

A large, semi-transparent graphic of the letters "SDN" is overlaid on a dark blue background. The background features a grid of squares, some of which contain faint, abstract images related to technology, such as circuit boards, network diagrams, and server racks. The "SDN" text is rendered in a bold, sans-serif font, with the "S" and "D" in a lighter blue and the "N" in a darker blue.

**ACHIEVING
OPERATIONAL
INSIGHT
IN SDN & NFV
ENVIRONMENTS.**

Software-Defined Networking (SDN) and its counterpart, Network Functions Virtualization (NFV), have ushered in a new way of thinking about application and service delivery. Leading service providers and enterprises are now modeling their digital infrastructure, redundancy and scale-out patterns from Web-scale and DevOps companies, and applying the methods discovered in that world to SDN and NFV platforms to operate voice and data networks.


In the past, constraints with network and storage capacity and performance dictated how organizations developed and deployed such applications and services. It was common for a development team to release the initial version of a new application, uncover constraints once the application launched in a production environment, and then re-code the application around those constraints.

Today, rather than allow the infrastructure to determine application architectures, developers are taking an “apps-down” approach to meeting demanding time-to-market pressures. The agility provided by SDN and NFV allows infrastructures to spin up additional resources on demand as required by the applications. As a result, the constraints of the past have been largely eliminated, paving the way for a “software-centric” or “developer-driven” infrastructure where success no longer depends upon premium, on-premises hardware.

However, achieving operational insight into these highly dynamic environments is near impossible with existing performance monitoring tools. That’s because they were never designed for today’s elastic, complex, digital infrastructure.

In order to achieve operational insight, you need to take the same top down approach now employed by developers. It’s no longer a simple game of alerting on component failures across your digital infrastructure. The management platform must deliver analytics that reveal the business impact of infrastructure performance on your applications and services. In addition, assuring service delivery with SDN and NFV requires you to understand both your physical and virtual infrastructure, map the rapidly changing dependencies, and correlate events.

This whitepaper addresses the new requirements for performance management platforms. It explains why the traditional way of thinking no longer applies to SDN and NFV infrastructures. It also outlines a three stage methodology for delivering operational insight into these complex, dynamic environments.





1

STEP 1: COLLECT AND VISUALIZE.

Any sound service and resource assurance strategy begins with collecting performance metrics that support application and service delivery. However, management tools optimized for static infrastructure do not perform well in software-defined environments because:

- Software-defined infrastructure doesn't respond to traditional protocols like SNMP
- You can't poll for metrics every 5 minutes when resources move so fast
- It's difficult to maintain relationships among constantly shifting resources

In addition, in an SDN environment, applications are loosely coupled to the resources that support them. This is in stark contrast to the recent past, when applications were tied to a single physical machine and had a lifespan in sync with that hardware. Now, applications swap out their resources on an ongoing basis.



To succeed in this new environment, performance management platforms must begin by supporting a range of metrics across physical and virtual infrastructures. Data collected from bare metal and hypervisors must sit alongside data from a heterogeneous set of SDN/NFV solutions, such as those from OpenStack, Cisco, VMware and others. If you're troubleshooting why a given application performs sluggishly, you'll spend hours stringing together reports across your digital infrastructure if you don't have visibility into physical and virtual performance from one place.

Many SDN and NFV solutions don't support performance monitoring using legacy protocols like SNMP. Instead, they expose their own unique APIs for performance data. These APIs can also provide context in the form of inventory, metadata, topology and dynamic device groups/relationships shared with the performance management platform.

A performance management platform that delivers the value you expect in an SDN/NFV environment will need to rapidly develop and ship collection integrations for an ever-expanding number of APIs. It must also lifecycle them quickly, as APIs from SDN suppliers can change significantly every 6 months.

Once you have the data, it's time to envision the metrics. This starts at the application and service level with visualizations that match your delivery model. Remember, we're no longer talking about an application as a fixed set of devices. Instead, think of the service as a collection of variables. For example, you may have thousands of performance indicators spread across hundreds of virtual components in your digital infrastructure. But these KPIs aren't coupled to physical hardware anymore. They constantly move to take advantage of less congested infrastructure for network or compute needs. Pinpointing a problem in this environment with traditional infrastructure performance management tools is like trying to find a needle in a haystack.

