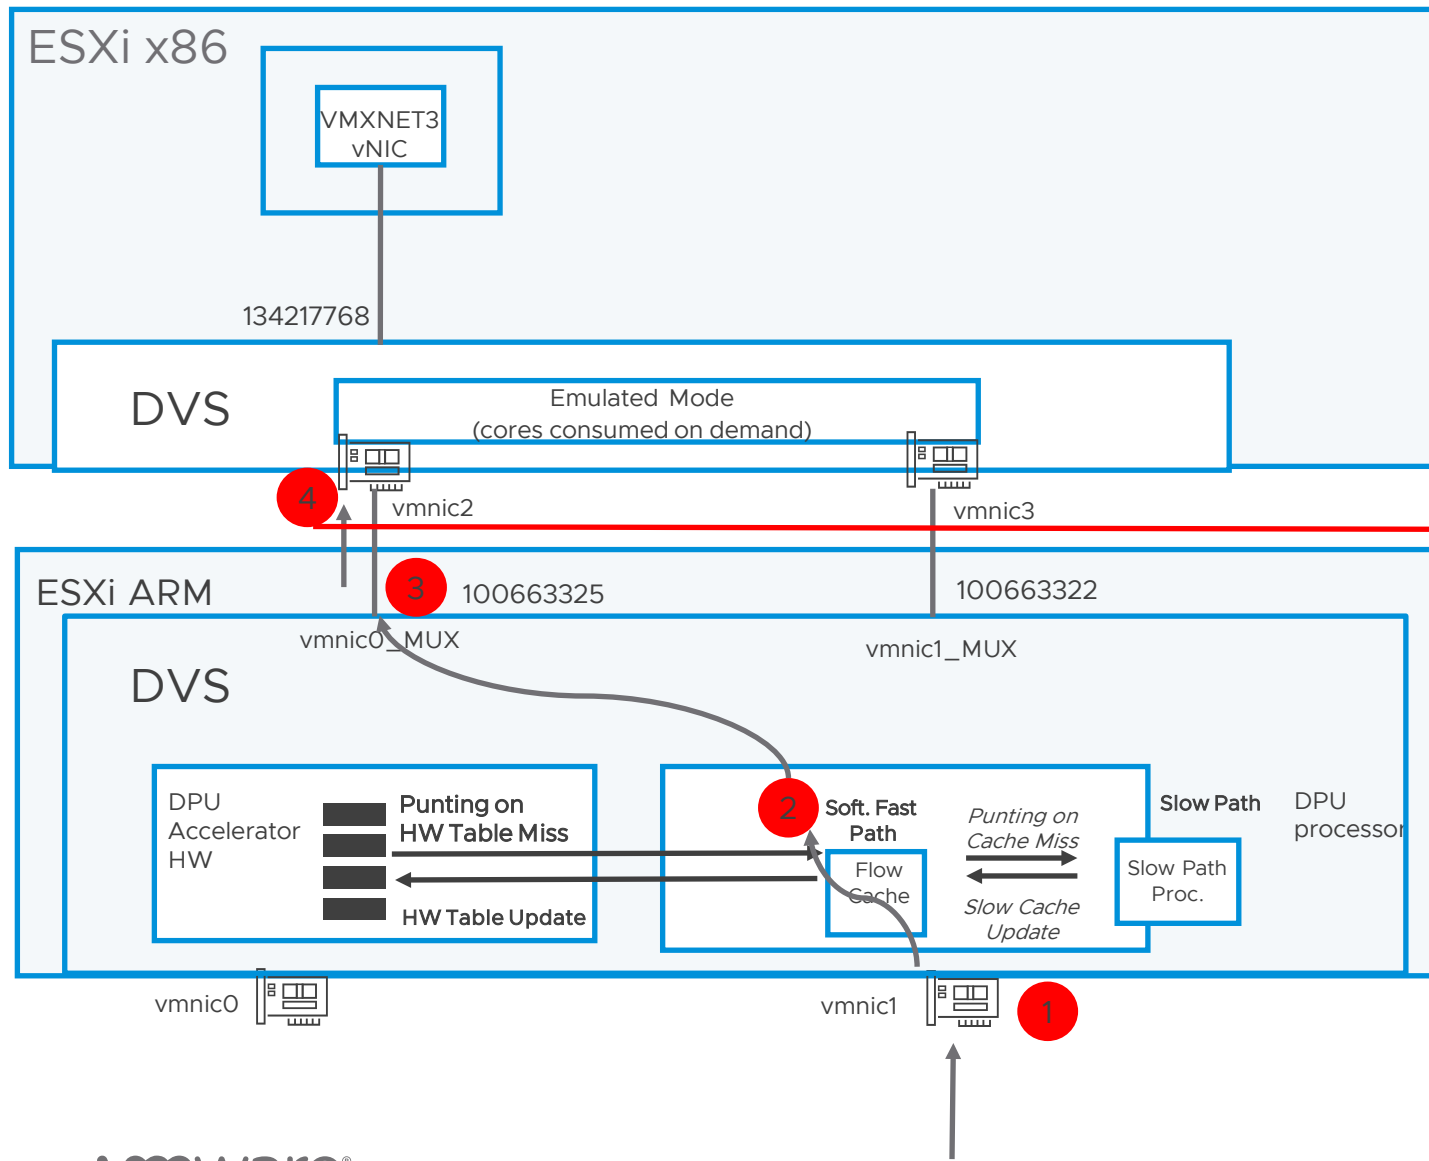
```
Port Info:
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
100663325    13      0 DvsPortset-1
08:c0:eb:5d:ae:08  vmnic0_MUX_PF16387

Command:
pktcap-uw -switchport 100663325 --dir 2 -o - | tcpdump-uw -
enr -

Packet:
00:50:56:a1:fc:f3 > 00:50:56:a1:33:24, ethertype IPv4
(0x0800), length 74: 172.16.31.26.57448 > 172.16.31.25.31:
Flags [S], seq 2198858042, win 64240, options [mss
1460,sackOK,TS val 3020130186 ecr 0,nop,wscale 7], length 0
```

Packet Walk in Emulated Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.25 HTTP

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
2281701394   4      0 DvsPortset-1
08:c0:eb:5d:ad:78  vmnic2
```

Command:

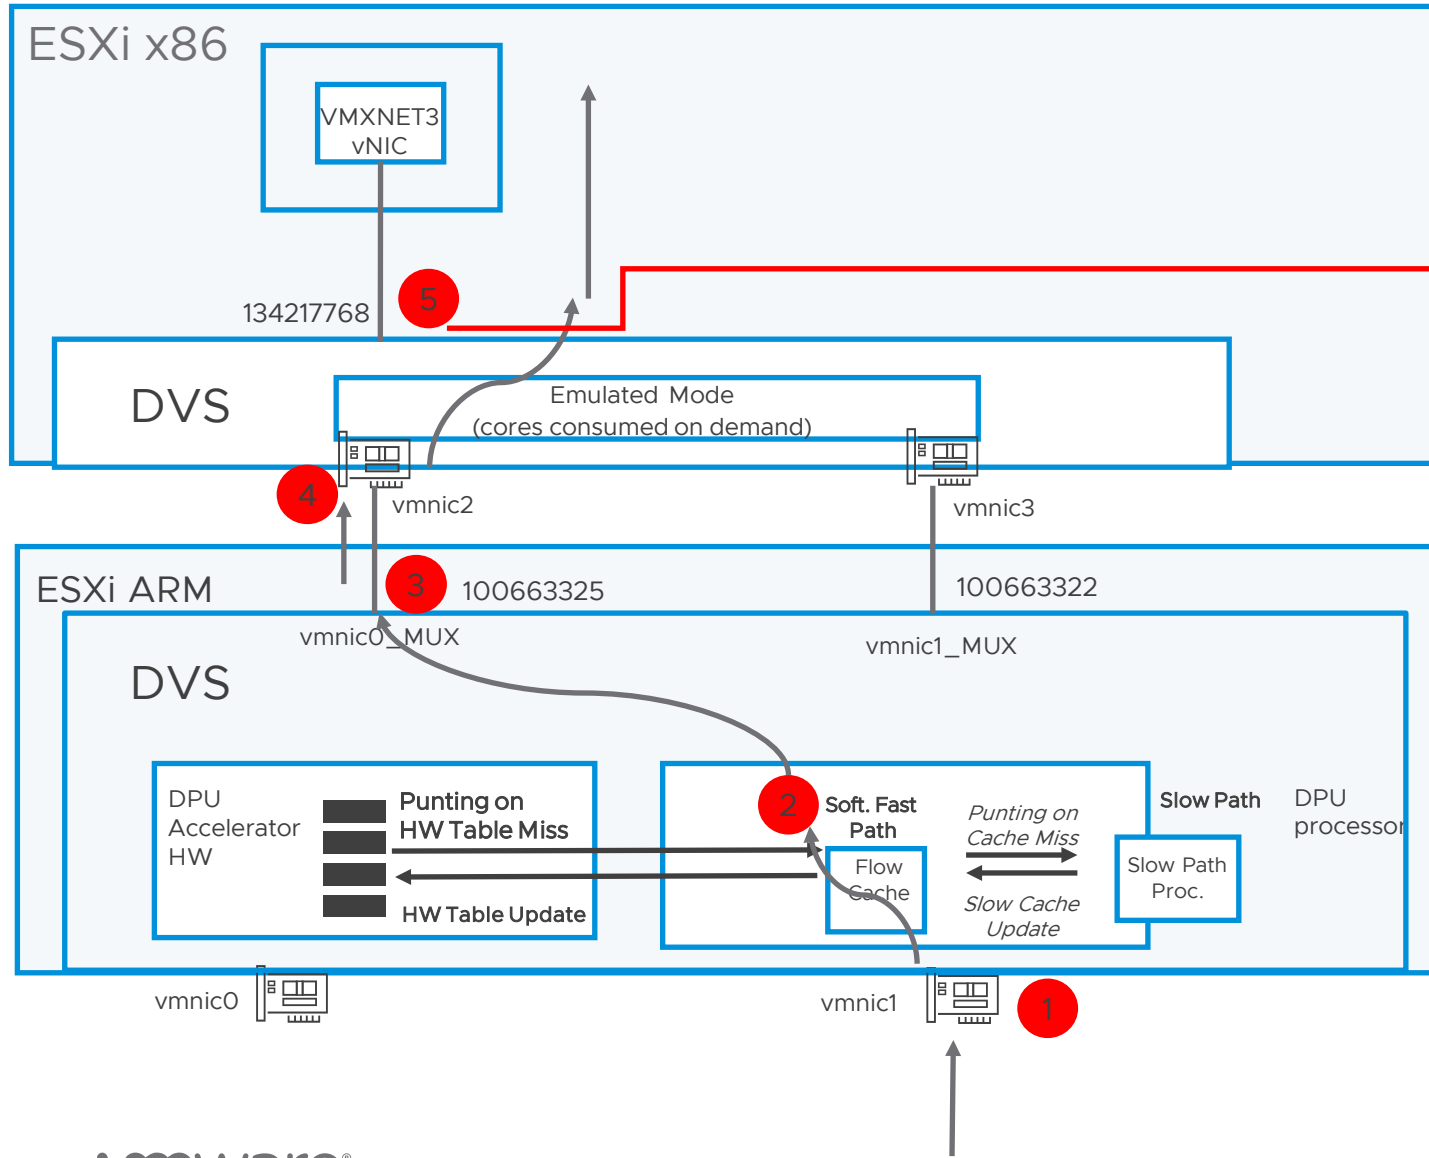
```
pktcap-uw --uplink vmnic2 --dir 2 -o - | tcpdump-uw -enr -
```

Packet:

