

Q: What is the overall purpose of VMready™?

A: VMready™ complements datacenter virtualization in performing two vital functions: first, it discovers Virtual Machines [VMs] that exist downstream of the downlink ports, giving visibility into the virtual environment in the physical server(s); it then tracks those VMs by their unique vMACs as they undergo VMotion events, potentially moving from one physical server to another, *regardless of which physical server or where in the VMready™ domain they end up.*

The VMs become the foundation for the grouping concept in VMready. This allows for much greater control over the allocation of network resources. Through VMready, the server administrator can assign network policies on a Group basis, as opposed to on a physical server basis. ACLs, QoS and VLAN assignments in VMready™ are done at the Group level, and when VMs move through some VMotion event, the ACLs, QoS and VLAN assignments move with them, ensuring that the appropriate network resources are allocated to the VMs, no matter where they reside in the network.

Q: What switches support VMready™?

A: The 10Gb switches are the ones that will initially support the VMready™ offering. A subset of VMready™ will be available for the 1Gb switches.

Q: Are there any restrictions on the switches that support VMready™?

A: Right now, the switches have to be 10Gb switches. In the next phase of VMready™, you will be able to expand the implementation to the other switches in the BLADE portfolio.

Q: What is the benefit of VMready?

A: VMready™ provides the ability to extend the capabilities and experience of the server administrators beyond the physical machine boundary. It provides for greater networking control, and greater network resiliency, by giving those administrators the ability to set individual VM network settings in the network Ethernet switches. Having been set once, those VM network settings can now move within the VMready™ domain in synchronization with the VM without further administrative actions, freeing the administrator for other tasks.

At the same time, VMready™ gives the network administrator greater visibility into the underlying structure of the virtual environment beyond the physical servers, making SLA configuration and monitoring easier and more effective.

VMready™ implies at least knowledge of a vendor's virtualization of the bare-metal server, but more an ability to identify individual VMs and then take action based upon them. In this release of VMready™ we will be able to identify specific VMs based upon their vNICs, and then take action based upon where those MACs are, and where they move to... Fundamentally, grouping becomes a function of the VMs, and not the physical downlink ports.

For example, we have a server blade with a dozen VMs running on it, coming through the vSwitch & then out the physical NIC. All those vNICs are propagated out the NIC, & we can learn them on the server port. Once we have learned those vMACs, then a number of different capabilities become possible: grouping by underlying VM; applying QoS by underlying VM; detecting server movement by virtue of VMotion and using that to trigger our NMotion, and; supporting mobility across chassis without having a dependency upon DHCP Option 82 (i.e.: Server Mobility, or MAC spoofing by chassis/slot assignment).

Q: How does VMready™ work?

A: VMready™ uses the Organizationally Unique Identifier [OUI] from the Media Access Control [MAC] address to identify VMs from the common virtualization vendors. It does this by detecting the MAC addresses appearing on the downlink ports, then ‘sniffing’ the OUIs of those MACs, to identify a MAC that is known to be associated with a VM from a registered vendor (with the IEEE, or IANA). The MAC is then put in a Port/VM table that carries the association of the VM with its original port. When the VM moves to a new port, it is discovered & the Port/VM table is updated to reflect the new location.

OUI (hex)	Vendor	Note:
00-0C-29	VMware	VM
00-50-56	VMware	VM
00-16-3E	XenSource	VM
00-03-FF	Microsoft	VM
00-0F-4B	Virtual Iron	VM
00-18-51	SWsoft	VM

Q: Does VMready™ use the ESX™ repository or any VMware information at all?

A: VMready™ does not access the ESX™ network repository data; it ‘learns’ the identities of the VMs by snooping the downlink ports on the switch. It then uses the vNIC to identify an entity that we can configure and manipulate *in the switch* for setting network characteristics. It is this vNIC that is used to track the VM throughout the VMready domain. The Q-in-Q mechanism is what is used to create, maintain and monitor groups.

VMready™ does not actively manipulate any of the VMware network resources (like the vSwitch). VMready™ manipulates the objects representing the VMs within the switches in a rack.

Q: What is a “VM Group” in VMready?


