

Cisco VM-FEX Best Practices for VMware ESX Environment

Deployment Guide

December 2011



Contents

1	Executive Summary.	3
	1.1 Target Audience 1.2 Introduction	. 3
2	Cisco UCS VM-FEX Best Practices	. 4
	2.1 Scale Considerations for static vNICs/vHBAs and Dynamic vNICs	. 4
	2.2 Defining Dynamic vNICs polices	. 9
	2.3 Service Profile creation for Full / Half Width blades	10
	2.4 VM-FEX with VMDirectPath (UPT)	11
	2.5 VMDirectPath Sizing	13
<u>3</u>	References	18

1 Executive Summary

1.1 Target Audience

The target audience for this guide includes, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy Cisco VM-FEX for VMware ESX Environment.

1.2 Introduction

Cisco Virtual Machine Fabric Extender (VM-FEX) is a Cisco technology that addresses management and performance concerns in a data center by unifying physical and virtual switch management. The Cisco VM-FEX collapses virtual and physical networking into a single infrastructure. This unified infrastructure enables data center administrators to provision, configure, manage, monitor, and diagnose virtual machine network traffic and bare metal network traffic.

The Cisco VM-FEX significantly reduces the number of network management points, enabling both physical and virtual network traffic to be treated in a consistent policy driven manner.

The VM-FEX software extends Cisco fabric extender technology to the virtual machine with the following capabilities:

- · Each virtual machine includes a dedicated interface on the parent switch
- All virtual machine traffic is sent directly to the dedicated interface on the switch
- · The software-based switch in the hypervisor is eliminated

Figure 1. Extension of Fabric Extender Technology with Fabric Interconnects Using VM-FEX



Figure 1 shows the extension of Fabric Extender technology with Fabric interconnects using VM-FEX. The Cisco Virtual Machine Fabric Extender (VM-FEX) technology extends Cisco Fabric Extender technology all the way to the virtual machine. Each virtual machine gets a dedicated interface on the parent switch (virtual Ethernet port). All virtual machine traffic is sent directly to the dedicated interface on the switch. VM-FEX eliminates the software based switch within the hypervisor by providing individual virtual machine virtual ports on the physical network switch. Virtual machine traffic is sent

directly to the upstream physical network switch, which takes full responsibility for virtual machine switching and policy enforcement. This approach leads to consistent treatment for all network traffic, virtual or physical. VM-FEX consolidates virtual and physical switching layers into a single layer and reduces the number of network management points by an order of magnitude. The following are the benefits of VM-FEX:

- Simplicity
 - One infrastructure for virtual and physical resource provisioning, management, monitoring and troubleshooting
 - · Consistent features, performance and management for virtual and physical infrastructure
- Robustness
 - Programmability, ability to re-number VLANs without disruptive changes
 - Trouble shooting & Traffic engineering VM traffic from the physical network
- Performance
 - VMDirectPath with vMotion provides near bare metal I/O performance
 - Line rate traffic to the virtual machine

2 Cisco UCS VM-FEX Best Practices

This guide covers the following topics along with the recommended settings for best performance and a trouble free environment.

- Scale considerations for static vNIC/vHBA and Dynamic vNICs
- Defining Dynamic vNICs polices
- Service Profile creation for Full / Half Width blades
- VM-FEX VMDirectPath (UPT) mode
- VMDirectPath Sizing

2.1 Scale Considerations for static vNICs/vHBAs and Dynamic vNICs

2.1.1 Discovery Policy

Chassis Discovery Policy defines the number of Fabric Extender (IOM) port links available for all the blades in Chassis. These port links are used to access the Fabric Interconnect for all operations.

The maximum number of Virtual Interfaces (VIF) that can be defined on Cisco VIC Adapters depends on the maximum number of port links available on the Fabric Extender. In a single chassis environment, the number of port links is determined by the chassis discovery policy. However, in environments where there is more than two chassis managed by a single Fabric Interconnect, the number of available port links can vary with the availability of ports on the Fabric Interconnect

Figure 2 shows the method to define Discovery Policy.