```
00:50:56:a1:fc:f3 > 08:00:5c:a1:32:24, ethertype IPv4
(0x0800), length 74: 172.16.31.26.57448 > 172.16.31.25.31:
Flags [S], seq 2198858042, win 64240, options [mss
1460,sackOK,TS val 3020130186 ecr 0,nop,wscale 7], length 0
```

Packet Walk in Emulated Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.25 HTTP

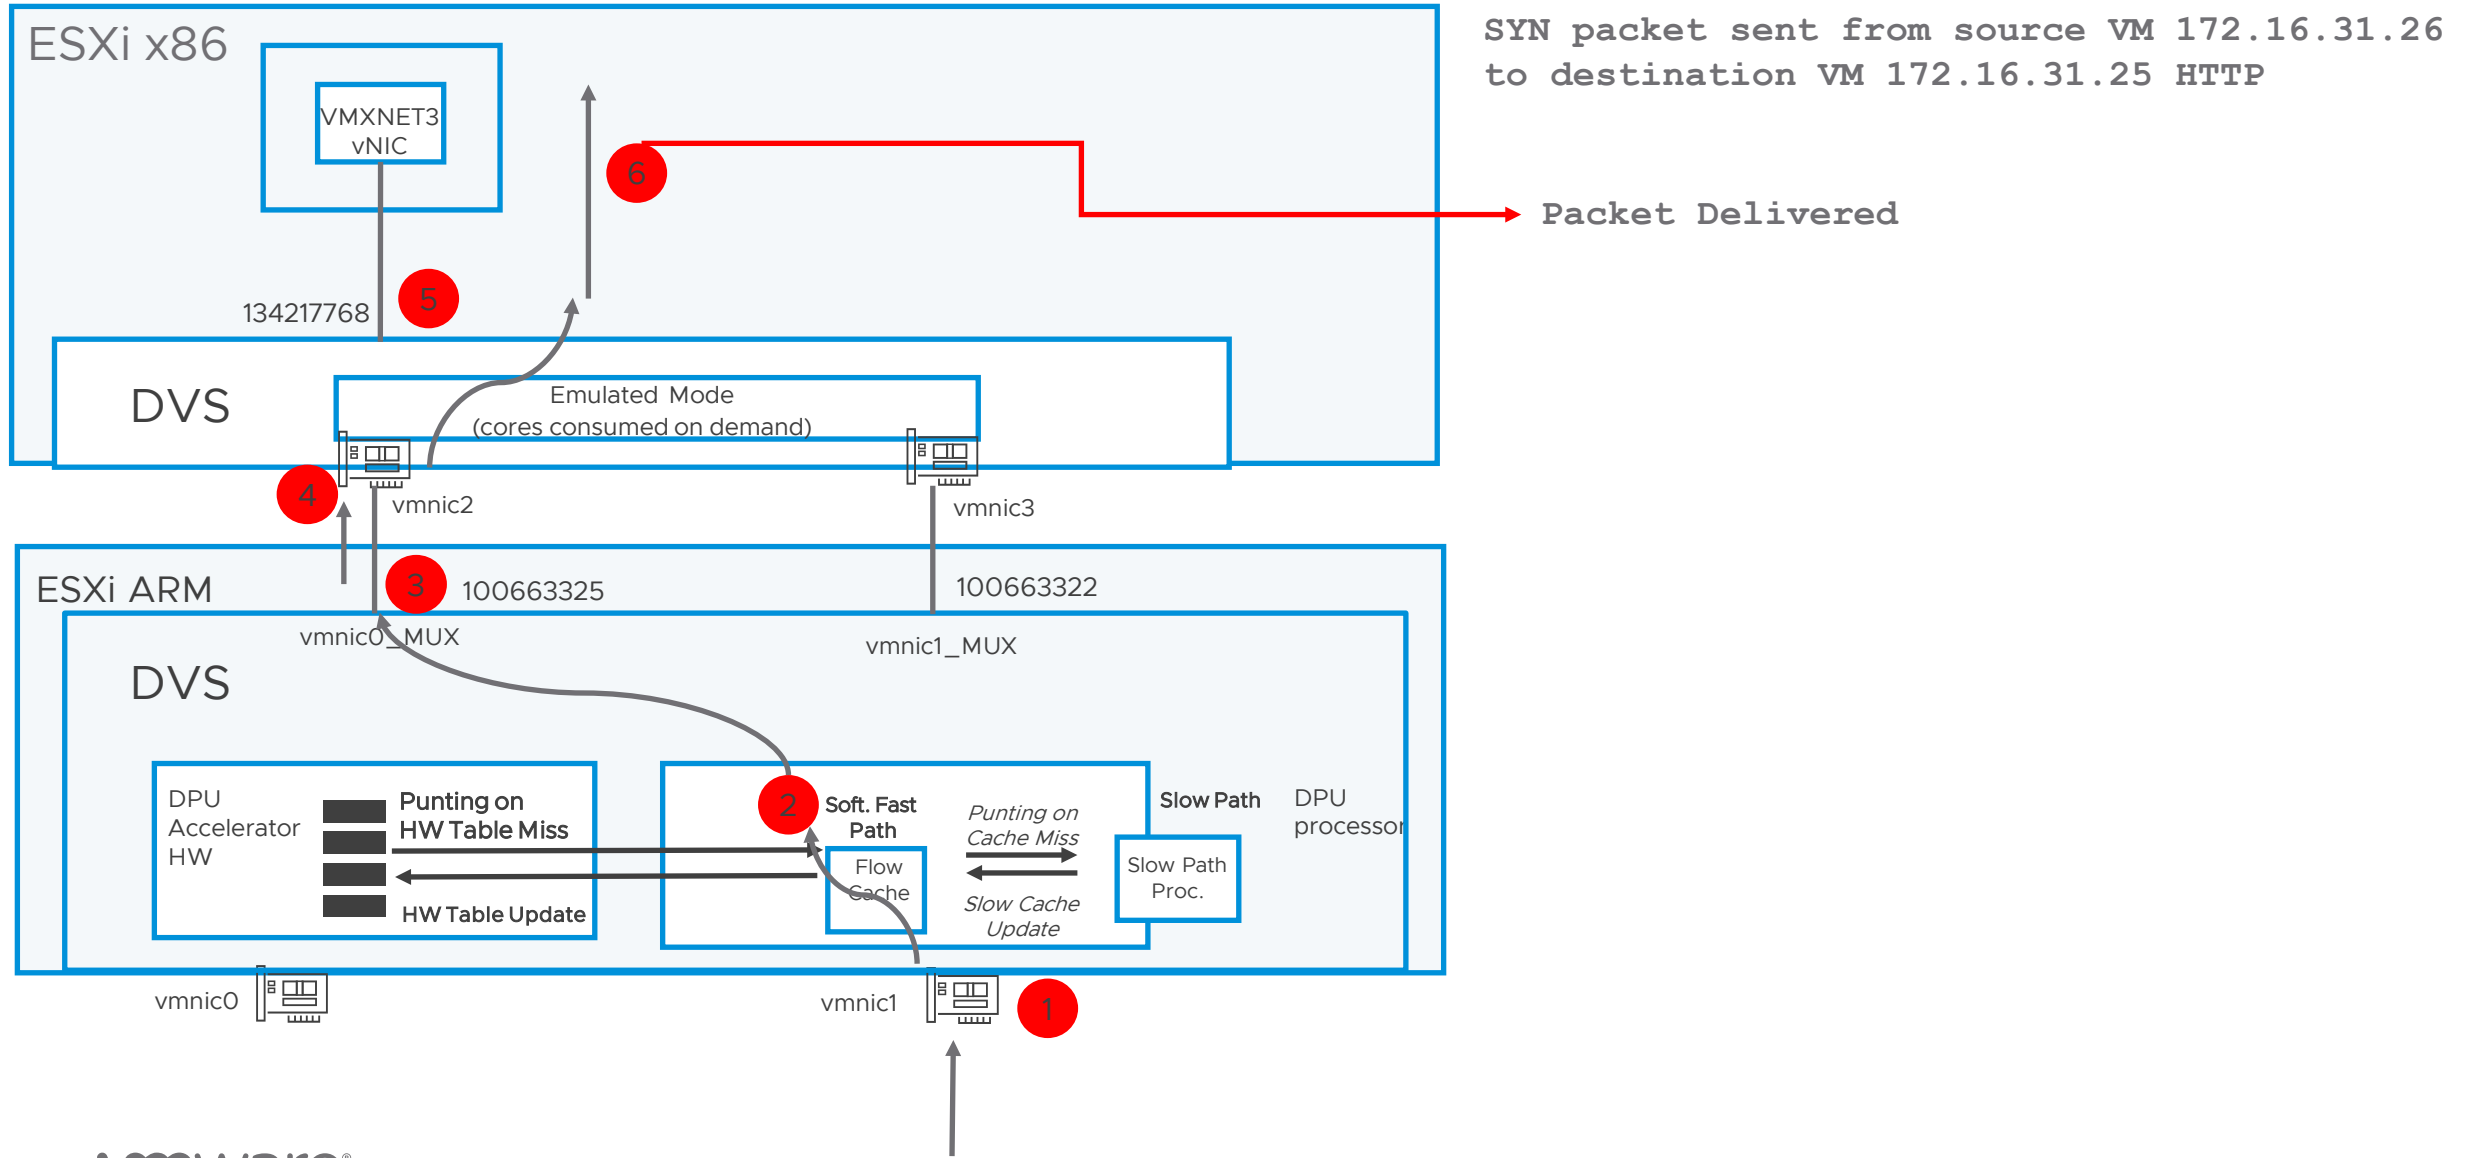
```
Port Info:
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
134217768    5      9 DvsPortset-1
00:50:56:a1:99:9e gs-ubuntu-01.eth0

Command:
pktcap-uw --uplink 134217768 --dir 2 -o - | tcpdump-uw -
enr -

Packet:
00:50:56:a1:fc:f3 > 00:50:56:a1:99:9e, ethertype IPv4
(0x0800), length 74: 172.16.31.26.57448 > 172.16.31.25.31:
Flags [S], seq 2198858042, win 64240, options [mss
1460,sackOK,TS val 3020130186 ecr 0,nop,wscale 7], length 0
```

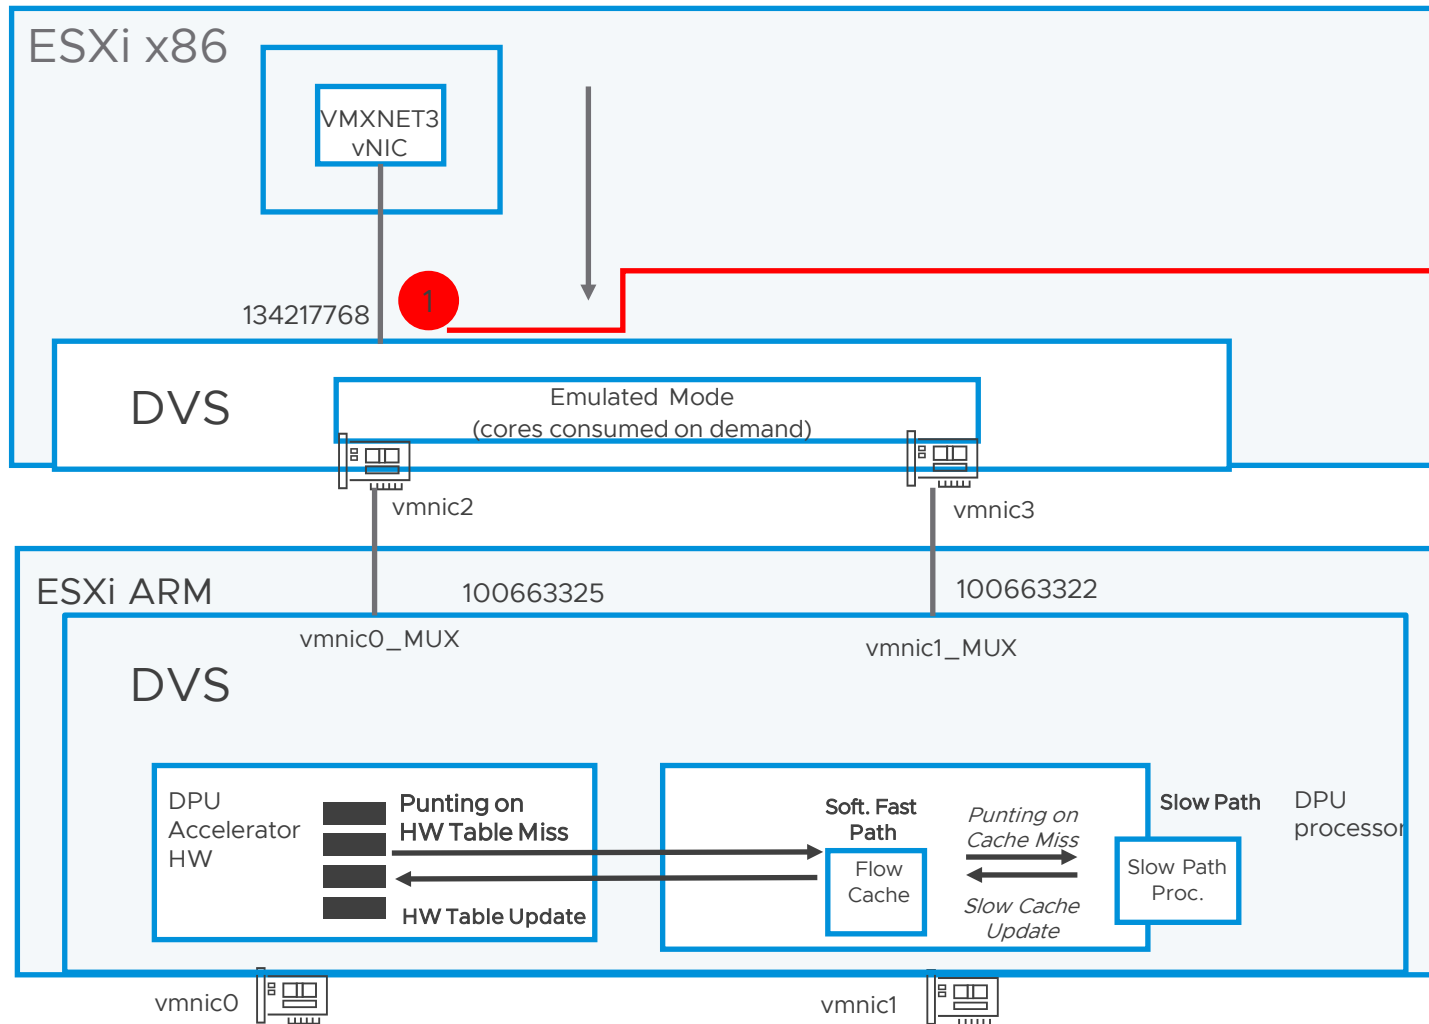
Packet Walk in Emulated Mode

Follow the RED Dot



Packet Walk in Emulated Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.25 HTTP to destination VM 172.16.31.26

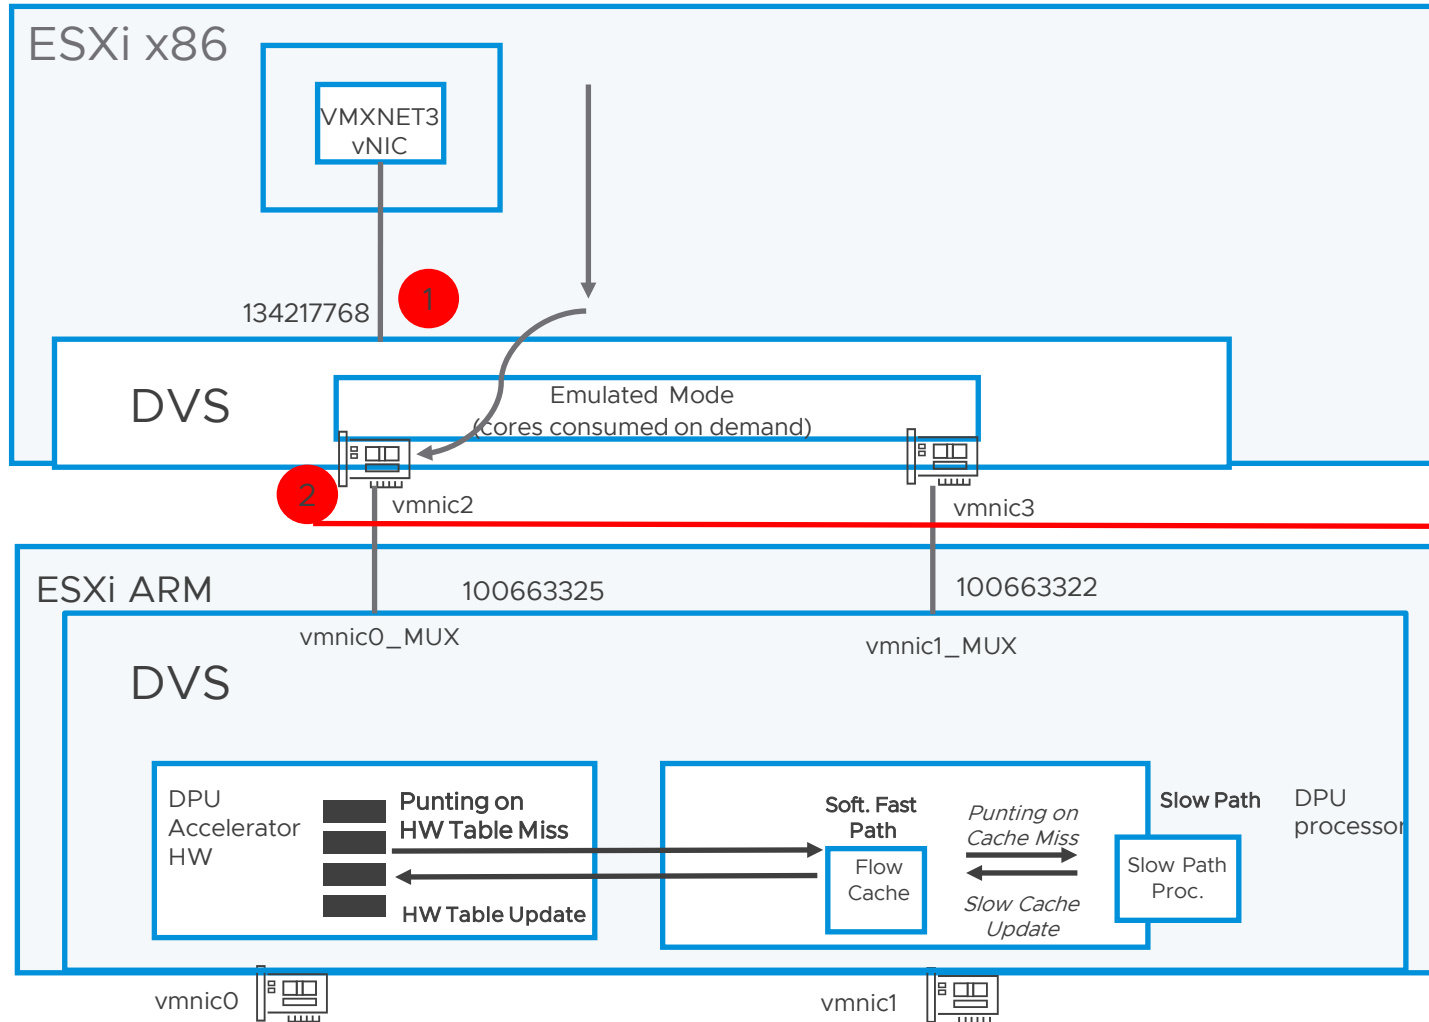
```
Port Info:
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
134217768    5      9 DvsPortset-1
00:50:56:a1:99:9e  gs-ubuntu-01.eth0

