All Your Packets Are Belong to &hellip; You?

Lori MacVittie, 2012-06-08

Yes, even the ones over there, in that there cloud, can be yours.

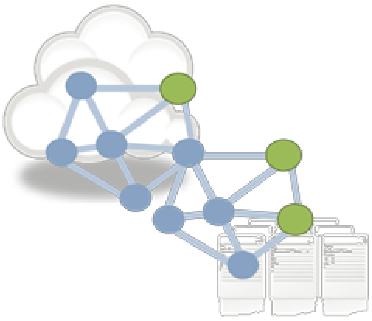

No one argues that networks have not exploded in terms of speeds and feeds in the past decade. What with more consumers (and cows), more companies going "online", and more content it'd be hard to argue that there's less traffic out there today than there was even a mere four or five years ago. The increasing pressure put on the network is often mentioned almost in passing, as though merely moving from 10Gbps to 40Gbps to 100Gbps will solve the problem. Move along now, nothing to see here but a higher flow of packets.

But that higher density of packets along with greater diversity of content coupled with distribution through cloud computing that's creating other issues for network services whose purpose it is to collect, analyze, and act upon those packets.

IDS, IPS, secure web gateways, voice analyzers, honeypots. There are myriad network infrastructure devices that are tasked with analyzing the content of packets flowing in and out of the data center that find it more and more difficult to scale along with the rapid growth of data on the network. Application Performance Monitoring (APM) systems, as well, often take advantage of port mirroring as a way to collect and analyze intra-system traffic to pinpoint configuration or network issues that may cause performance degradation.

These systems need one thing: all your (relevant) packets. The problem is that on most switches, you can designate only a couple of ports as egress span ports and you may have three, four or more devices and systems that need those packets. And Heaven forbid you have a desperate need to later tap into the switch to troubleshoot an urgent issue.

The answer in the past has been some highly complex network topologies that are difficult to maintain and not easy to extend when the next system needing all your packets is deployed. Additionally, cloud-deployed applications and systems are not easily included, even though organizations desire the same level of visibility and analysis of those packets as is found in the data center.

One answer to these issues is found in what Gartner is calling Network Packet Brokers. One such provider in this space is VSS Monitoring, which recently introduced a new set of solutions to resolve this lack of visibility both in the data center and within the cloud.

VSS MONITORING

VSS Monitoring has been around since 2006, shipping aggregation and related management products. Now it's introduced several new products that assist in the goal of collecting packets across the increasingly cloudy landscape and getting them to the right place at the right time, a market being referred to as "Network Packet Brokers (NPB)". Gartner analysts describe these solutions as consisting of "devices that facilitate monitoring and security technologies to see the traffic which is required for those solutions to work more effectively. They could be called "monitoring switches" "matrix switches" (Application Aware Network Performance Monitoring (NPM) and Network Packet Broker (NPB) research).

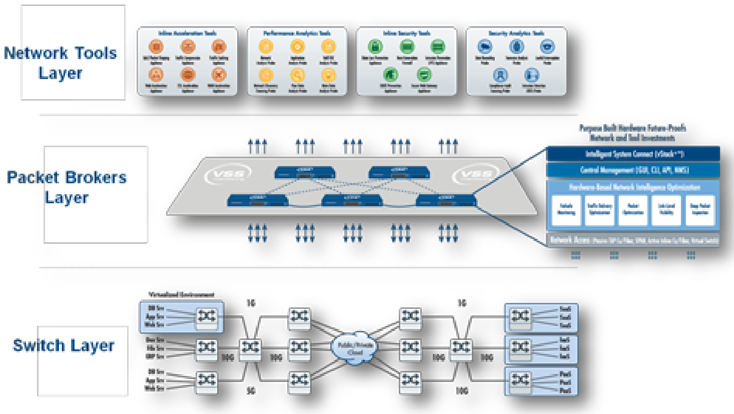
Network Tools Layer
Packet Brokers Layer
Switch Layer
Purpose Built Hardware Future-Proofs Network and Tool Investments
Intelligent System Connect (vStack+)
Central Management (GUI, CLI, API, NMS)
Hardware-Based Network Intelligence Optimization
VSS

NPB solutions must be able to perform many-to-many port mapping using a GUI or CLI, filter packets at L2-4, and perform packing slicing and deduplication as well as aggregation and intelligent distribution. This last criteria is an important one, as it allows operators to filter out noise when directing packets to reduce the requirement that analyzers and systems process (and ultimately discard) irrelevant traffic.

VSS Monitoring has introduced a set of solutions that meet (and in some cases exceed) the requirements laid out by Gartner (VSS supports L2-7 filtering) and that further expand the scope of such solutions into cloud computing environments:

- New packet broker appliances -- vBrokers™
- Expanded system-level scalability – vMesh™
- Topology-level unified management console – vMC™

VSS achieves this inter-cloud monitoring capability by leveraging a proprietary L2 bi-directional protocol for its interconnects called vMesh. Its vBrokers are purpose-built appliances that can interconnect with one another using vMesh to form a virtual network tool optimization fabric . These vBrokers can be deployed across LAN, WAN segments and in a wide variety of cloud network infrastructure environments using the vMesh architecture effectively forming an overlay network over which packets are shared. From there, it's a matter of dragging and dropping policies and configuration via its vMC unified management console to access network packets on demand and properly direct them based on organizational needs. the VSS' new vMesh technology can scale out to up to 256 devices and 10,000 and more ports.