radic Summary	📑 Ġ 🍩 🚨 New 👻 🙀 Options 🛛 🚱 🕕 📥 Pending Activities 🛛 🙆 Exit.	cisco
	>> 🛱 Equipment	Se Equipment
2 43 8	🛱 Main Topology View 💷 Fabric Interconnects 🥪 Servers 🧹 Thermal 🎕 Decommissioned 🚔 Firmware Management 📅 Policies 💦 Faults	
Equipment Servers LAN SAN VM Admin	Global Policies Autoconfig Policies Server Inheritance Policies Blade Server Discovery Policies SEL Policy Power Groups	1
Filter: All Classis Classis Classis Structure Structure Structure Structure Structure Rack-Mounts Structure FX Structure Fairs Structure Structure Structure Filter: Fairs	Chassis Discovery Folky Action: Units Rack Server Discovery Link Action: Immediate. Liga: Advanueleged Scrub Policy: end: set> • Power Policy Redundancy: Non Redundant. n+1 Grid MAC Address Table Aging Aging Time: Never Mode Default other Clobal Power Allocation Policy Alocation Method: Manual Blade Level Cap Policy Driven Chassis Group Cap	
		and the second se

Figure 2. Chassis Discovery Policy Setting Window

Table 1 provides the scaling information on the number of VIFs that can be created with the following considerations:

- With or without Jumbo frames
- A combination of Fabric Interconnects (6100 / 6200) and Fabric Extenders (2104 / 2208)
- · Number of static and dynamic vNICs and vHBAs on Cisco VIC Adapters
- ESX version 4.1 and above

Fabric Interconnec t	Fabric Extender s	Chassis UPLINK S	Num VIC Car d	Num Static vNICs	Num vHBA s		Maximu m Dynamic vNICs		Total VIFs (Static Dynam c) vNICs FC vHBAs	-/ 2 ni 2 / 1 5 1	ESX I.0 J3 – I.1 – I.0i J3 – I.1iU Max /IFs	ESX 5 Max VIFs	5.0
6100	2104	1	1	2	2	9		13	3		58	58	
										MTU 9000	4	MTU 9000	58

 Table 1.
 Scaling Information of Virtual Interfaces

Fabric Interconnec t	Fabric Extender s	Chassis UPLINK S	Num VIC Car d	Num Static vNICs	Num vHBA s	Maximu m Dynam vNICs	I Total VIFs C (Stati Dynau c) VNICs FC VHBA	c / 4 mi 4 s / 4 s 4	ESX 4.0 J3 – 4.1 – 4.0i J3 – 4.1iU I Max /IFs	ESX Max VIFs	5.0
								MTU 1500	54	MTU 1500	
6100	2104	1	2	4	2	7	13	5	58	58	
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	
6100	2104	2	1	2	2	24	28	5	58	58	3
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	
6100	2104	2	2	4	2	22	28	5	58	11	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6100	2104	4	1	2	2	54	58	5	58	58	3
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	
6100	2104	4	2	4	2	52	58	5	58	11	6
								MTU 9000	4	MTU 9000	58

Fabric Interconnec t	Fabric Extender s	Chassis UPLINK S	Num VIC Car d	Num Static vNICs	Num vHBA s	Maximu m Dynamic vNICs	Total VIFs (Static Dynar c) vNICs FC vHBA	; / 2 ni 2 / 2 s 1	ESX I.0 J3 – I.1 – I.0 J3 – I.1 Max /IFs	ESX Max VIFs	5.0
								MTU 1500	54	MTU 1500	58
6100	2208	8	1	2	2	112	116	5	8	110	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6100	2208	8	2	4	2	110	116	5	8	11	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2104	1	1	2	2	57	61	5	68	11	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2104	1	2	4	2	54	61	5	8		116
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2104	2	1	2	2	112	116	5	8	11	6
								MTU 9000	4	MTU 9000	58

Fabric Interconnec t	Fabric Extender s	Chassis UPLINK S	Num VIC Car d	Num Static vNICs	Num vHBA s	Maximu m Dynamic vNICs	Total VIFs (Static Dynar c) vNICs FC vHBA	:/ 4 ni 4 / 4 s 1	ESX 1.0 J3 – 1.1 – 1.0i J3 – 1.1iU Max /IFs	ESX Max VIFs	5.0
								MTU 1500	54	MTU 1500	58
6200	2104	2	2	4	2	110	116	5	8	110	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2104	4	1	2	2	112	116	5	8	11	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2104	4	2	4	2	110	116	5	8	11	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2208	1	1	2	2	57	61	5	8	11	6
								MTU 9000	4	MTU 9000	58
								MTU 1500	54	MTU 1500	58
6200	2280	1	2	4	2	55	61	5	8	11	6
								MTU 9000	4	MTU 9000	58