Command:
pktcap-uw --switchport 134217768 --dir 2 -o - | tcpdump-uw -enr -

Packet:
00:50:56:a1:99:9e > 00:50:56:a1:99:9e, ethertype IPv4 (0x0800), length 74: 172.16.31.25.80 > 172.16.31.26.80:
Flags [S.], seq 2193816220, ack 2198858043, win 65160, options [mss 1460,sackOK,TS val 4278092283 ecr 3020130186,nop,wscale 7], length 0
```

Packet Walk in Emulated Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.25 HTTP to destination VM 172.16.31.26

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
2281701394   4      0 DvsPortset-1
08:c0:eb:5d:ad:78  vmnic2
```

Command:

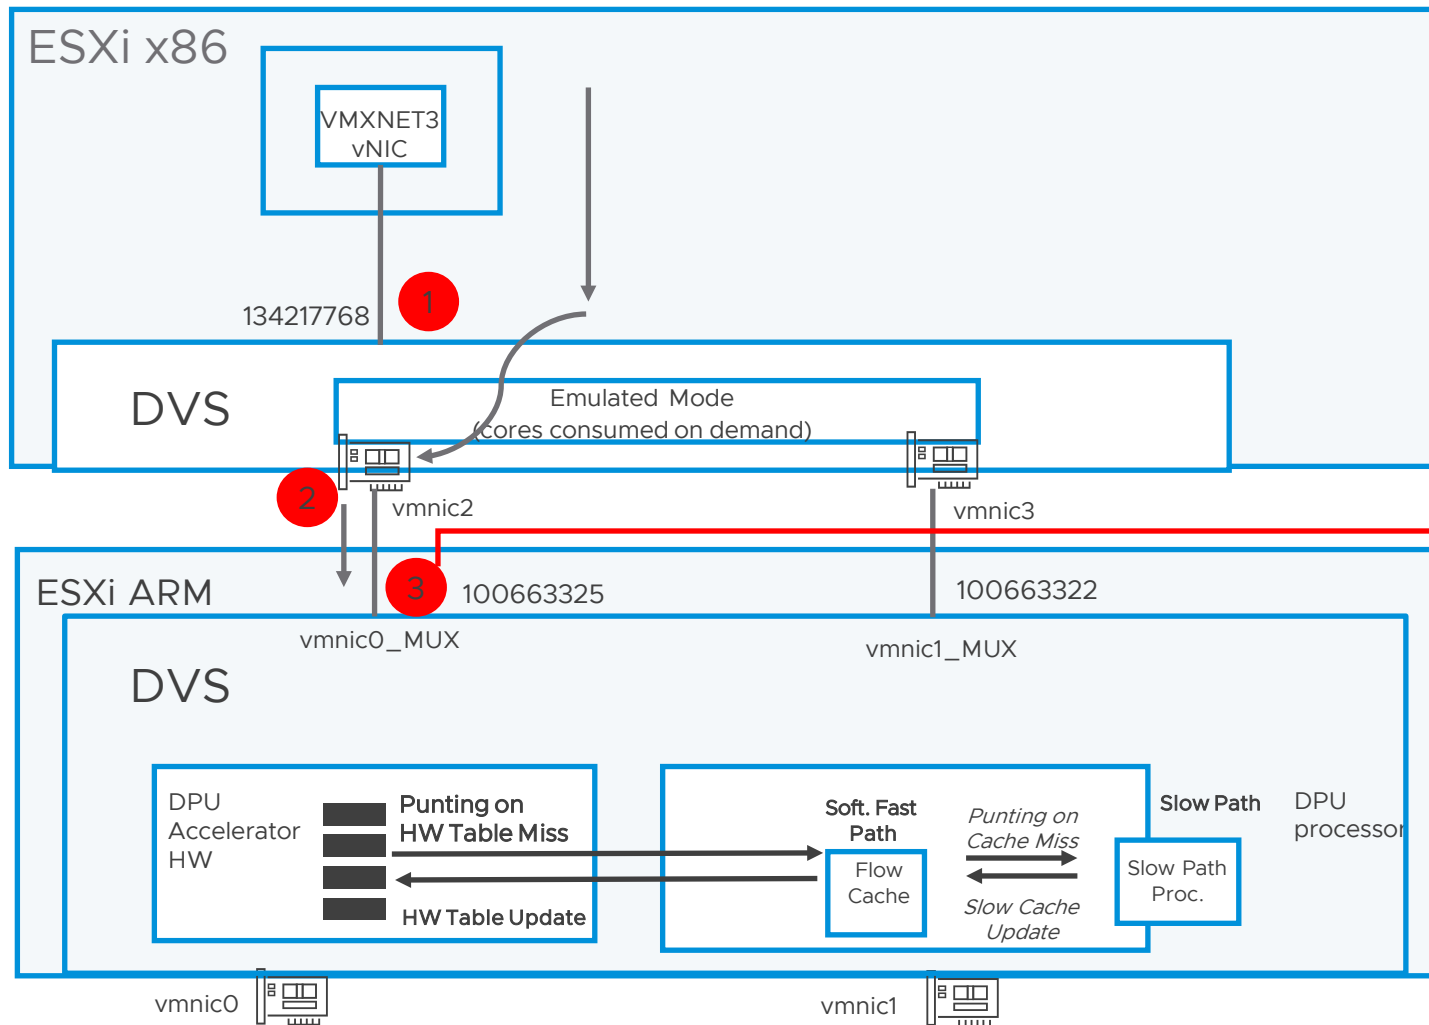
```
pktcap-uw --uplink vmnic2 --dir 2 -o - | tcpdump-uw -enr -
```

Packet:

```
00:50:56:a1:99:9e > 08:00:56:a1:99:9e, ethertype IPv4
(0x0800), length 74: 172.16.31.25.80 > 172.16.31.26.80:
Flags [S.], seq 2193816220, ack 2198858043, win 65160,
options [mss 1460,sackOK,TS val 4278092283 ecr
3020130186,nop,wscale 7], length 0
```

Packet Walk in Emulated Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.25 HTTP to destination VM 172.16.31.26

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
100663325    13      0 DvsPortset-1
08:c0:eb:5d:ae:08  vmnic0_MUX_PF16387
```

Command:

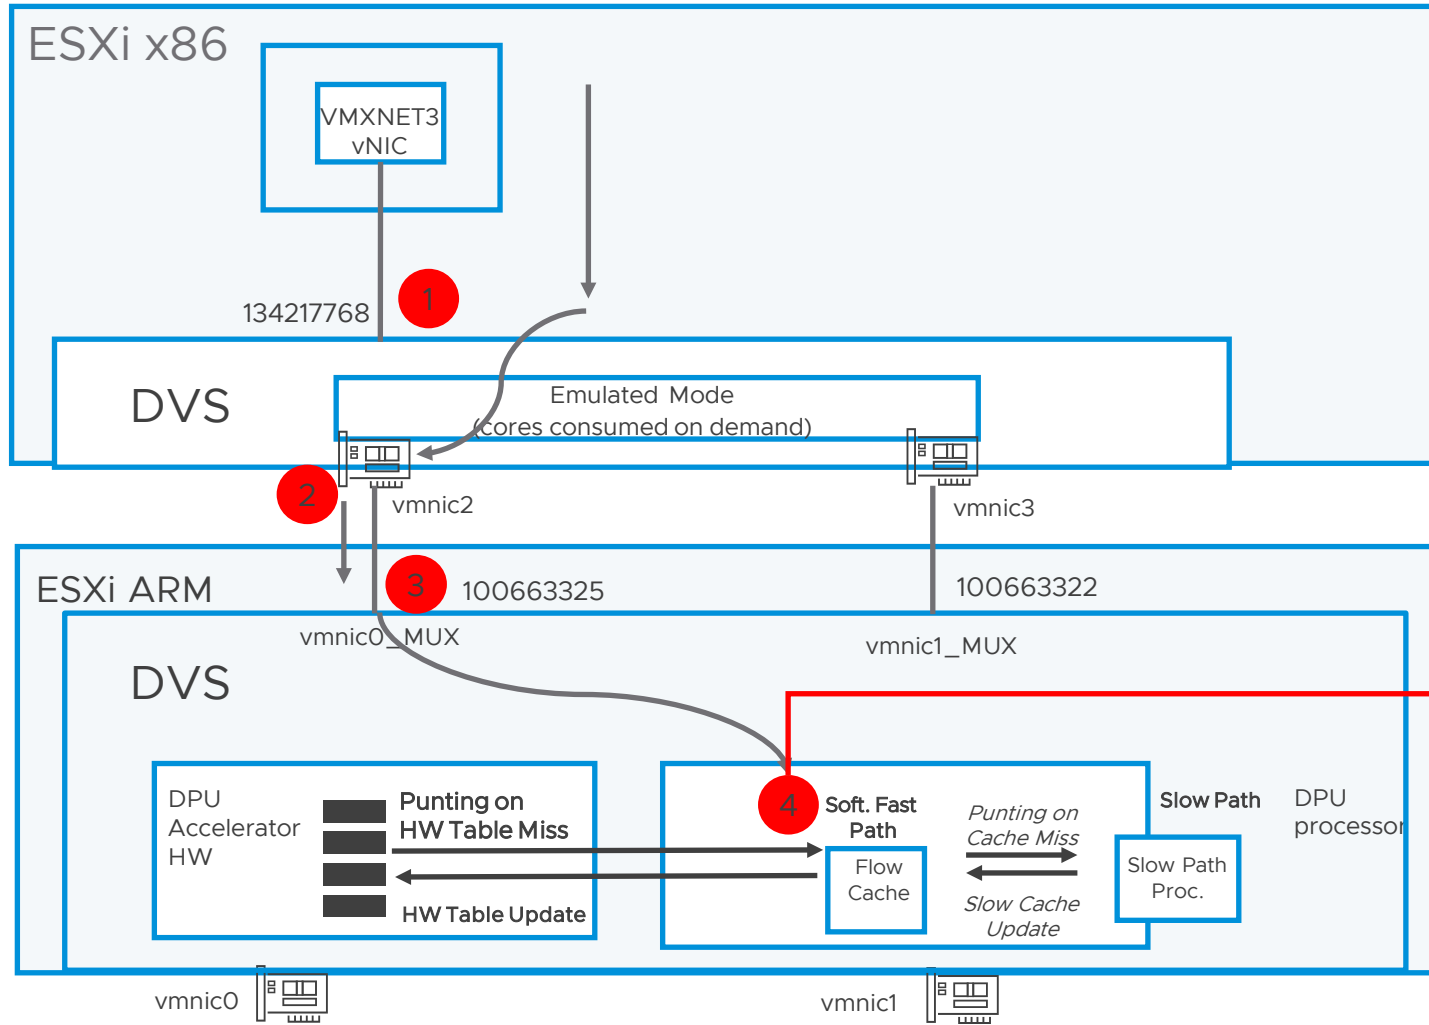
```
pktcap-uw -switchport 100663325 --dir 2 -o - | tcpdump-uw -
enr -
```

Packet:

```
00:50:56:a1:99:9e > 08:c0:eb:5d:ae:08, ethertype IPv4
(0x0800), length 74: 172.16.31.25.80 > 172.16.31.26.80:
Flags [S.], seq 2193816220, ack 2198858043, win 65160,
options [mss 1460,sackOK,TS val 4278092283 ecr
3020130186,nop,wscale 7], length 0
```


Packet Walk in Emulated Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.25 HTTP to destination VM 172.16.31.26

Command:

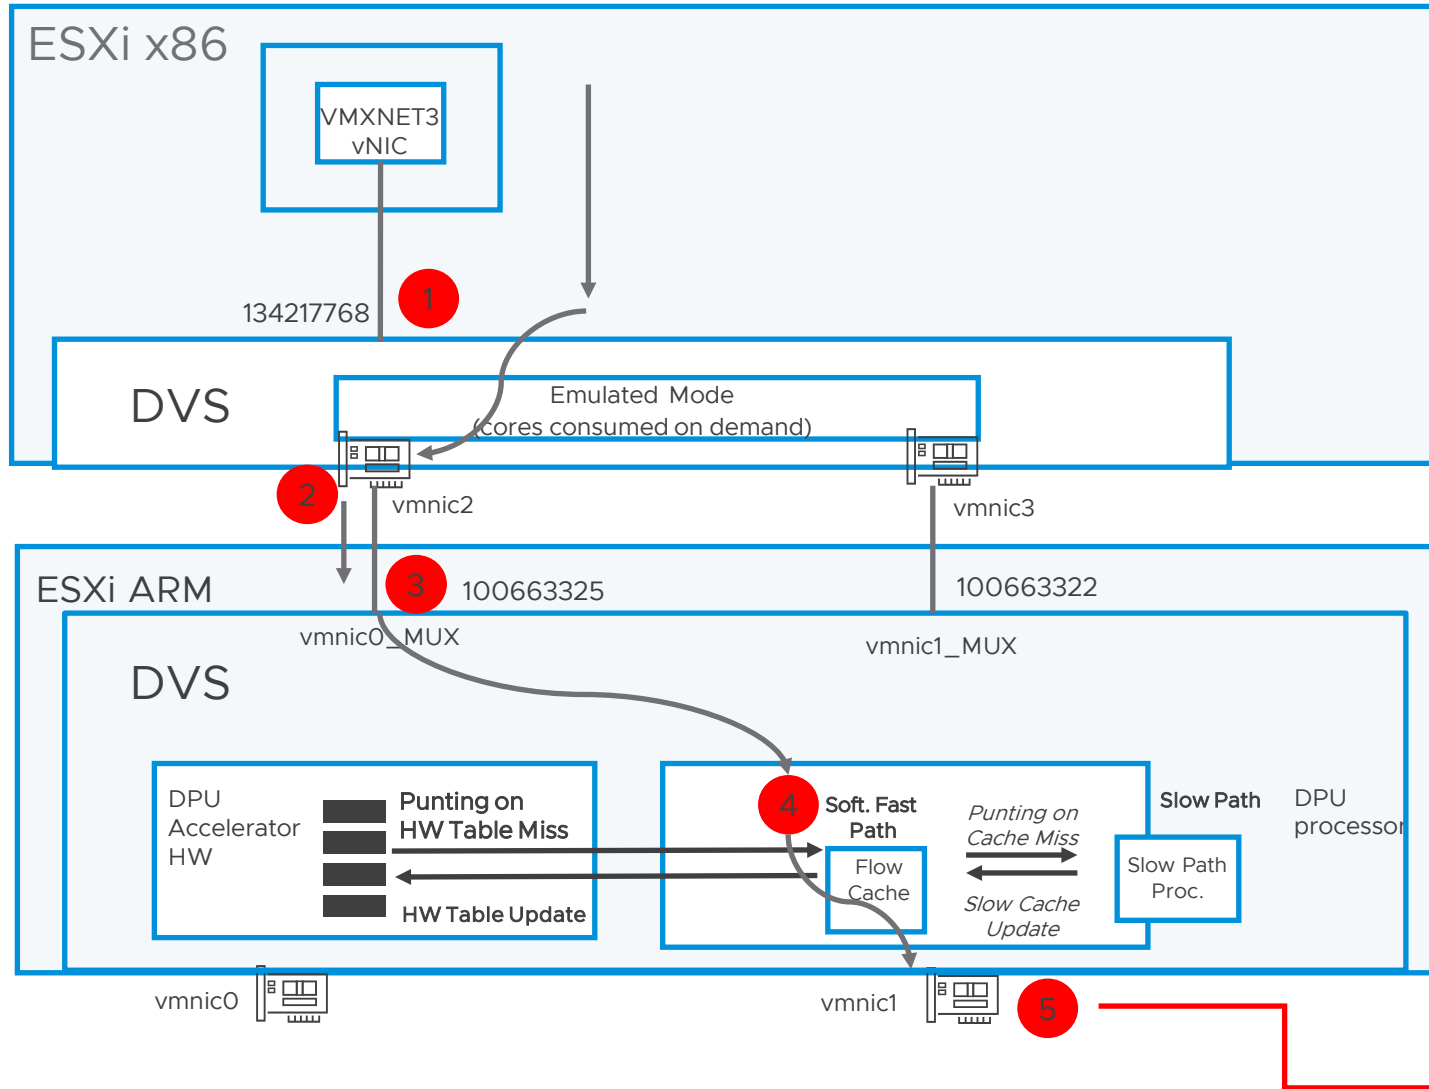
```
nsxcli ens flow-table dump
```

Flow Cache:

FT	dstMAC	srcMAC	VLAN	srcPort	dstIP	proto	VNI	srcPort/type	dstPort/code	Actions	hwHits
L4	00:50:56:a1:99:9e	00:50:56:a1:fc:f3	0	3	172.16.31.26						6
										bmap:0x80800080	
										inval(s):150 cg:2123 dp:0x7 len:704; GENEVE DECAP; DFW on dstPort; 0 0	
L4	00:50:56:a1:fc:f3	00:50:56:a1:99:9e	0	7	172.16.31.25						6
										bmap:0x20000c0	0
										inval(s):138 cg:2123 dp:0x3 len:814; DFW on srcPort; VNI: 74752; GENEVE ENCAP VNI: 74752; 0 0	

Packet Walk in Emulated Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.25 HTTP to destination VM 172.16.31.26

```

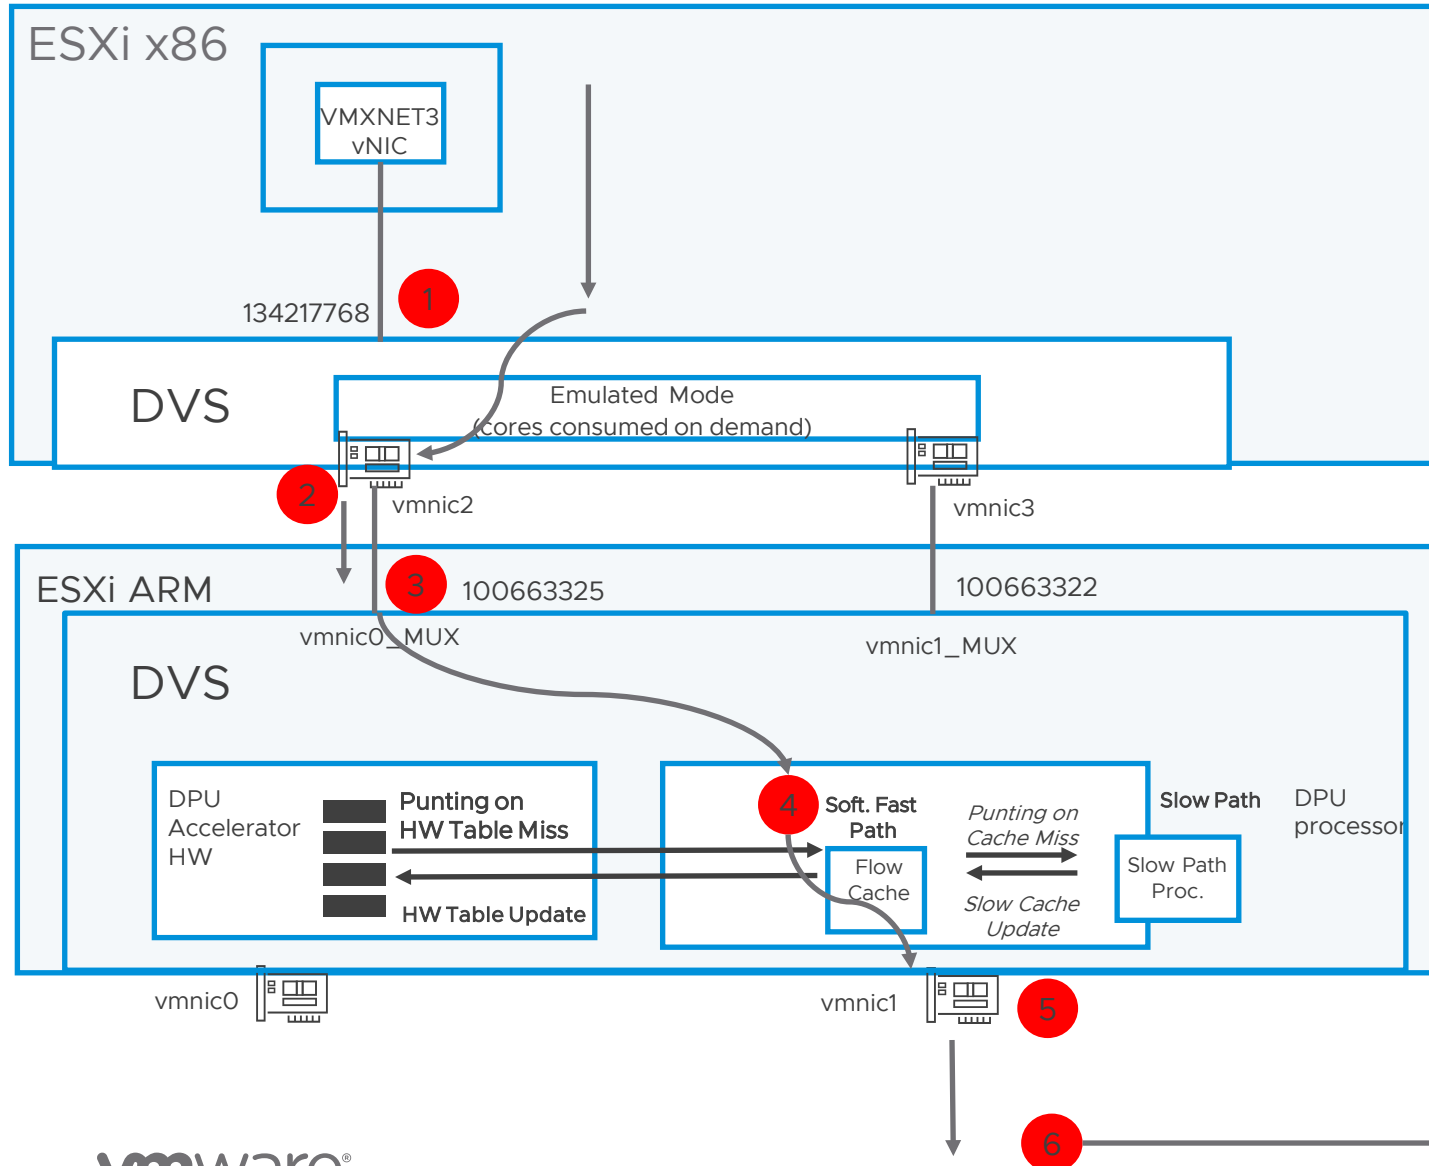
Port Info:
[root@localhost:~] net-stats -l
PortNum          Type SubType SwitchName      MACAddress
ClientName
2248146967       4      0 DvsPortset-1
08:c0:eb:5d:ae:0f vmnic1

Command:
pktcap-uw --uplink vmnic1 --dir 2 -o - | tcpdump-uw -enr -

Packet:
00:50:56:67:eb:13 > 00:50:56:63:03:03, ethertype IPv4
(0x0800), length 132: 192.168.4.60.56554 >
192.168.4.51.6081: Geneve, Flags [none], vni 0x12400, proto
TEB (0x6558), options [8 bytes]: 00:50:56:a1:99:9e >
00:50:56:a1:99:9e, ethertype IPv4 (0x0800), length 74:
172.16.31.25.80 > 172.16.31.26.57444: Flags [S.], seq
2193816220, ack 2198858043, win 65160, options [mss
1460,sackOK,TS val 4278092283 ecr 3020130186,nop,wscale 7],
length 0
    
