



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

SESSION 3 (Business Track): **Virtual Security Gateway**

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:

Track	Date	Session Title
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	Journey to the Cloud w/ N1KV: vCloud Director & Long Distance vMotion
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 19-April event:

QUESTION	ANSWER
Do we need N1k to deploy VSG?	<ol style="list-style-type: none"> 1. Yes. You need N1k to deploy VSG. Nexus 1000v and VSG are tightly integrated, and that is what makes this a powerful product. Because of this tight integration, any security policy you define in your VSG is assigned to a port-profile in N1k, and any new VM with this port-profile is automatically provisioned with the right security policies without any manual intervention. 2. This also helps in separation of duties between the security administrator and the network administrator.
Do we need to install one VSG per each ESX-host?	<ol style="list-style-type: none"> a. No. Each VSG can support multiple VEMs or ESX-hosts. This is possible because of the distributed VSG architecture, and is the biggest differentiator for VSG. b. This decoupling has multiple advantages <ol style="list-style-type: none"> i. You don't want to overload your ESX-host for running switching, then firewall rules, then IPS, IDS, malware filtering, load-balancing etc. – which leaves very little computing power for what it is supposed to do – run applications. With the vPath architecture in N1k, you can offload all these network services to appliances that sit outside your server. ii. When you want to instantiate a firewall on each ESX-hosts, managing these separate instances becomes a night-mare, cumbersome and highly error-prone. It also creates single-points of failure. iii. There will be no clear role separation between the server admin and the security admin. Each time a security policy has to change, the server admin has to change on each server. c. This also allows for a highly scalable and highly available design. If you have more VMs or if you need more availability, all you need to do – is instantiate more VSGs.
What do you mean by attribute-based firewall rules?	<ul style="list-style-type: none"> • We have a very sophisticated policy engine as part of VSG, and unlike traditional firewall rules, with VSG - users can define their own attributes and write policies based on these attributes. For example, you can create a logical zone for all your test VMs with one simple rule, and use this zone as one of the attributes in your security rules. • This attribute-based rule-engine is very important and required for the dynamic nature of virtualized environments where VMs can be instantiated with the click of a button, and where these VMs can move from one physical machine to another pretty easily.
How many VMs can each VSG support? How many	Like mentioned earlier, VSG employes a scale-out architecture,

attributes can we define?

and is highly scalable. From firewall policy rules and scalability, we have another session next week where we are going to provide detailed information around policy creation and the associated limits.