Fabric Interconnec t	Fabric Extender s	Chassis UPLINK S	Num VIC Car d	Num Static vNICs	Num vHBA s	Maximu m Dynamic vNICs	Total VIFs (Static Dynan c) vNICs FC vHBAs	:/ 4 ni 4 / 4 5 1	ESX 4.0 J3 – 4.1 – 4.0i J3 – 4.1iU I Max /IFs	ESX S Max VIFs	5.0
									54		58
								MTU		MTU	
								1500		1500	

Note: These numbers are applicable to Cisco UCS 2.0 and ESXi 5.0 and they are consistent with Cisco UCSM 2.0 configuration limits. For more information on the configuration limits see : http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/configuration_limits/2.0/b_UCS_Configuration_limits_2_0/b_UCS_Configuration_limits_2_0.html. These numbers have been tested with Cisco UCSM 2.0 and ESXi 5.0, however they may change in future versions.

2.2 Defining Dynamic vNICs polices

It is recommended to choose the VMWarePassThrou adapter policy when defining Dynamic Ethernet vNIC policy for the Cisco VIC Adapter in the Service Profile.

The VMWarePassThrou adapter policy parameters are derived from testing performed under various workloads and are tuned for the Cisco UCS system to provide better performance on ESX with the VM-FEX environment.

Cisco UCS applies the VMWarePassThrou adapter policy parameters to Cisco VIC adapters during Service Profile association. Later, when the ESX OS is loaded on the blade associated with the Service Profile, the Ethernet driver of Cisco VIC automatically sets these values in the OS. Manual configuration and tuning is not required on the host.

You can choose to manually tune these parameters depending on your current requirements. However, improper tuning of the parameters can impact performance, cause OS instability and affect the maximum number of VIC Interfaces (Static / Dynamic vNICs).

- Configure VMXNET3 driver resource in the Guest OSs in terms of number of interrupts and number of queues (WQs, RQs, CQs).
- The values set for WQs, RQs, CQs and Interrupts on the Dynamic vNIC Adapter Policy must match the values for these queues and interrupts in VMXNET3 driver.
- Consider the resource configuration of all the VMXNET3 instances of the VMs destined to run on a server.
- Configure appropriately the adapter policy within the dynamic vNIC policy in the ESX Host; in
 order for a VMXNET3 instance to go into High Performance mode, the adapter policy must have
 enough queue and interrupts resources to satisfy the VMXNET3 instance configuration.

The Figure 3 shows the method to define Dynamic Adapter Policy.



Cisco Unified Computing System Manager	j2-151-h07		- 2
Fault Summary	🕝 🔘 🖬 New - 🔛 Options 🛛 🚱 🕕 📥 Pending Activities 🗌 🔟 Exit		alı ci
2 43 8 1	>> 🦏 Servers + 📓 Policies + 👗 root + 📓 Adapter Policies + 🍯 Eth Adapter Policy V	VMWarePassThru 🖉 Eth Adapter Polic	cy VMWarePassTh
Fourment Servers LAN SAN VM Admin	General Events		
Filter: Al	Actions Properties		
+	Delete Name: YMWarePassThru		
	Show Policy Usage Description: Recommended adapter settings for VMWa	are pass-thru (dynamic VNic)	
(*) ■ 823972-Ch03-8108-8551 (*) ■ 50,UEST=820-0-M2-8.08 (*) ■	Resources Transmik Queues: 1 Ring Size: 256 Ring Size: 1 Receive Queues: 1 Ring Size: 512 Ring Size: 512 Completion Queues: 1 Interrupts: 1		
Constant Sprice Template SP-template-with-f A Sub-Organizations Solicies A root S Adapter Policies	Options Transmit Checksum Offload: Oisabled O Enabled Receive Checksum Offload: Oisabled O Enabled		
SS Eth Adapter Policy Linux SEth Adapter Policy VMWare	TCP Segmentation Offload: O Disabled ③ Enabled		
Eth Adapter Policy Windows	TCP Large Receive Offload: O Disabled Enabled		
Setti Adapter Policy default	Receive Side Scaling (RSS): O Disabled Enabled		
	Faiback Timeout (Seconds): 5 [0-600]		
- S FC Adapter Policy default	Interrupt Made: O Msi X O Msi O Intx		
	1.1	<u>×</u>	
Boot Policies Ger und rinnen Doducer		Save Changes	Reset Values
3			

2.3 Service Profile creation for Full / Half Width blades

2.3.1 Full Width Blade with Dual Cisco VIC Adapters

It is the best practice to create Vcon placement policy on Full width blades installed with dual Cisco VIC Adapter and have Static and Dynamic vNICs defined in the Service Profile.

