



## Nexus 1000V Portfolio: Spring '11 Public Webinar Series: Q&A

**SESSION 4 (Business Track): *Journey to the Cloud with Nexus 1000V and Long Distance vMotion***

**Authors:** Nexus 1000V Product Management and Technical Marketing Engineering Team

This Q&A supports one session in a series of Webinars offered to our customers and partners in the spring of 2011. The session covered in this Q&A is highlighted in yellow below:

<b>Track</b>	<b>Date</b>	<b>Session Title</b>
Business	22-Mar	Nexus 1000V/1010 Overview and Update
Business	5-Apr	Virtual Network Services: Security (VSG), Appl. Acceleration (vWAAS), Monitoring (NAM)
Business	19-Apr	Virtual Security Gateway (VSG) Overview
Business	3-May	<b>Journey to the Cloud w/ N1KV: vCloud Director &amp; Long Distance vMotion</b>
Business	17-May	Secure Virtual Desktop with Nexus 1000V & VSG
Technical	29-Mar	Nexus 1000V v1.4 New Features and Installation/Upgrade Overview
Technical	12-Apr	Nexus 1010 Deployment & Best Practices
Technical	26-Apr	Virtual Security Gateway Installation & Basic Configuration
Technical	10-May	Nexus 1000V Advanced Configuration
Technical	24-May	Nexus 1000V Troubleshooting

The following questions, and corresponding answers, came from our 3-May event:

QUESTION	ANSWER
When using port-profiles to create network pools, can I reuse the same VLAN in more than one port-profile?	Since each port group must be layer 2 isolated, a vCloud a requirement, each port-profile must have a unique VLAN ID assigned to it. SO the answer is no.
Can I use more than 1 VSM as part of my vCloud Director deployment?	Yes you can. As long as that VSM is already part of your current vCenter or one of your vCenters under the vCloud inventory, you will be able to use its networking.
When you mentioned using ERSPAN across L3 boundaries, what might the configuration look like?	You would simply need to select the VLAN as the source when creating your monitor session, and then you can have a remote utility in a completely different network. You would still need to configure ERSPAN on each VEM, but that can all be done ahead of time. Some customer use host profiles to install the VEM with its associated vmk interfaces.
What is the Criteria to take the solution and turn it into a Cisco Validated Design ?	Our criteria in SASU includes no open severity 1 software defects, any severity 2 defects must have a defined alternative that allows the test case to execute and meet expectations, and all defined priority 1 and priority 2 test cases must execute and meet expectations. Further, the products must be supportable, that is the decision makers at Cisco and our partners will support our solutions and innovations.
Does the Virtualized Workload Mobility use case have a security component ?	Yes, we've added the Virtual Security Gateway as our host based security platform to isolate and protect virtual machines in single and multi-tenant environments. By the way, this new use case is Secure Virtualized Workload Mobility and not only will your port profile migrate, but so will your virtual machines security profile. It is also included in our CVD.
You have Layer 2 as the LAN extension using OTV (Overlay Transport Virtualization). Does the Nexus 1000v have this layer 2 requirement?	No it doesn't. The Nexus 1000v has control plane communication between the VSM and the VEM and it can communicated via Layer 2 OR a Layer 3 session. Our solution, Virtualized Workload Mobility, chose layer 3 for few reasons, but I particularly like L3 because it helps

	reduce layer 2 broadcasts from VSM to VEM and helps maintain VSM to VEM communication when your layer 2, aka OTV, session goes down for support, maintenance, migration, new deployments, etc.
--	--