A: The group aggregates sets of destination stations representing a virtual machine (or machines, and their virtual NICs), to which a physical L2 Switch running VMready™ tries to forward an incoming packet.

A configuration can be defined by specifying 1) VM group number, which should be configured-VMready-wide unique, 2) a set of the MAC addresses representing relevant VMs on their vNICs and 3) an uplink port number or LAG on a physical L2 switch running VMready, which the relevant packets should traverse.

Packet forwarding based on VM Group definitions is achieved by maintaining and referencing the corresponding outer tag in the Q-in-Q technology adopted on a physical L2 switch.

VMready, consequently, guarantees the seamless and timely movement of a set of L2 network port attributes, aka, VLAN membership, ACL, and QoS from an origin downlink port to a destination downlink port on a physical L2 switch, in synchronization with the movement of a virtual machine or its vNIC station beyond a boundary of a physical server.

For example, in this diagram there are four ESX™ servers, each with two VMs. The color of the VM indicates its VM Port Group. The Content Servers [Red, ‘CS’] are placed in one VMready™ group; the Security servers [Blue, ‘IPS’ and ‘FW’] are placed in another, and; the Authentication / Authorization / Accounting servers [Green, ‘AAA’] are in the last group. As can be seen, for



each VMready™ group there are VMs associated with uplink trunk ports.

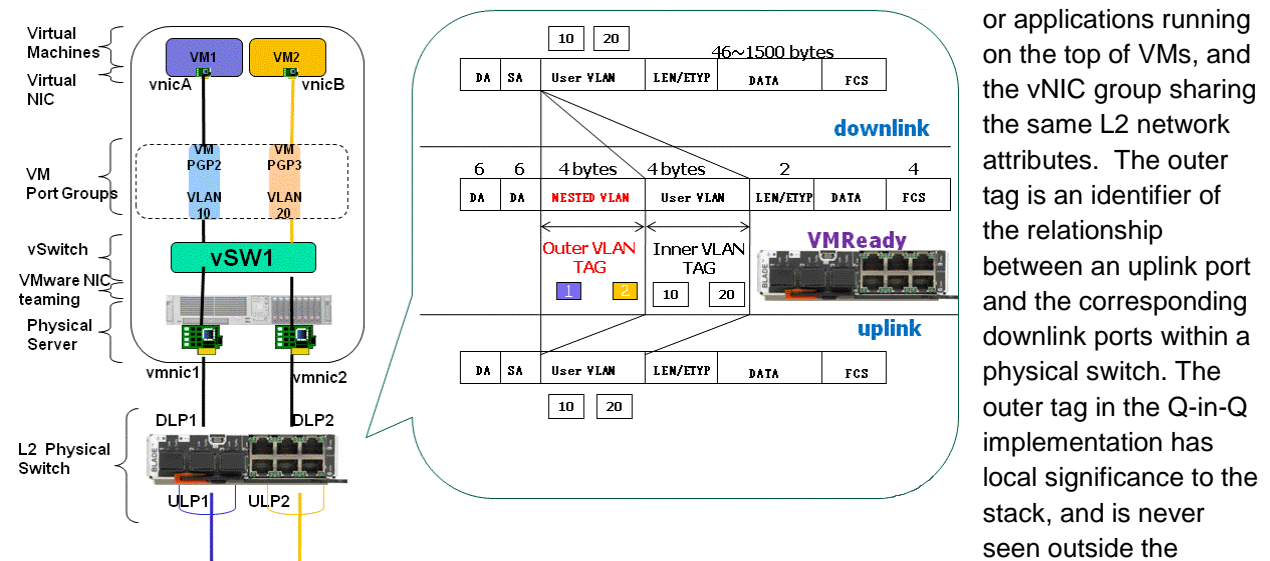
Q: What is the implementation of Q-in-Q in VMready?

A: The main purpose of Q-in-Q is to be transparent to any user VLANs. This allows the users to put their VMs in VLANs as per their needs (departmental separation, resource allocation etc.) without any additional configuration in the VMready™ capable physical switches.