To create the Vcon placement policy, you need to create vNIC/vHBA Placement Policy and place 2 static vNICs on Fabric Interconnects on Vcon 1 & 2 with fabric failure option enabled from Cisco UCS 2.0 onwards. Then, you need to create Dynamic vNIC policy with the number of vNICs defined and define Adapter policy (refer to section 2.2 on page 7). By default protected option A-B is selected, which is also the recommended option.

Note: Ensure not to enable fabric failover on static vNICs before applying Vcon placement policy with Cisco UCS versions 1.4 and below.

The following are the benefits of adopting the configuration best practices:

- Cisco UCS will automatically place the entire Dynamic vNICs evenly across the Physical Cisco VIC adapters. This provides a better network load balancing by allowing Dynamic vNICs to pin to specific Fabric Interconnect LAN Uplink ports.
- By placing Dynamic vNICs evenly on dual Cisco VIC adapters you can effectively utilize 10G ports on both the adapters to perform network related operations.

 When there is a failure along the Fabric Interconnect path including the LAN Ethernet Uplink Interfaces, protected mode enables all Dynamic vNICs to fail over to other Fabric Interconnect.

2.3.2 Half Width Blade with Single Cisco VIC Adapter

Half width blade with single Cisco VIC adapter does not need any Vcon placement policy to be defined in the Service Profile since two Static vNIC / vHBAs and Dynamic vNICs are placed on a single adapter by default.

In the VM-FEX Port Profile, you need to configure the Dynamic Ethernet interface's actual VLAN, QOS, pinning etc. Unlike static vNICs, Dynamic vNICs are not configured upfront.

2.4 VM-FEX with VMDirectPath (UPT)

VM-FEX VMDirectPath is a feature where the network IO bypasses the Hypervisor's network kernel and communicates directly with the Cisco VIC adapter. This helps in offloading the host CPU and Memory cycles to handle VM's networks.

The conditions that need to be met for the Virtual Machines on VM-FEX to use the VMDirectPath are as discussed the following sections.

2.4.1 High Performance Option

VMDirectPath with vMotion (supported with EX 5.0) on VM-FEX is the high performance mode of VM-FEX. You can enable this mode by selecting the High Performance radio button in the port profile as shown in Figure 4.

Fault Summary	🗄 🕝 🔘 🛄 New 👻 📝 Optic	ions 😢 🕕 🕍 Pending Activities 🔟 Exit
♥ ▲ ▲ 0 10 7 2	>> Al + 🚍 Port Profiles + 🚍	Port Profile Mgmt
Equipment Servers LAN SAN VM Admin	General VM LAN Profile Client	tts Virtual Machines Events
Filter: Al 💌	Actions	Properties
	Create Prohie Client	Name: Mgmt
	-I Modify VLANs	Description: Mangment
B-GAN Clusters	🕆 Delete	QoS Policy: Platinum 👻
Port Profiles		Network Control Policy: default
Port Profile Data-1 Port Profile Data-2		Max Ports: 64
Port Profile Mgmt Port Profile Ymotion		Host Network IO Performance: O None O High Performance
Port Profile vNic-811		Pin Group: <not set=""></not>
Generation Generation Generation Generation		

Figure 4. High Performance Mode Setting in the Port Profile Management Window

Figure 5 shows the method to apply port profile for the Virtual Machine in vCenter Manger. After port profile is applied Directpath I/O Gen2 status automatically goes to active state. The Figure 6 shows the DirectPath I/O Gen2 status is Active after the profile is applied.