```

Packet Walk in Emulated Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.25 HTTP to destination VM 172.16.31.26

Packet Walk Over DPU/SmartNic UPTv2 Mode

Enable UPTv2 Mode

Edit Settings | gs-ubuntu-01

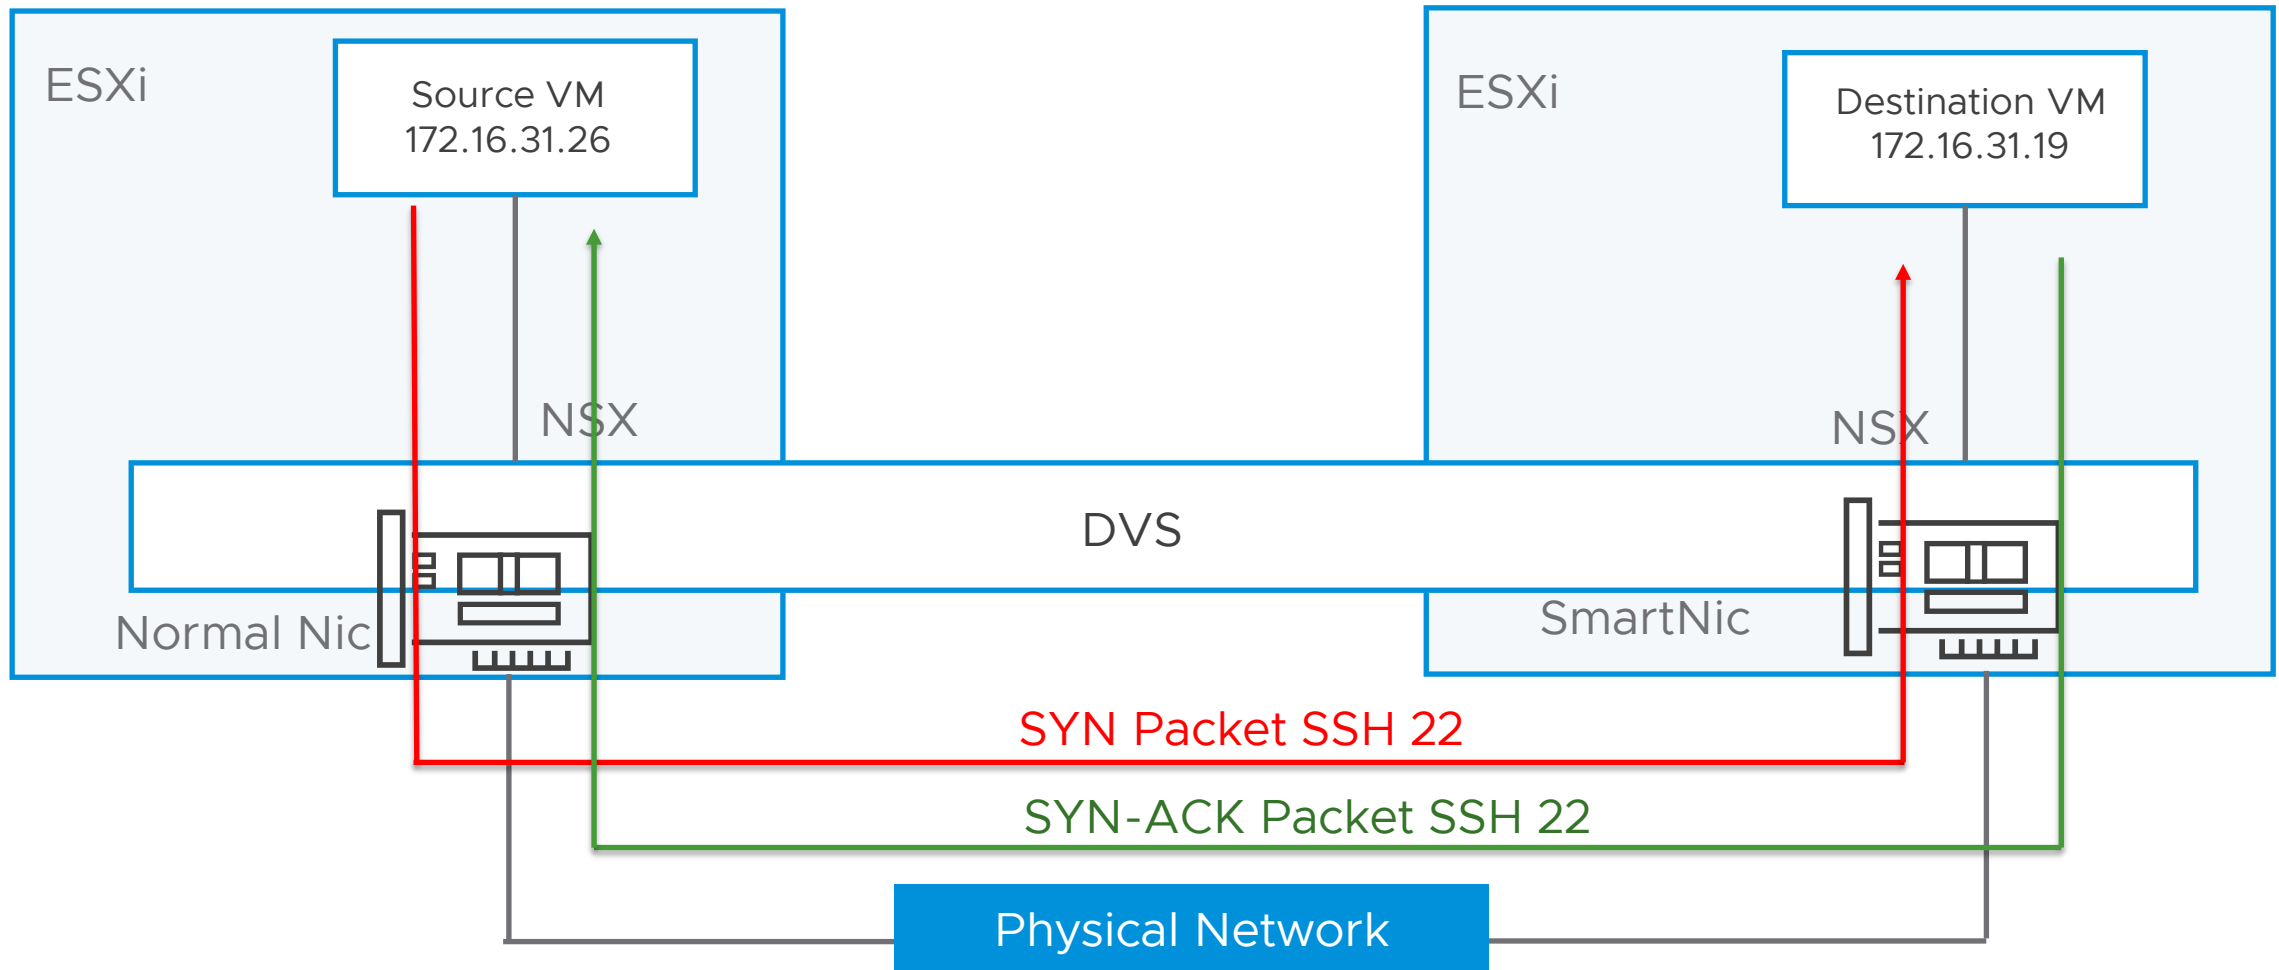
Virtual Hardware | VM Options | Advanced Parameters

ADD NEW DEVICE ▾

> CPU	2	ⓘ
> Memory *	2	GB ▾
> Hard disk 1	16	GB ▾
> SCSI controller 0	LSI Logic Parallel	⋮
▾ Network adapter 1 *	cluster1-segment1	Connected
Status	<input checked="" type="checkbox"/> Connect At Power On	
Port ID	c380b93d-ee7f-4eb5-92	
Adapter Type	VMXNET 3	▾
UPT Support	<input checked="" type="checkbox"/> Use UPT Support	
MAC Address	00:50:56:a1:99:9e	Automatic ▾
> CD/DVD drive 1	Client Device	Connect At Power On
> Video card	Specify custom settings	▾
> SATA controller 0	AHCI	⋮
> Security Devices	Not Configured	
> Other	Additional Hardware	

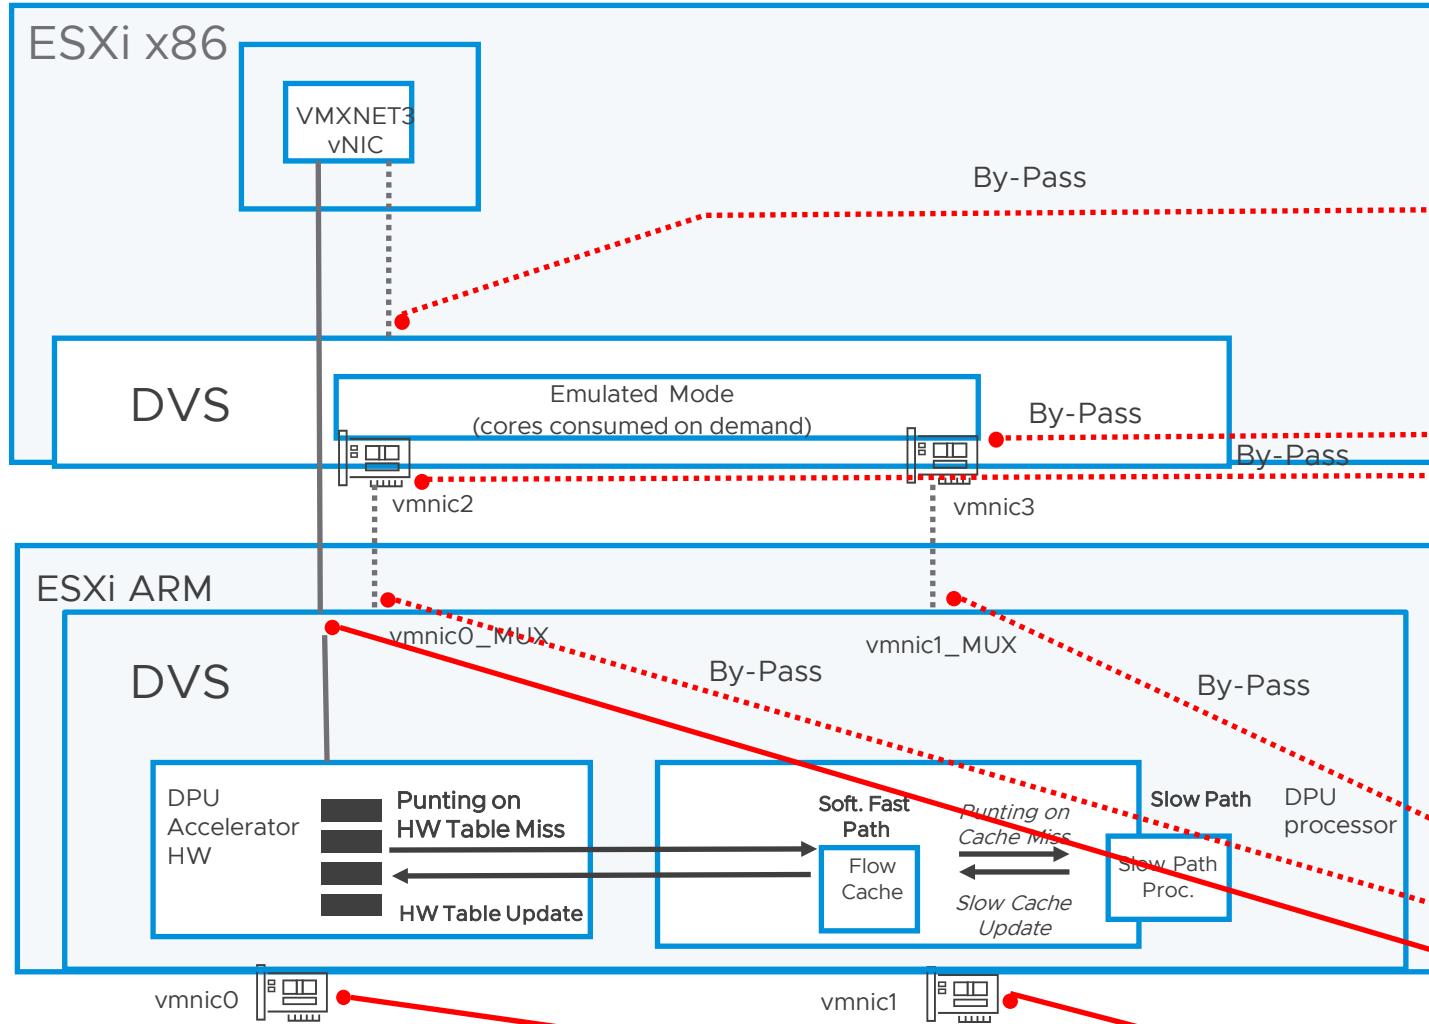
CANCEL OK

Datapath UPTv2 Mode



Mapping of ports in UPTv2 Mode

Follow the RED Dot



ESXi x86

PORT-ID USED-BY
67108872 Management
67108874 vmk7379
100663307 Management
100663317 Shadow of vmnic0
100663319 Shadow of vmnic1
100663320 vmk0
100663322 vmk2
100663323 vmk3
134217740 Management
134217745 Shadow of vmnic3
134217747 Shadow of vmnic2
134217764 2192756:cluster1-ubuntu-vm8.et
167772189 Management
167772191 Shadow of vusb0
167772192 vmk1
2214592521 vmnic128
2248146964 vmnic0
2248146966 vmnic1
2281701392 vmnic3
2281701394 vmnic2

TEAM-PNIC DNAME
n/a esxioVswitch0
vmnic128 esxioVswitch0
n/a DvsPortset-0
n/a DvsPortset-0
n/a DvsPortset-0
vmnic0 DvsPortset-0
vmnic1 DvsPortset-0
vmnic1 DvsPortset-0
n/a DvsPortset-1
n/a DvsPortset-1
n/a DvsPortset-1
vmnic2 DvsPortset-1
n/a vSwitchBMC
n/a vSwitchBMC
vusb0 vSwitchBMC
- esxioVswitch0
- DvsPortset-0
- DvsPortset-0
- DvsPortset-1
- DvsPortset-1

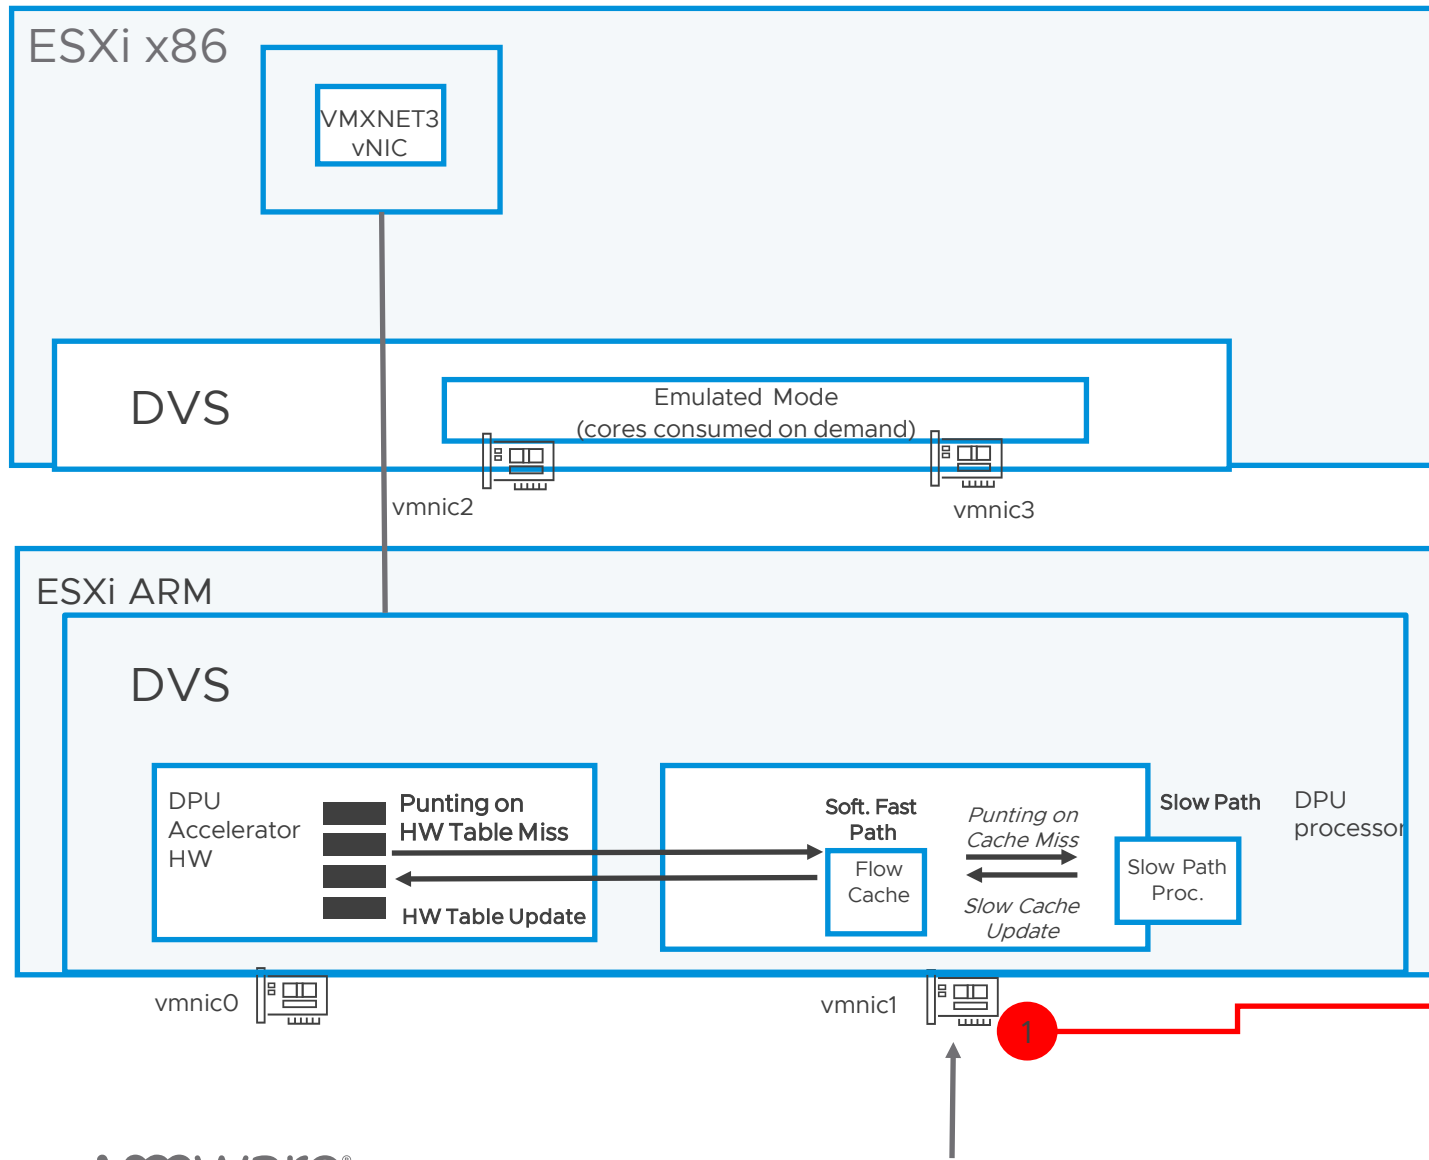
ESXi ARM

PORT-ID USED-BY
67108871 Management
67108873 vmk7379
100663306 Management
100663310 vmk10
100663311 vmk11
100663313 vmk50
100663320 Shadow of vmnic1
100663321 vdr-vdrPort
100663322 vmnic1_MUX_PF20483
100663324 Shadow of vmnic0
100663325 vmnic0_MUX_PF16387
100663329 cluster1-ubuntu-vm8.eth0.7c520
134217747 Management
134217749 Shadow of vmnic2
134217750 vmk255
2214592520 vmnic129
2248146967 vmnic1
2248146971 vmnic0
2281701396 vmnic2

TEAM-PNIC DNAME
n/a
vmnic129 esxioesxioVswitch0Vswitch0
n/a DvsPortset-1
vmnic1 DvsPortset-1
vmnic0 DvsPortset-1
void DvsPortset-1
n/a DvsPortset-1
vmnic1 DvsPortset-1
n/a DvsPortset-1
vmnic0 DvsPortset-1
vmnic1 DvsPortset-1
n/a vSwitchESXioBMC
n/a vSwitchESXioBMC
vmnic2 vSwitchESXioBMC
- esxioVswitch0
- DvsPortset-1
- DvsPortset-1
- vSwitchESXioBMC

Packet Walk in UPTv2 Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.19 SSH

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
2248146967   4    0    DvsPortset-1
08:c0:eb:5d:ae:0f  vmnic1
```