Both a vNIC on a VM and a port group on a vSwitch can have a single VLAN definition. In this context a group in VMready™ should be able to synchronize itself with the users' VLAN definition. For example if the user configures a port group with VLAN 100 in the vSwitch then the corresponding port on the physical switch will have to be enabled for tagging so that the switch can accept packets with VLAN 100. But with Q-in-Q mechanism the user VLAN is transparent to the VMready™ and hence no configuration is needed on the physical switch.

Q: Can Q-in-Q be clarified any further?

A: The definition of a group in VMready™ is an association between the transaction path(s) for VMs



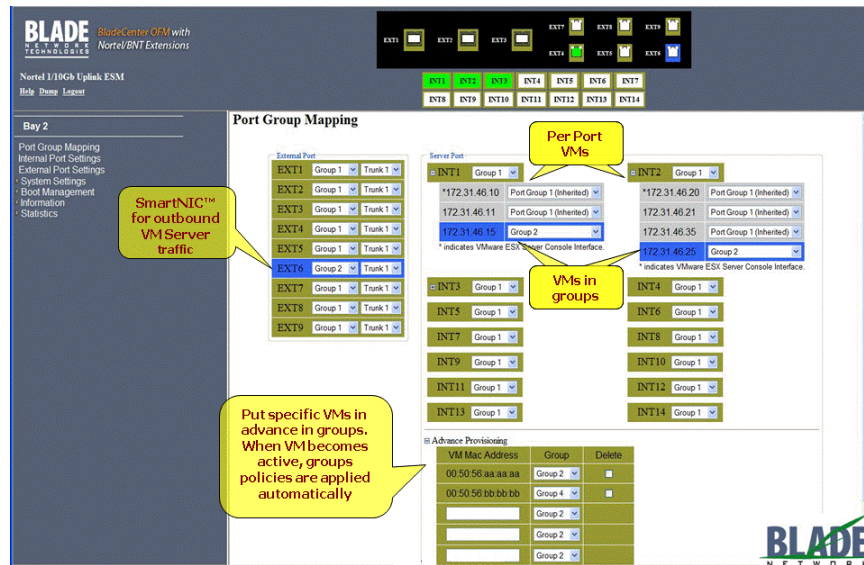
or applications running on the top of VMs, and the vNIC group sharing the same L2 network attributes. The outer tag is an identifier of the relationship between an uplink port and the corresponding downlink ports within a physical switch. The outer tag in the Q-in-Q implementation has local significance to the stack, and is never seen outside the

physical ports which belong to the group.

A VM group is the set of destination stations representing virtual machines or their virtual NICs, to which a physical L2 Switch running VMready™ tries to forward an incoming packet. In VMready, packet dispatching is not necessarily the same as a users' VLAN traffic path. Packet forwarding based on VM Group definitions is achieved by maintaining and referencing the corresponding outer tag in the Q-in-Q technology on the physical L2 switch that represents the underlying group.

Q: What is required for a VMready™ Configuration?

A: VMready™ has to take care of multiple groups on a single down link port on a BLADE switch by using a Virtual Machine's vNIC MAC address. VMready™ software learns the vNIC MAC address



automatically by snooping on the VM traffic. We present a list of vNIC MAC addresses and then let the user create a group by choosing the vNIC MACs from the list. For every vNIC MAC a unique index is also presented so that the user need not enter the whole MAC address but just a small index number.

However, they do not have to know their complete vNIC MAC address info first to proceed with the group definition with the complete MAC address on a vNIC. Even if they don't know the

exact MAC address info on a vNIC, users can define a group in VMready™ at a configuration step by pre-configuring the vMAC addresses resulting from the hypervisor configuration. VMready™ administrators generally know the vMACs from the configuration of the virtual machine environment, and by entering these addresses, when the virtual machines begin transmitting network traffic; VMready discovers the VM and associates the network policies appropriately.

Q: Does 'snooping' have to occur before a VMready™ group can be configured?

A: Users can get complete MAC addresses via the ESX™ Service Console and they can pre-configure VMready™ with the information if necessary for the time being. It's also possible to differentiate Service Console port groups and virtual machine port groups because Service Consoles generates heart beat packets in a certain manner.