2 vo	enter.UCSAD	Lcom - vSphere Client					_ @ ×
File	Edit View 1	inventory Administration	Plug-ins Help				
		Home 🕨 👸 Invento	ry 🕨 🎁 Hosts and Ousters			Search Inventory	۹.
		🚗 l 🚓 🛛 🚓 😽	alenter and		-		
_	🕜 Witness	-Server2 - Virtual Mac	hine Properties		3		
18 6	Hardware	Options Resources Pro	ofiles VServices	Virtual Machine Version: 8			
	E Show	All Devices	Add Remove	Connected	ms Console Permissions Maps Sorage Vi slose tab	two Update Manager	-
	Hardware		Summary	If Connect at power on			
	Meno	ory.	32768 MB	Adapter Type			
	Video	• a.card	Video card	Current adapter: VMDNET 3	Virtual Machines		
	G 174CE	I device	Restricted				
	🧿 ५८५१	controller 0	LSI Logic SAS	Power Appress			
	Bard Hard	disk 1	Virtual Disk.	process and a second	uster		
	CD/D	WD drive 1	Client Device	G Automatic C Manual			
	Propp	vy drive 1 vyk adaptar 3	Manet (DTC-W"1) (Deate	DirectPath I/O Gen. 2			
	-	na na analagonan la	regime of the transport of const	Status: Inactive 0	Host		
				To activate DirectPath I/O Gen. 2, go to the Resources tab			
				and select Memory Settings to reserve all guest memory.			
				Network Connection	Datacenter		
				Network label:	p vCenter Server		
				Mgmt (PTS-VC1)	Client.		
				VM Network			
				Data-2 (PTS-VC1)			
				Added as PESAVC1 (PTS-VC1)	e Further		
				(motion(HPmcI))			
					in more about virtual machines		
					in how to install an operating system		
							-
Trigge					Object	or Name contains: -	Clear X
Objec	Male	-		or court	dead for		
~~~/~~					and the second s		
				4	6		
- 🛃 T	asks 👷 Al	arms			License Period 108 days remaining	To release oursor, press CTRL+ALT	Administrator

#### Figure 5. Adding Port Profile for Virtual Machines

Figure 6. DirectPath I/O Gen.2 Status

Perf-10 - Virtual Machine Pro	perties	×
Hardware   Options   Resources   R	Profiles vServices Add Remove	Virtual Machine Version: 8 Device Status Connected
Hardware Memory CPUs Video card VMCI device SCSI controller 0 Hard disk 1 CD/DVD drive 1 Network adapter 1 Floppy drive 1	Summary 1024 MB 1 Video card Restricted Paravirtual Virtual Disk [UPT-Datastore] rhel-se Mgint(PTS-VC) Client Device	Image: Connected         Image: Connect at power on         Adapter Type         Current adapter:       VMXNET 3         MAC Address         Image: Image
		Port: 1732 Switch to advanced settings

#### 2.4.2 Virtual Machine Configuration with VMDirectPath

The Virtual Machines using the VMDirectPath Dynamic vNICs should have memory reserved. The memory reserved must not exceed the physical available memory on ESX Host. The Adapter Type for VM vNICs using the VMDirectPath should be VMXNET3.

Figure 7 shows the method to enable memory reserve.

to D0	28.51 50	1	V Satesares All	scalles, Tur
	Power			
	Guest		30	fied Hat ent
	Snapehot	•	per	10001
Let .	Open Console		- manager -	1024 MB
10	Edit Setting Sci	_	Overhead	66-99 MB
09#	Migrate		Tools:	Notinstate
8	Cone		ssesi	
	Template		- iei	
	Fault Tolerance	<b>.</b> •	(e)	-N/4
	vH Storage Profile			Poweres of
	Add Permesson Ctri+i	ć i	apla:	172.35.178.
	Alarm		HA Protection:	ID NA C
	Report Performance		ada	
	Renone			
	Open in New Window Ctrl+Alt+R Remove from Inventory Delete Your Dos	K.	a Settinge vate	

Setting the Memory Reserve

Figure 7.

You need to make sure that the virtual machine memory is 'reserve' and is configured with 1 GB. This is a mandatory configuration, as shown in Figure 8.

Figure 8. Resource Allocation Window to Reserve All Guest Memory

dware Options Resources Profiles VServices	Virtual Machine Version:
dware [ Opbons Resources ] Summary U 0 MHz emory 0 MB example of the second of the	

# 2.5 VMDirectPath Sizing

#### 2.5.1 Windows Guest Virtual Machines with RSS Enabled

The following formula gives the maximum number of Dynamic vNIC with VMDirectPath active mode when all the Virtual Machines run Windows Guest with VMXNET3 emulated driver and have RSS is enabled.

4 * (Static vNICs) + 4 *(vHBAs) + (Num of Dynamic vNICs) + (Num of Dynamic vNICs with VMDirectPath * (Max(Num TQs, Num RQs) + 2) = < 128