Command:

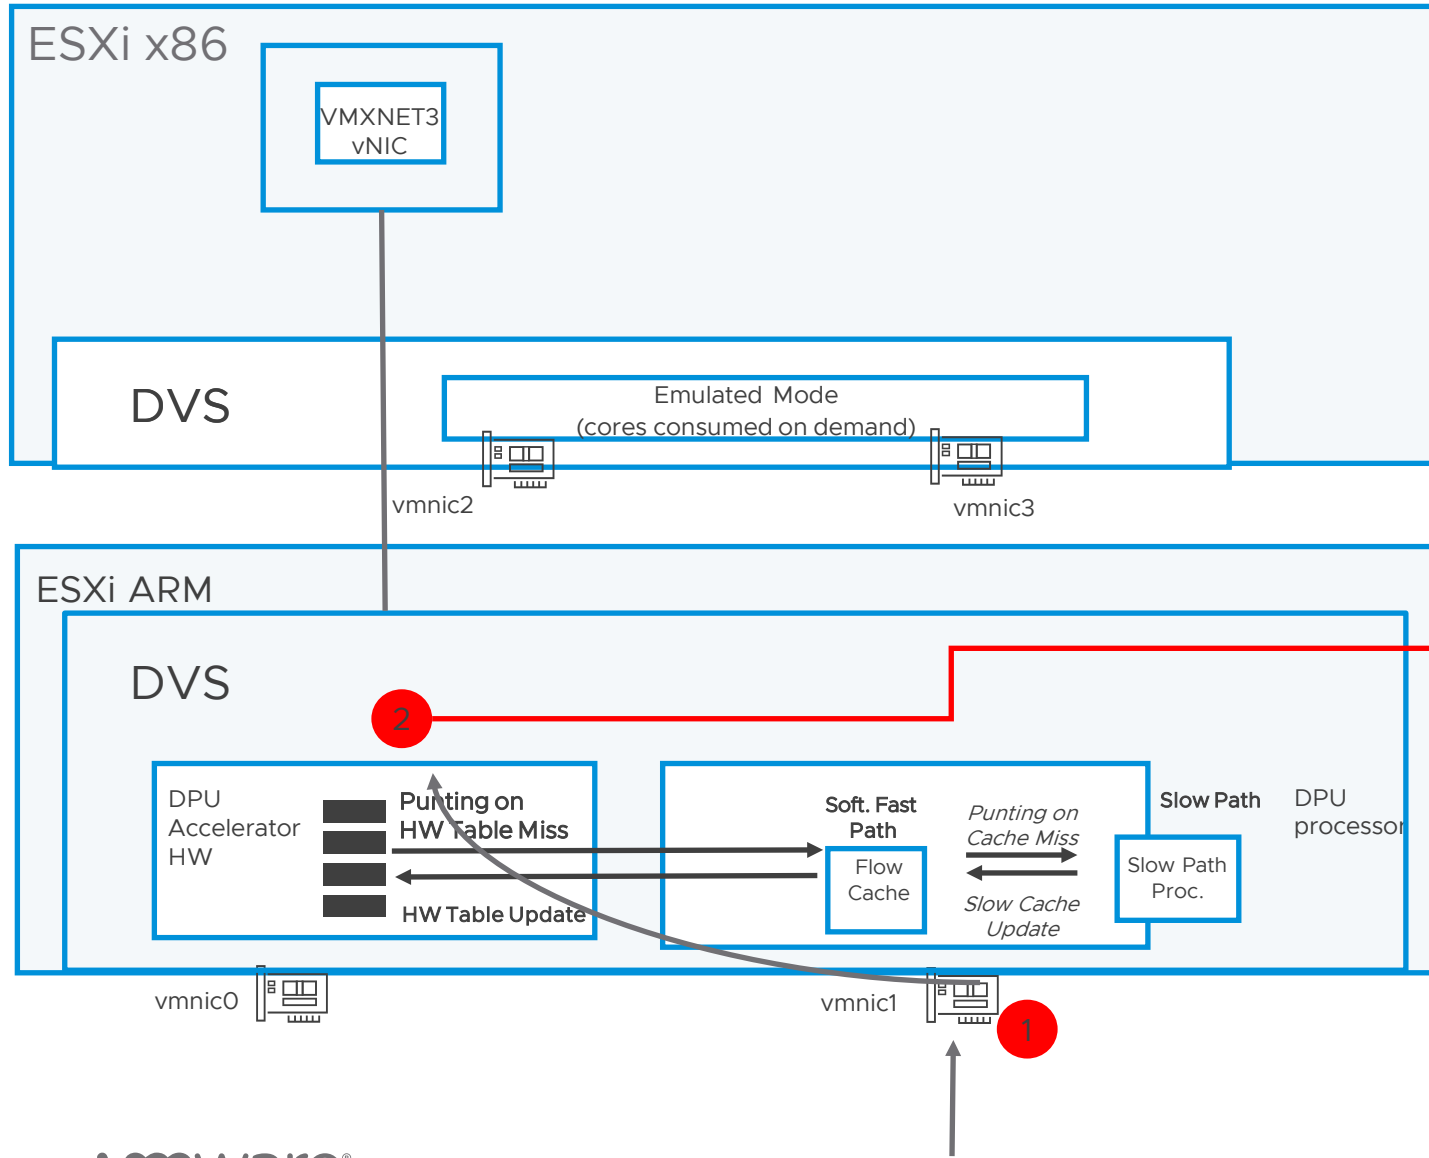
```
pktcap-uw --uplink vmnic1 --dir 2 -o - | tcpdump-uw -enr -
```

Packet:

```
00:50:56:63:04:5c > 00:50:56:63:04:5c, ethertype IPv4 (0x0800), length 132: 192.168.4.57.49583 > 192.168.4.60.6081: Geneve, Flags [none], vni 0x12400, proto TEB (0x6558), options [8 bytes]: 00:50:56:a1:fc:f3 > 00:50:56:a1:fc:f3, ethertype IPv4 (0x0800), length 74: 172.16.31.26.60186 > 172.16.31.19: Flags [S], seq 402077945, win 64240, options [mss 1460,sackOK,TS val 279309001 ecr 0,nop,wscale 7], length 0
```


Packet Walk in UPTv2 Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.19 SSH

Command:

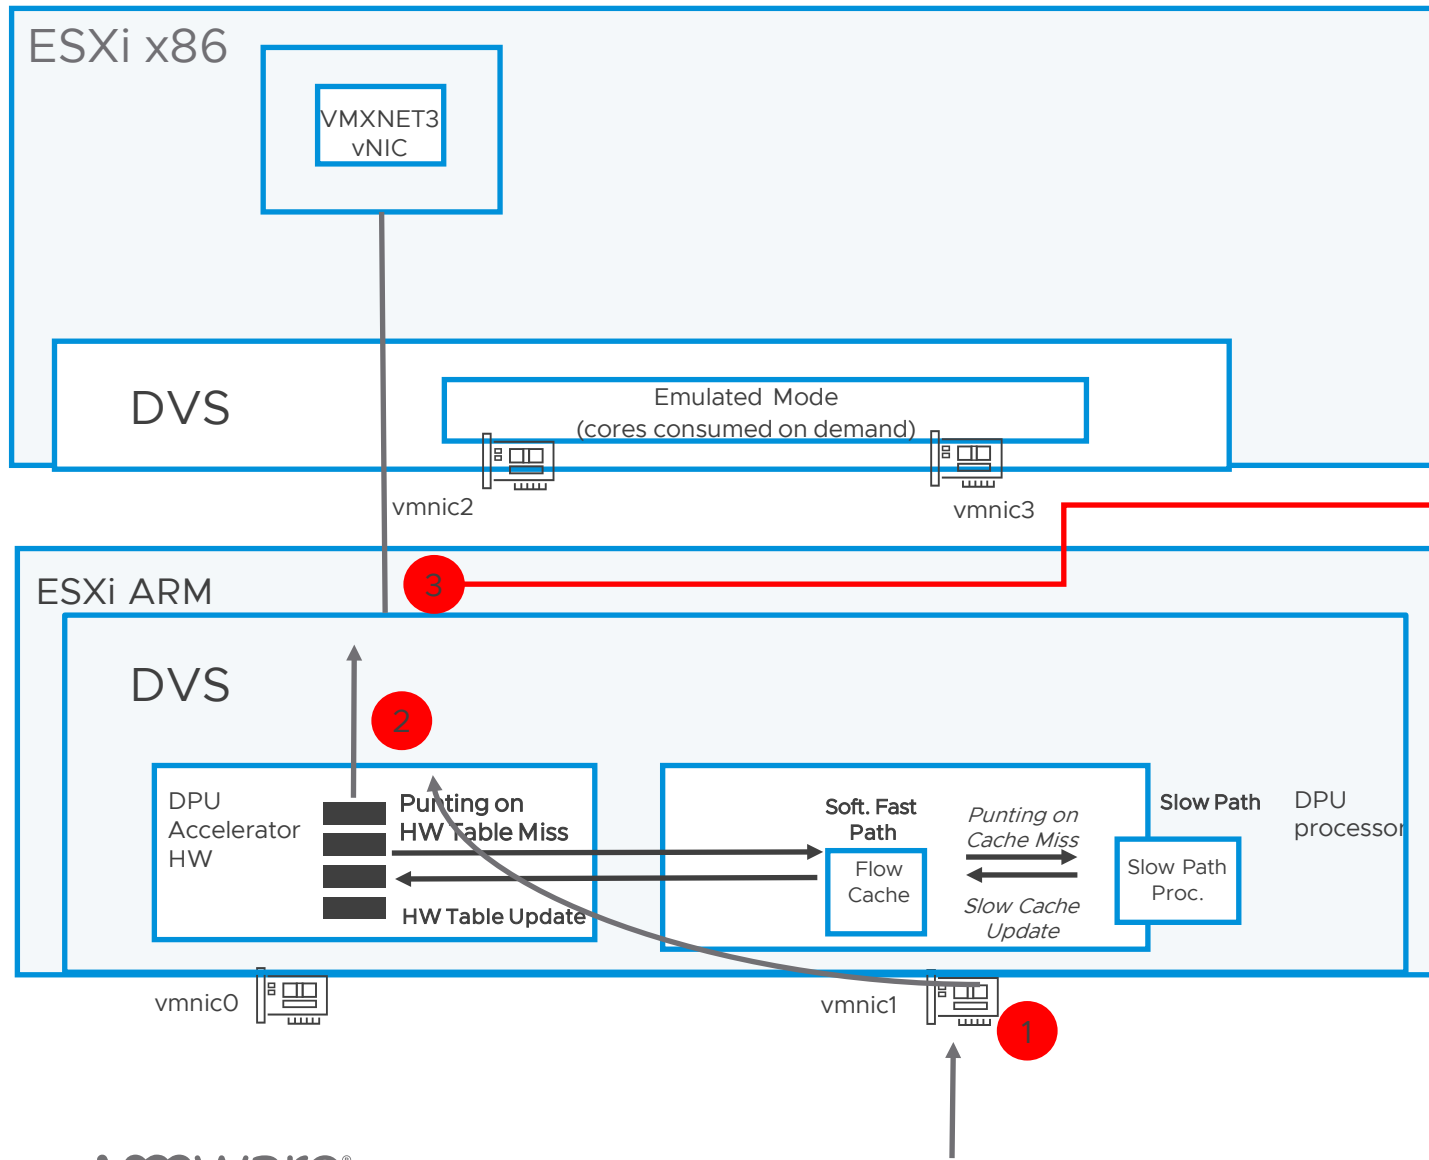
```
nsxdataplanecli ens flow-table dump
```

Flow Cache:

FT	dstMAC	srcMAC	VLAN	srcPort	dstIP	proto	VNI	srcPort/type	dstPort/code	Actions	hwHits
L4	00:50:56:a1:87:fe	00:50:56:a1:fc:f3	0	3	172.16.31.26	TCP	74752	60186	0	bmap:0x80800080	6
inval(s):95 cg:2115 dp:0xa len:704; GENEVE DECAP; DFW on dstPort; 0 0											
L4	00:50:56:a1:fc:f3	00:50:56:a1:87:fe	0	10	172.16.31.19	TCP	22	0	0	bmap:0x20000c0	0
cg:2115 dp:0x3 len:814; DFW on srcPort; VNI: 74752; GENEVE ENCAP VNI: 74752; 0 0											