VSS also provides an Open XML API that encourages integration. Configuration, remote management, metrics, etc… can be achieved via this API. VSS solutions today are not supported by common provisioning and automation frameworks (Chef, Puppet, OpenStack) although that is something that may very well be supported in the future.

Still, the ability to reach out into the cloud and direct packets to DC-hosted infrastructure services providing analysis, security, or other functions solves a major issue with managing cloud-deployed applications: visibility.

SDN versus NETWORK PACKET BROKERS

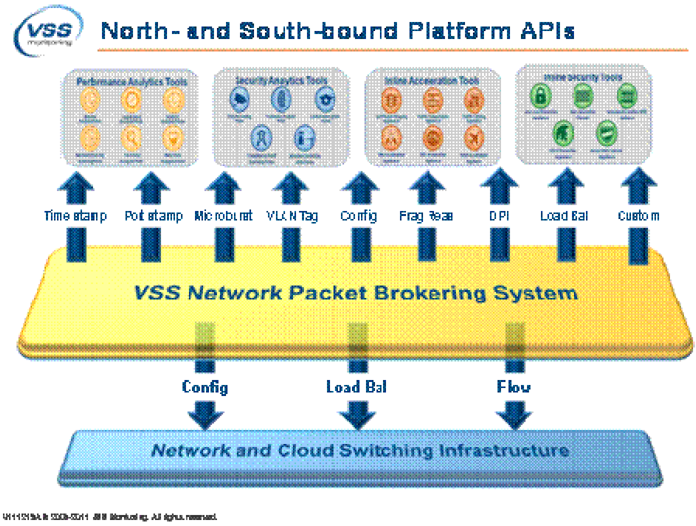
VSS monitoring
North- and South-bound Platform APIs
Performance Analytics Tools
Security Analytics Tools
Inline Acceleration Tools
Inline Security Tools
Time stamp
Port stamp
Microburst
VLAN Tag
Config
Frag Reas
DPI
Load Bal
Custom
VSS Network Packet Brokering System
Config
Load Bal
Flow
Network and Cloud Switching Infrastructure
2009-2011 VSS Monitoring. All rights reserved.

At first read, this sounds a lot like a suggested SDN (Software-Defined Networking) use case (found on SDN Central) that posits the use of OpenFlow as a Virtual Patch Panel. However, on deeper inspection there are some distinct differences between the two solutions.

While both are focused on solving what is essentially a port forwarding problem (port spanning is really just a case of directing ingress packets on one port to more than one egress port) SDN is (today) more disruptive a solution both in the enterprise and in the cloud. While it's true that with both solutions you need some means to direct ingress packets to the desired egress port, VSS' solution does not require that the switches in question be OpenFlow enabled (which may be problematic in cloud environments). Additionally, the forwarding mechanism available with OpenFlow is simple forwarding – packet in, packet out. While a more sophisticated forwarding algorithm could certainly be employed, this would require specific code. VSS, on the other hand, enables intelligent forwarding of actionable packets, reducing the amount of irrelevant traffic any given infrastructure solution might need to process. Voice analyzers, for example, need only see VoIP, SIP and related traffic. Such a system doesn't need to inspect a JSON exchange, nor will it – the packets will be inspected and discarded. Using a more intelligent approach, VSS can intervene and eliminate the overhead associated with inspecting and discarding non-actionable traffic. This offload-like capability improves the capacity and performance of packet analyzing systems.

Further more, VSS offers a single-pane of glass management system for monitoring and managing its packet brokers, while an OpenFlow-enabled solution currently does not. This is certainly an area of exploration for SDN and OpenFlow-enabled devices and future value-add for those banking on SDN; admittedly the technology is still very much in its nascent phase and maturation will bring more mature, robust solutions not only in core device support but in management and niche-market solutions.

The other issue is deployment in the cloud, as a virtual device. The good news is that Open vSwitch is embedded in many hypervisors and is available as a package for a variety of Linux-based systems. The bad news is that in some cloud environments (like Amazon) these approaches may not be possible to deploy and/or take advantage of, thus rendering an SDN-OpenFlow approach more or less toothless. VSS' packet broker, vBroker, supports a broad set of physical and virtual environments (i.e. physical and virtual span ports, ability to filter and remove VN-Tags, etc) which enables a wider set of cloud environments to take advantage of the capabilities.

That's not to say the two couldn't be combined, either. In fact, VSS could be described as "SDN for networking monitoring", though VSS itself has not chosen to represent its solution this way. But essentially it's acting in the same manner as SDN – simply confined to a specific area of functionality – monitoring. As I posited in the past, I suspect we'll continue to see these kinds of "pockets of SDN" capabilities pop up to resolve some pressing issues that simply can't be addressed by traditional networking methods – or at least can't be addressed efficiently or in an acceptably rapid manner. In such an architecture (one comprised of controllers at strategic points of control) VSS Monitoring is certainly positioned to act as the control point for managing a broadly distributed monitoring network.