SL NO	Guest OS	RSS Enabled	Static vNICs	Static vHBAs	Max(Num TQs, Num RQs)	Dynamic vNICs	Max VMDirectPath Dynamic vNICs
1	Windows	Yes	2	2	2	22	22
2	Windows	Yes	2	2	4	16	16
3	Windows	Yes	2	2	8	10	10
4	Windows	No	2	0	2	24	24
5	Windows	No	2	0	4	17	17
6	Windows	No	2	0	8	10	10

# Table 2. Scaling of Dynamic vNIC with VMDirectPath, Virtual Machines Running on Windows Guest with VMXNET3 Emulated driver and RSS Enabled

#### 2.5.2 Windows Guest Virtual Machines with RSS Disabled

The following formula gives the maximum number of Dynamic vNIC with VMDirectPath active mode when all the Virtual Machines run Windows Guest with VMXNET3 emulated driver and RSS is disabled.

4 * (Static vNICs) + 4 *(vHBAs) + (Num of Dynamic vNICs) + (Num of Dynamic vNICs with VMDirectPath * 2) = < 128

 
 Table 3.
 Scaling of Dynamic vNIC with VMDirectPath, Virtual Machines Running on Windows Guest with VMXNET3 Emulated Driver and RSS Disabled

SL NO	Guest OS	RSS Enabled	Static vNICs	Static vHBAs	Dynamic vNICs	Max VMDirectPath Dynamic vNICs
1	Windows	No	2	2	37	37
2	Windows	No	2	0	40	40

#### 2.5.3 Linux Guest Virtual Machines with Multi Queues Enabled

The following formula gives the maximum number of Dynamic vNIC with VMDirectPath active mode when all Virtual Machines run Linux Guest with VMXNET3 emulated driver and multi-queue is enabled.

4 * (Static vNICs) + 4 *(vHBAs) + (Num of Dynamic vNICs) + (Num of Dynamic vNICs with VMDirectPath) * (Num of queue pairs + 1) = < 128

 Table 4.
 Scaling of Dynamic vNIC with VMDirectPath, Virtual Machines Running on Linux Guest with VMXNET3 Emulated

 Driver and Multi-Queue Enabled
 Scaling of Dynamic vNIC with VMDirectPath, Virtual Machines Running on Linux Guest with VMXNET3 Emulated

SL NO	Guest OS	Multi-Queue Enabled	Static vNICs	Static vHBAs	Number Of	Dynamic	Max VMDirectPath
-------	----------	------------------------	--------------	-----------------	-----------	---------	---------------------

					Queue Pairs	vNICs	Dynamic vNICs
1	Linux	Yes	2	2	1	37	37
2	Linux	Yes	2	2	2	28	28
3	Linux	Yes	2	2	4	18	18
4	Linux	Yes	2	2	8	11	11
5	Linux	Yes	2	0	1	40	40
6	Linux	Yes	2	0	2	30	30
7	Linux	Yes	2	0	4	20	20
8	Linux	Yes	2	0	8	12	12

#### 2.5.4 Linux Guest Virtual Machines with Multi Queues Disabled

The following formula determines the maximum number of Dynamic vNIC with VMDirectPath active mode when all Virtual Machines run Linux Guest with VMXNET3 emulated driver and multi-queue is disabled.

4 * (Static vNICs) + 4 *(vHBAs) + (Num of Dynamic vNICs) + (Num of Dynamic vNICs with VMDirectPath + 1) = < 128

 
 Table 5.
 Scaling of Dynamic vNIC with VMDirectPath, Virtual Machines Running on Linux Guest with VMXNET3 Emulated Driver and Multi-Queue Disabled

SL NO	Guest OS	Multi-Queue Disabled	Static vNICs	Static vHBAs	Number Of Queue Pairs	Dynamic vNICs	Max VMDirectPat h Dynamic vNICs
1	Linux	Yes	2	2	1	56	56
2	Linux	Yes	2	0	1	60	60

#### 2.5.5 Linux Guest Virtual Machines with Advanced Configuration

In a Linux Guest Virtual Machine, you can change the interrupt mode of the vmxnet3 from MSI-X to MSI. The MSI-X interrupt mode which is the default interrupt mode when VMDirectPath is enabled, consumes four interrupts. In comparison, the MSI interrupt mode consumes only one interrupt, but causes the multi-queue feature to be turned off.

The following formula gives the maximum number of Dynamic vNIC with VMDirectPath mode active when all Virtual Machines run Linux Guest with VMXNET3 emulated driver and MSI Interrupt is turned on.

4 * (Static vNICs) + 4 *(vHBAs) + (Num of Dynamic vNICs) = < 128

You need to ensure that the Guest VMs are shutdown before enabling VMDirectPath and changing the interrupt mode from MSI-X to MSI. To change the interrupt mode from MSI-X to MSI, you need to edit the configuration file (.vmx). All supported UPT Linux Guests vNICs (RHEL 6, SLES 11, SLES 11 SP1) can be forced to use MSI interrupt by editing *.vmx file and inserting the following line.