Packet Walk in UPTv2 Mode

Follow the RED Dot



SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.19 SSH

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
100663329    13   14 DvsPortset-1   cluster1-ubuntu-vm8.eth0.7c520dd0
```

Command:

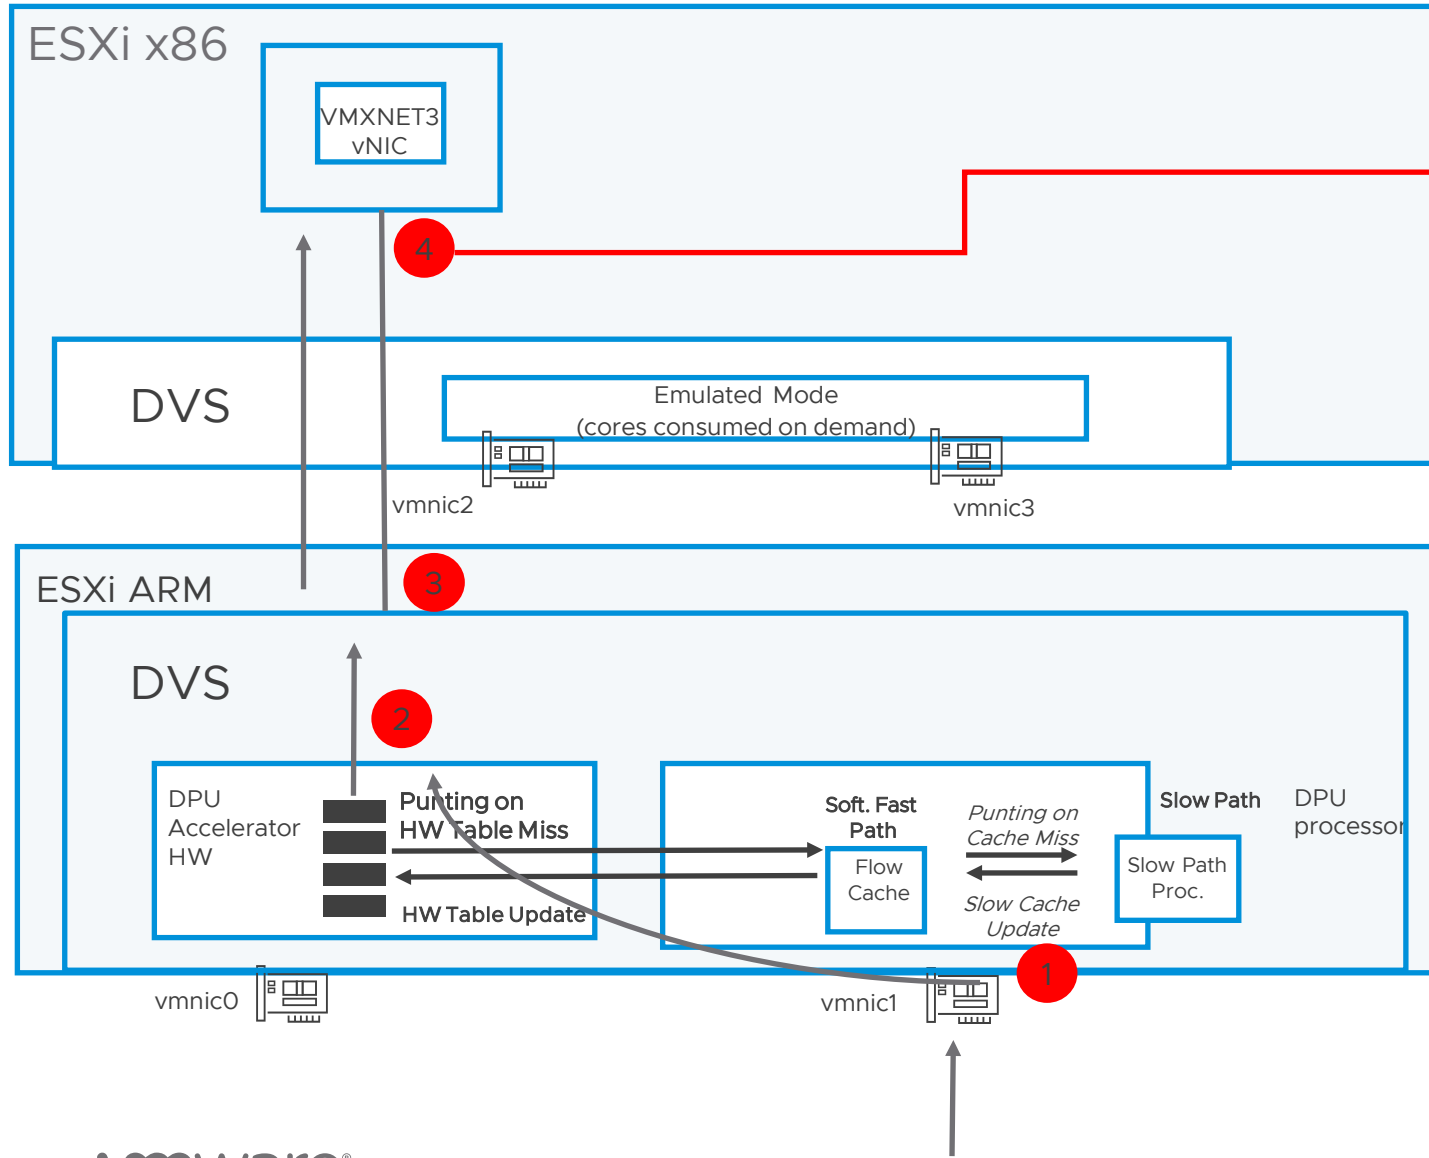
```
pktcap-uw --switchport 100663329 --dir 2 -o - | tcpdump-uw -enr -
```

Packet:

```
00:50:56:a1:fc:f3 > 00:50:56:a1:fc:f3, ethertype IPv4 (0x0800), length 74: 172.16.31.26.60186 > 172.16.31.19: Flags [S], seq 402077945, win 64240, options [mss 1460,sackOK,TS val 279309001 ecr 0,nop,wscale 7], length 0
```

Packet Walk in UPTv2 Mode

Follow the RED Dot

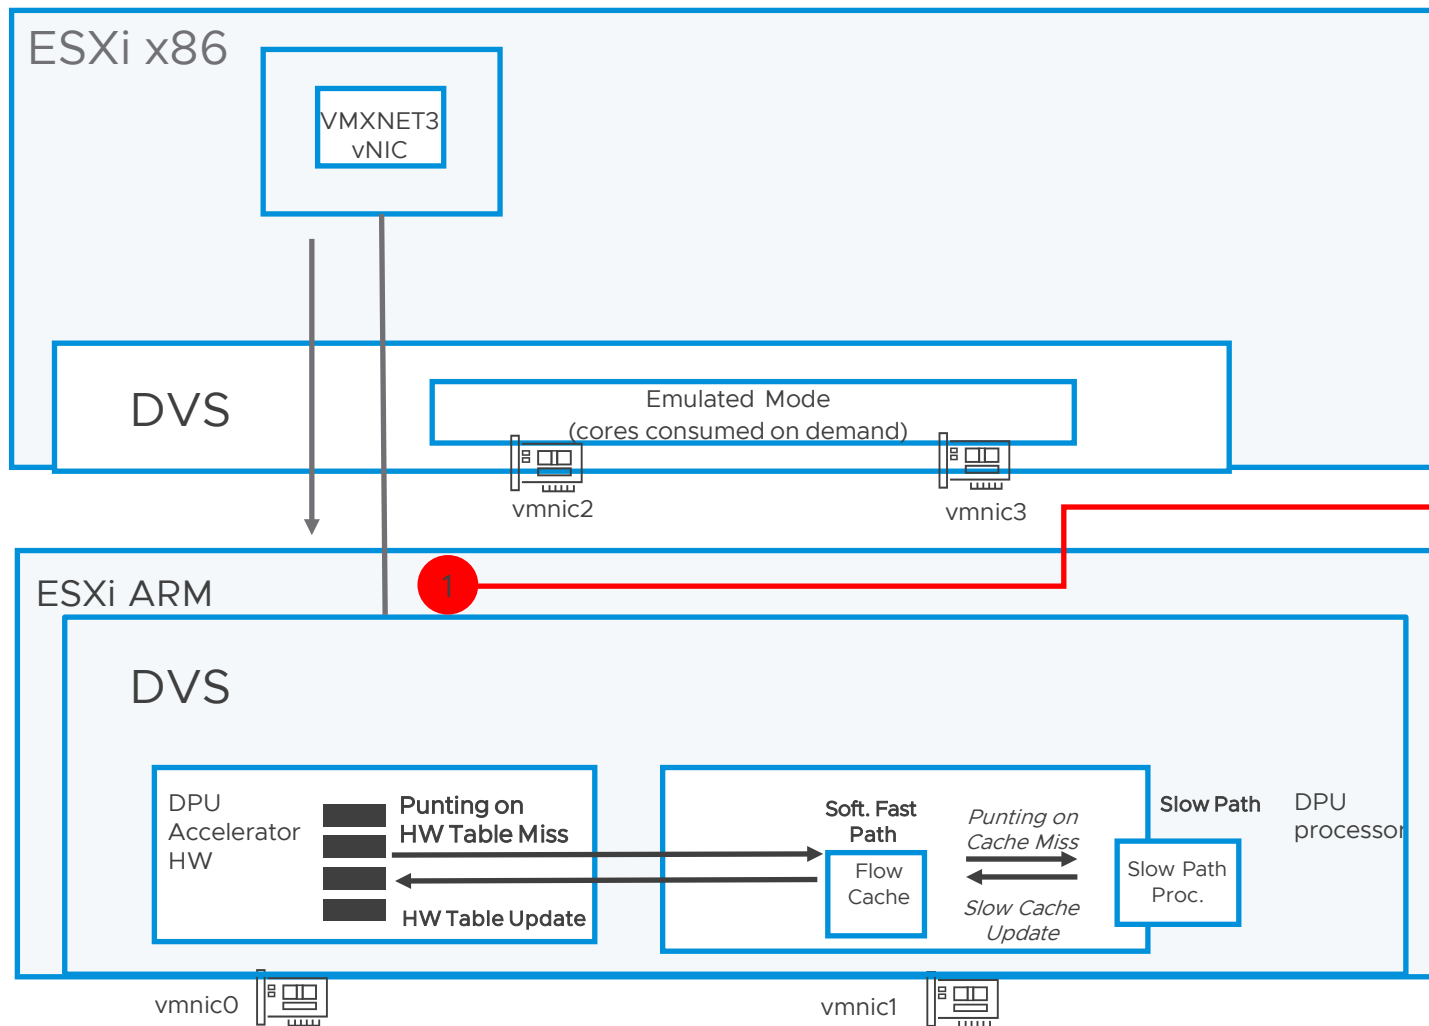


SYN packet sent from source VM 172.16.31.26 to destination VM 172.16.31.19 SSH

Packet Delivered.

Packet Walk in UPTv2 Mode

Follow the RED Dot



SYN-ACK packet sent from source VM 172.16.31.19 SSH to destination VM 172.16.31.26

Port Info:

```
[root@localhost:~] net-stats -l
PortNum      Type SubType SwitchName      MACAddress
ClientName
100663329    13   14 DvsPortset-1
00:50:56:a1:87:fe cluster1-ubuntu-vm8.eth0.7c520dd0
```

Command:

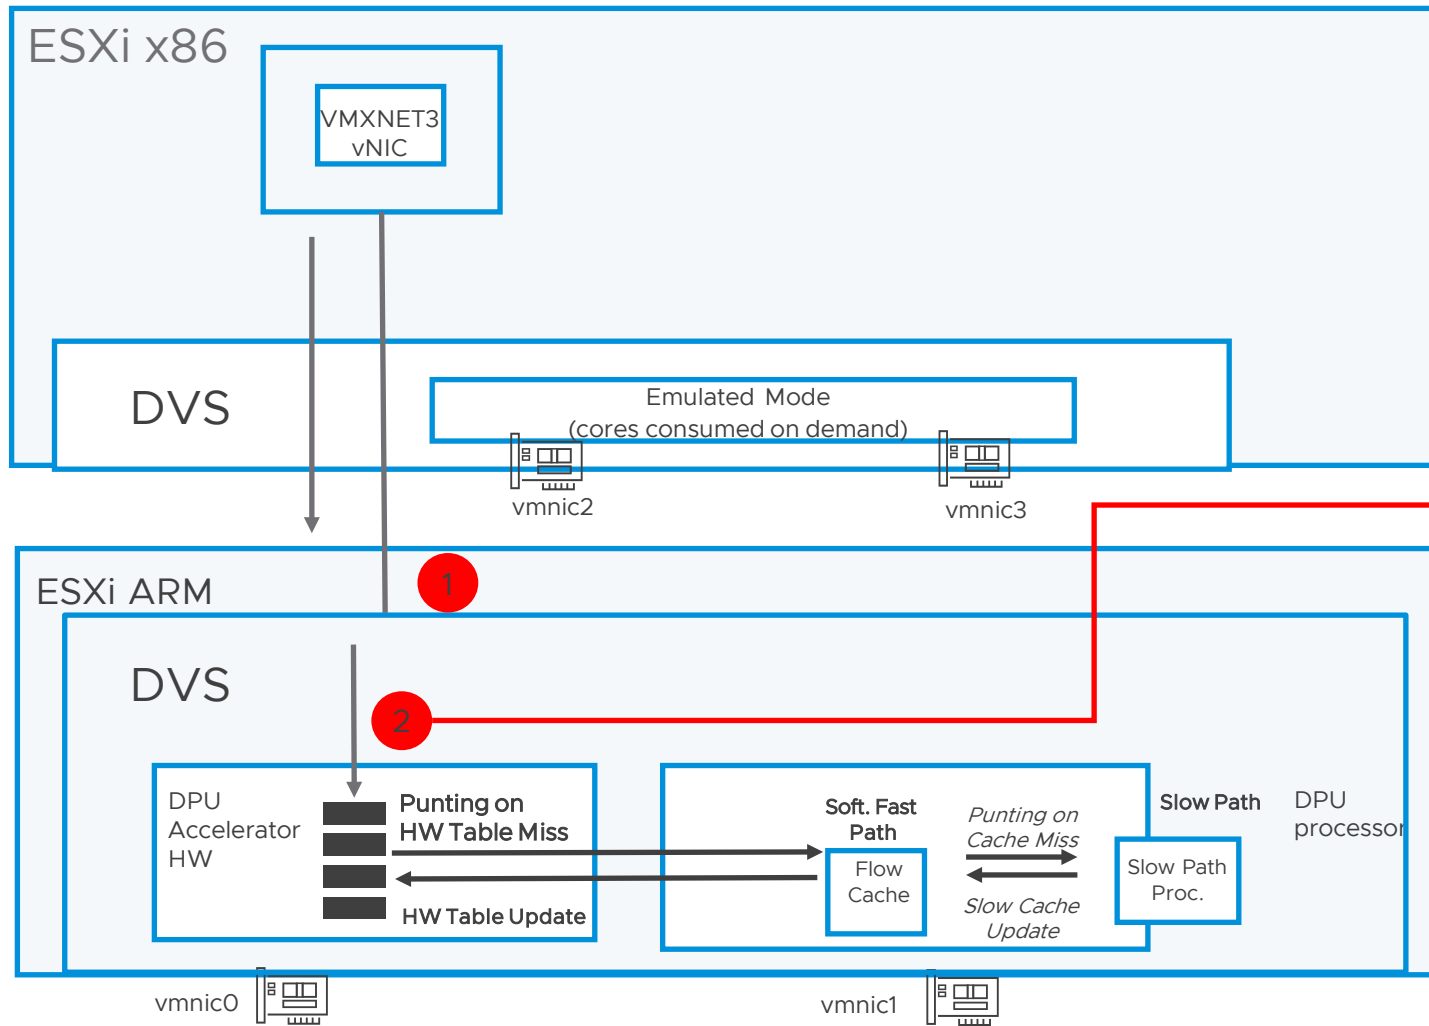
```
pktcap-uw --switchport 100663329 --dir 2 -o - | tcpdump-uw -
enr -
```

Packet:

```
00:50:56:a1:87:fe > 00:50:56:a1:87:fe, ethertype IPv4
(0x0800), length 74: 172.16.31.19.22 > 172.16.31.26.22:
Flags [S.], seq 3759600490, ack 402077946, win 65160,
options [mss 1460,sackOK,TS val 3501164822 ecr
279309001,nop,wscale 7], length 0
```

Packet Walk in UPTv2 Mode

Follow the RED Dot



SYN-ACK packet sent from source VM
172.16.31.19 SSH to destination VM
172.16.31.26

Command:

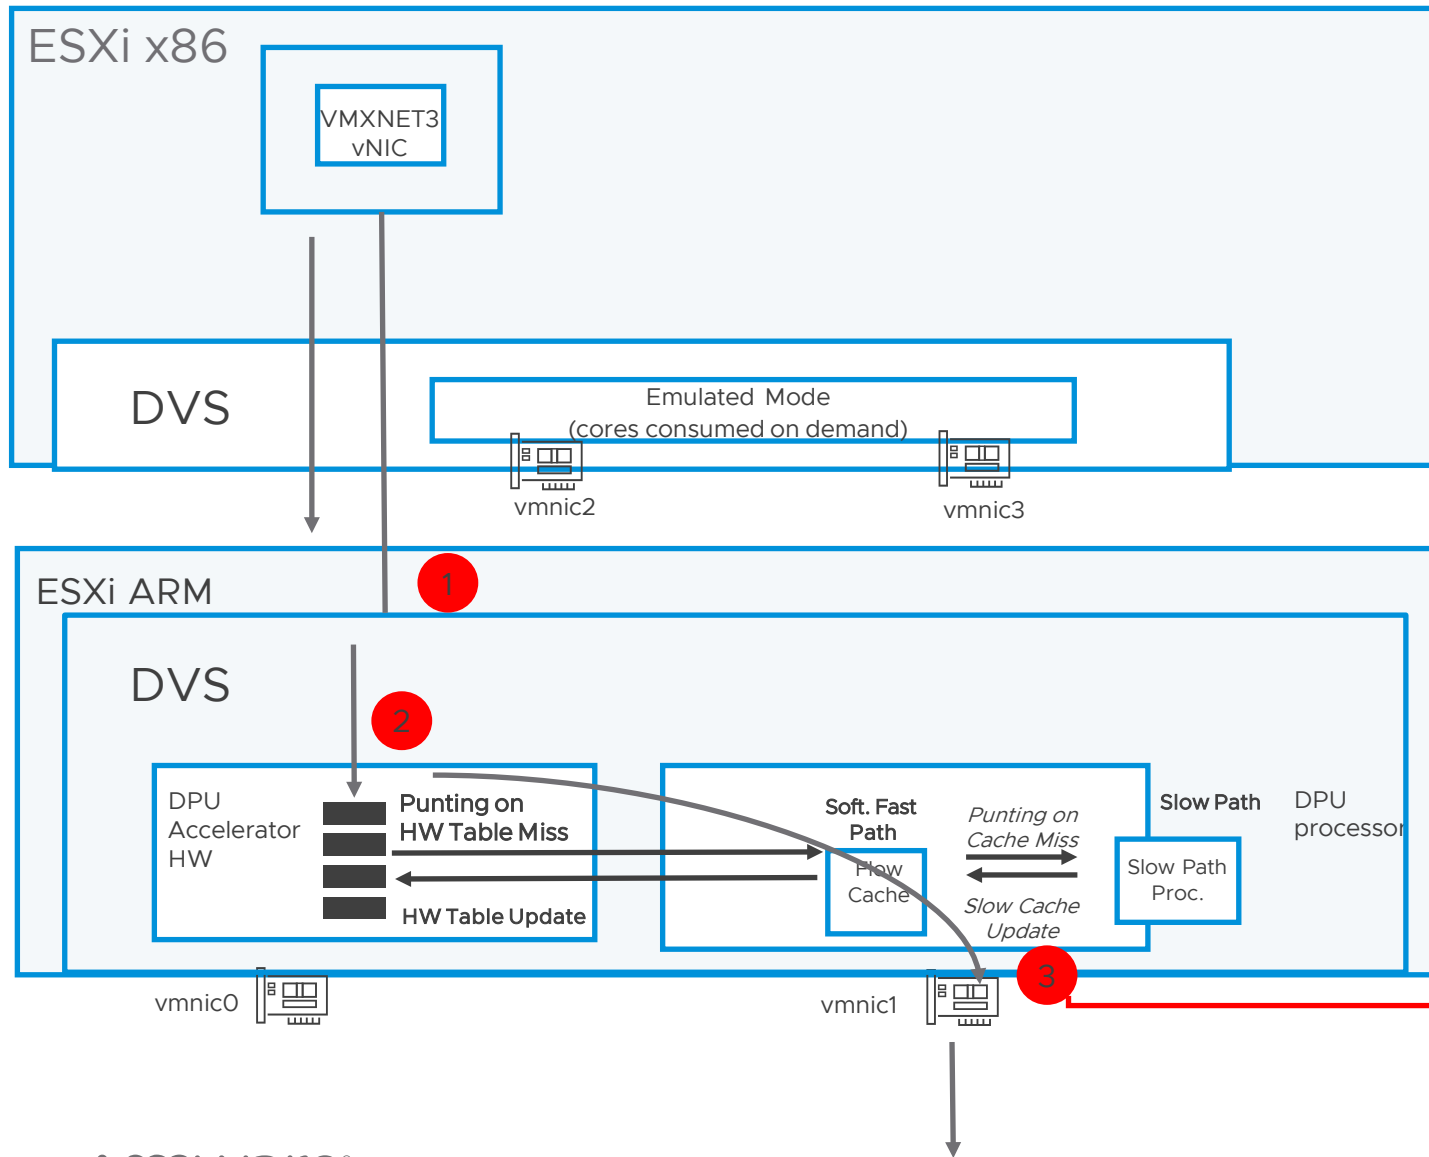
```
nsxdataplanecli ens flow-table dump
```

Flow Cache:

FT	dstMAC	srcMAC	VLAN	srcPort	dstIP	srcIP	proto	VNI	srcPort/type	dstPort/code	Actions	hwHits
L4	00:50:56:a1:87:fe	00:50:56:a1:fc:f3	0	3	172.16.31.26	172.16.31.19	TCP	74752	60186	0	bmap:0x80800080	6
inval(s):95 cg:2115 dp:0xa len:704; GENEVE DECAP; DFW on dstPort; 0 0												
L4	00:50:56:a1:87:fe	00:50:56:a1:87:fe	0	10	172.16.31.19	172.16.31.19	TCP	22	0	0	bmap:0x20000c0	0
cg:2115 dp:0x3 len:814; DFW on srcPort; VNI: 74752; GENEVE ENCAP VNI: 74752; 0 0												

Packet Walk in UPTv2 Mode

Follow the RED Dot



Port Info:

```
[root@localhost:~] net-stats -l
PortNum          Type SubType SwitchName      MACAddress
ClientName
2248146967       4      0 DvsPortset-1
08:c0:eb:5d:ae:0f vmnic1
```

Command:

```
pktcap-uw --uplink vmnic1 --dir 2 -o - | tcpdump-uw -enr -
```

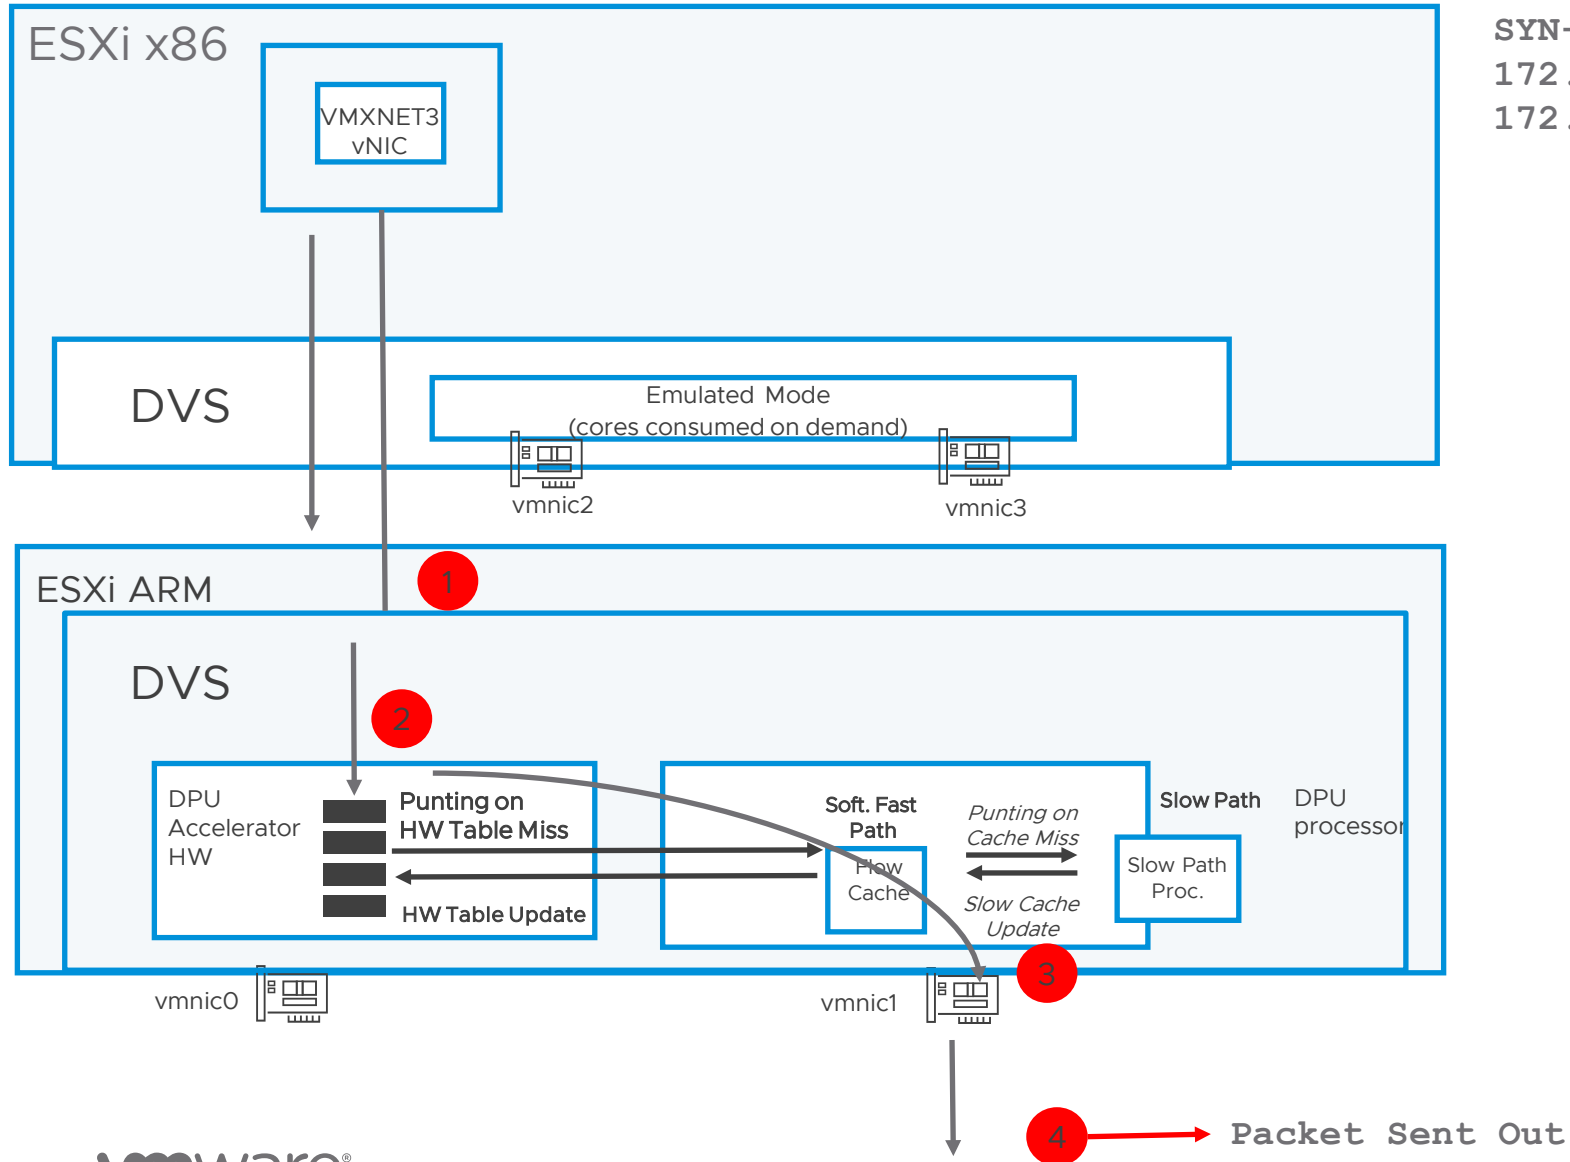
Packet:

```
00:50:56:67:eb:13 > 00:50:56:63:04:5c, ethertype IPv4
(0x0800), length 132: 192.168.4.60.55616 >
192.168.4.60:8080 : Geneve, Flags [none], vni 0x12400, proto
TEB (0x6558), options [8 bytes]: 00:50:56:a1:87:fe >
00:50:56:a1:87:fe, ethertype IPv4 (0x0800), length 74:
172.16.31.19.22 > 172.16.31.25:8080: Flags [S.], seq
3759600490, ack 402077946, win 65160, options [mss
1460,sackOK,TS val 3501164822 ecr 279309001,nop,wscale 7],
length 0
```

Packet Walk in UPTv2 Mode

Follow the RED Dot

SYN-ACK packet sent from source VM
172.16.31.19 SSH to destination VM
172.16.31.26



UPTv2 vs Emulated mode

vSphere and NSX Features Preserved – DPU based
Recommended with EDP – Standard

UPTv2 (Uniform Passthrough)

- Near zero CPU Consumption
- Complete Guest Memory Reservation
- Requires VMXNET3 Version 7 on GuestOS
- High Performance, Low Latency Data Path
- UPTv2 VMs scale varies with VFs resources

Emulated (Default Mode)

- Higher CPU footprint than UPTv2
- No Guest Memory Reservation
- No driver dependencies on GuestOS
- Performance accelerated by HW in DPU
- Higher Workload Density

Thank You