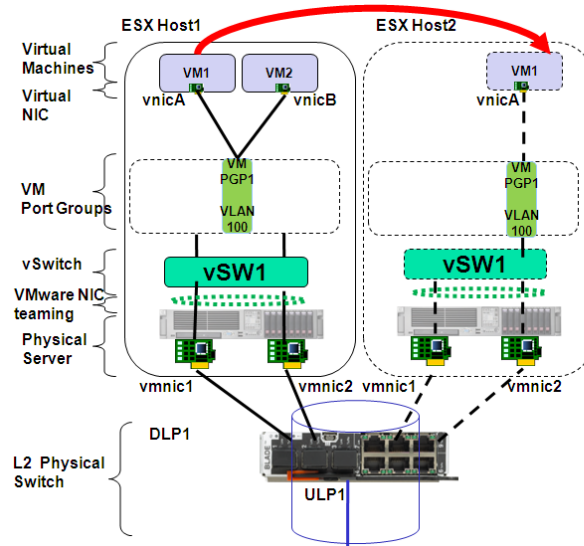
Q) When a new VM starts up on an ESX™ host due to VMotion or VMwareHA, does VMready™ configure vSwitches or VM port groups?

A) VMware's Virtual Center configures the VMs & vSwitches. VMready™ re-discovers the VM when it appears on the downlink port associated with the destination physical server of the VMotion. Generally, the hypervisor initiates a network connection on behalf of the VM just prior to the VM coming active. That way, when the VM is ready to begin transmitting traffic, the network switch is ready and has all the proper attributes in place.

Q: How does NMotion Discovery work?

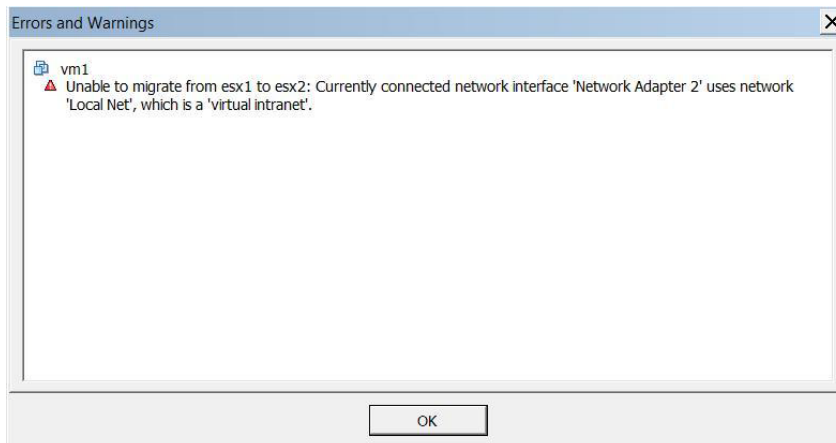
A: By capturing RARP (or some other network discovery traffic) originated on behalf of a destination VM by the hypervisor during a motion event, VMready™ completes moving the associated set of L2 network attributes (e.g.: ACL, QOS and VLAN information for a group) among downlink ports on the BLADE switch, in the synchronization with the movement of VM, before a VM or an application on a top of the VM starts communicating with their remote peer.

In general, a MAC table is maintained per VLAN identified by the 'outer' tag. A destination VM during an NMotion event causes MAC address learning by the VMready™ switch by causing the generation of some Ethernet administrative frame. VMready™ captures the frame (and the SMAC), parses the frame and discovers the vNIC, and moves the set of L2 network attributes to the new downlink port. The administrative traffic is generated on behalf of the moving VM by the hypervisor (in the case of VMware at around 90% VMotion complete) to do route setup prior to the VM sending traffic.



Q: What happens to a vSwitch with no External Links?

A: What has been configured is what VMware calls a "Virtual Intranet" (for example). An attempt to perform a "live" VMotion operation of that VM from one ESX™ host to the other failed will fail with an error message (please see the attached screenshot). In any event, the VMready™ solution would come into the picture only if one of the [downlink] switch ports is involved in the communication path.



Q: How many groups are supported by VMready™?

A: VMready™ supports up to 32 groups. It also supports two uplink trunks per group.