```
ethernetX.intrMode = "2"
```

```
vmfs/volumes/4c9ffcdb-1d1c1488-a327-0025b500008d/Perf-1 # cat Perf-1.vmx | grep
ethernet
```

```
ethernet0.present = "true"
ethernet0.virtualDev = "vmxnet3"
ethernet0.dvs.switchId = "04 f6 05 50 6f 1f e4 cf-d3 1c 0d e8 f7 24 50 b8"
ethernet0.dvs.portId = "1710"
ethernetX.intrMode = "2"
ethernet0.dvs.portgroupId = "dvportgroup-223"
ethernet0.dvs.connectionId = "753846454"
ethernet0.addressType = "static"
ethernet0.address = "00:50:56:3F:0A:01"
ethernet0.pciSlotNumber = "192"
ethernet0.dvs.portId = "1710"
```

 Table 6.
 Scaling of Dynamic vNIC with VMDirectPath Mode Active, Virtual Machines Running on Linux Guest with VMXNET3 Emulated Driver and MSI Interrupt Turned On

SL NO	Guest OS	Interrupt Mode MSI	Static vNICs	Static vHBAs	Dynamic vNICs	Max VMDirectPath Dynamic vNICs
1	Linux	Yes	2	2	112	112
2	Linux	Yes	2	0	112	112

#### 2.5.6 Configuration

This section helps you to understand the method to enable or disable Receive Side Scaling (RSS) on a Windows. Figure 9 shows the method to enable / disable RSS on a Windows Guest Virtual Machine.

Figure 9. Enabling RSS in VMXNet3 Driver

vmxnet3 Ethernet Adapter #9 Propertie	25	×
General Advanced Driver Details Pow The following properties are available for thi the property you want to change on the left on the right.	wer Management s network adapter. , and then select its	Click value
Property: Jumbo Packet Large Rx Buffers MAC Address Max Tx Queues Maximum number of RSS Processor Offload IP Options Offload tagged traffic Offload TCP Options Priority / VLAN tag Receive Throttle RSS RSS Base Processor Number Rx Ring #1 Size Rx Ring #2 Size	Value: Disabled Disabled Enabled	
	ОК	Cancel

• Linux

NICs support multiple receive and transmit descriptor queues referred to as multi-queue. On the receive side different packets are sent to different queues to distribute processing among CPUs. The NIC distributes packets by applying filters to each packet. These packets are distributed in separate queues which are processed by separate CPUs. This mechanism is also known as Receive Side Scaling (RSS). The multi-queue feature is enabled by default in Linux with VMXNET3 driver version vmxnet3 1.016.0-k or higher on the guest virtual machine.

#### 2.5.7 VMDirectPath Guest Virtual Machine OS and VMXNET3 Supported Versions

Table 7 provides the VMXNET3 driver versions which support VMDirectPath mode for different Guest Virtual Machine Operating Systems.

SL NO	Guest OS	VMXNET3 Driver Version
1	Windows 2008 SP2	1.2.22.0
2	Windows 2008 R2	1.2.22.0
3	RHEL 6.0	1.0.14.0-k

 Table 7.
 VMXNET3 Driver Versions

SL NO	Guest OS	VMXNET3 Driver Version
4	SLES11 SP1	1.0.14.0-k
5	SLES11	1.0.36.0

# **3 References**

Hardware and Software Interoperability Matrix:

http://www.cisco.com/en/US/docs/unified computing/ucs/interoperability/matrix/r hcl B rel2 0.pdf

Cisco UCS Manager VM-FEX for VMware GUI Configuration :

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/b_GUI_V Mware_VM-FEX_UCSM_Configuration_Guide.pdf



Americas Headquarters Cisco Systems, Inc. San Jose, CA Asia Pacific Headquarters Cisco Systems (USA) Pte. Ltd. Singapore Europe Headquarters Cisco Systems International BV Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)