Q: What is the control protocol for VMready™ trunking?

A: VMready™ allows the use of both static and dynamic trunking. On the uplink trunks, the default control protocol is the Link Aggregation Control Protocol [LACP], but static control is an option. The downlink trunks use static control of members.

Q: What limitations are there for trunking in VMready™?

A: VMready™ does not allow for the mixture of 1Gb and 10Gb links in the same trunk. They must all be the same bandwidth. Further, a maximum of eight links can be combined to form a trunk (excluding stack trunks). All links must be full-duplex in the trunk. VMready™ requires that all downlink trunks have at least two members, with a maximum of eight.

Q: Is the Spanning Tree Protocol supported in VMready™?

A: No, none of the STPs [STP, MSTP, and RSTP] are supported in VMready™. The reason for this is that in the grouping concept, the underlying VMs and their uplink port(s) are all part of the same broadcast domain, and therefore there is no potential for a loop condition. It is not until the traffic exits the switch that there is a potential for an STP loop, and it is expected that the upstream aggregation layer prevent the loop condition. On the downstream side, hypervisors prohibit looping conditions in their virtual switch implementations, and so this is not a problem, either.

Q: How does BLADEHarmony Manager (BHM) support SmartConnect with VMready switches?

A: BHM provides additional management functionality beyond what is available for other BLADE switches.

Switch IP Address	VSG	VLAN	Switch#	Virtual MAC	VM IP	VM Name	Hypervisor
172.20.7.13	32		3	00:50:56:40:04:c8	172.16.20.1		
172.20.7.13	32		1 (Detached)	00:50:56:48:ee:56	172.16.40.1		
172.20.7.13	32		1 (Detached)	00:50:56:49:8c:c5	172.16.50.1		
172.20.7.13	32		2 (Detached)	00:50:56:4b:23:cb	0.0.0.0		
172.20.7.13	32		2 (Detached)	00:50:56:70:c8:21	172.16.50.2		
172.20.7.13	32		4	00:50:56:77:07:7b	172.16.20.2		
172.20.7.13	32		1 (Detached)	00:50:56:96:63:45	0.0.0.0		
172.20.7.13	None	10	3	00:50:56:96:05:99	172.16.251.4		
172.20.7.13	32		2 (Detached)	00:50:56:96:54:d8	0.0.0.0		
172.20.7.13	None		3	00:50:56:96:09:3c	0.0.0.0		
172.20.7.13	None		3	00:50:56:96:38:a8	0.0.0.0		
172.31.46.50	None		1	00:50:56:41:34:45	172.31.46.20		172.31.46.20 (Service Con
172.31.46.50	None		1	00:50:56:46:f7:4f	172.31.46.10		172.31.46.10 (Service Con
172.31.46.50	None		1	00:50:56:72:be:fb	172.31.46.21		172.31.46.20 (VMkernel)
172.31.46.50	None		1	00:50:56:76:ff:97	172.31.46.11		172.31.46.10 (VMkernel)
172.31.46.50	None		2	00:50:56:9c:02:4f	172.31.41.106	vm6	172.31.41.50
172.31.46.50	None		2	00:50:56:9c:19:58	172.31.41.101	50VM1	172.31.41.50

- 1) Centralized reporting of both the Port- and VM-based Virtual Switch Groups across all the switches in the network. The user can see, in one snapshot, the Virtual Switch Grouping configurations across the network.
- 2) Detailed drill-down per switch on the Virtual Switch Group configuration for the switch or stack of switches.
- 3) BHM optionally interfaces with the virtualization vendor's management application, such as VMWare's Virtual Center to provide additional information about the virtual machines that are discovered by the switches. This includes mapping the VM's vNIC MAC to the virtual machine name and the hypervisor in which it resides. The benefit of this is that the reports are in familiar terms to the user.

©2009 BLADE Network Technologies, Inc. All rights reserved. Information in this document is subject to change without notice. BLADE Network Technologies assumes no responsibility for any errors that may appear in this document. All statements regarding BLADE's future direction and intent are subject to change or withdrawal without notice, at BLADE's sole discretion. www.bladenetwork.net.

MKT090131