Instead, operators of these management systems need a workflow that "remixes" all of their performance data (end user experience, metrics, flows and logs) into the context of the service or applications to which they're responding. The performance management solution should connect the dots, show you what's unique and needs investigating, and map it to your business applications and services.

Historically, this was accomplished by baselining every metric in your infrastructure and then alerting when the system detected a deviation from that baseline performance. But traditional baselines don't work well in these highly elastic environments. First, resources move too fast in virtualized environments. You can't wait 10 weeks to establish a true baseline of "normal" performance. If something happened hours ago, it may already be irrelevant. Second, it's normal for accordion-like fluctuation to occur among software-defined resources.

Visualization of performance issues requires a re-tooling of how we think about infrastructure problems. In the past, if you had a pool of 100 servers, you'd want to know which server is not performing like it typically did. But with SDN, the question you must ask is, "Which server is not behaving like the other 100 servers in the pool?"



2

STEP 2: ANALYZE AND INTERACT.

The traditional impact of down infrastructure is less relevant in a software-defined world. You may have an entire building (on-premises or in the cloud) of server racks at your disposal. You're not so concerned with a virtualized router or application building block going down. It's more important to understand when you need more reliable pieces of infrastructure to sustain the business growth. And scaling your applications is now more of a budgetary concern than an engineering challenge. You can scale as high as you can afford.

To operate in this mode, you need a performance management platform that shows the impact of your digital infrastructure on the business. You need to understand if this accordion of things allows you to achieve the SLAs you have in place. From there, you can make smarter decisions about how much you want to spend to maintain or grow revenues in different regions.

To properly analyze performance in SDN/NFV deployments, you need to act like a doctor questioning a patient. Start with the broadest questions, and then narrow down until you've eliminated the majority of options and honed in on one or two potential causes. For the operations staff, this means starting at the application with a strong understanding of the customer experience. For instance, why does one instance perform well and the other doesn't? In order to figure this out, you need visibility into how services are provisioned across disparate infrastructure and paths.

Working from the customer experience down is very different than working from the traditional device level up. But it's also the way most executives think. When was the last time an executive called on his network team to tell them a specific router was slow? Never, of course. They want to know why the application is slow. You need to turn infrastructure speak into a language that executives can understand, and that means starting the diagnosis at the *experience* level, not the device level.

And you need to do it fast, because time is valuable.

Let's face it – our brains can't keep up with pattern recognition algorithms. Instead, the performance management platform must put your infrastructure in context, allowing you to make rapid business decisions.

In a software-defined world, your goal should be to help align the CIO and his team to business objectives. In doing so, you can begin to move away from constantly reacting to blinking red lights across your infrastructure.



3

STEP 3: OPTIMIZE AND AUTOMATE.

With SDN and NFV architectures, service assurance is provisioned with the application. Performance monitoring needs to be provisioned within that loop as well, in order to optimize and automate the infrastructure for performance and cost.

It starts with operational leaders and IT partnering to encode business workflows, processes and logic into software. The goal is to make digital infrastructure that is fast and reactive (the purpose of pursuing SDN/NFV in the first place).

Andre Fuetsch, senior VP of architecture and design for AT&T, understands the need to react in near real-time to demands for capacity. AT&T has been at the forefront of deploying SDN and NFV solutions in the form of their Network On-Demand service, which the company claims has already reduced provisioning cycle times up to 95% in less than a year. It was a move driven by necessity, as AT&T realized traditional networking could no longer keep up with traffic growth.

In a 2015 interview, Fuetsch explained the volatility of traffic today: "Now the traffic capacity engineer has to worry about the latest Kardashian streaming video, and who has been watching House of Cards. And the biggest event we worry about is when Apple releases their next big iOS upgrade because you never know who is going to push the button to download now. So events like that really add a lot of volatility and unpredictability to our network."

Leaders like Fuetsch don't want to plan for capacity needs ever again, especially the unpredictable events that drive massive surges in traffic. They want the infrastructure to react and automatically scale out when an application or service demands. And then scale back when demand decreases.

Because the performance management platform is closest to the infrastructure, it can – and should – help make decisions and provide automated SDN intelligence.

To enable the real time, closed-loop goals of delivering services that self-optimize, the solution must be able to notify external systems of anomalies discovered in the analyze stage. OPNFV describes the receiver of this notification as a Consumer, which is developed by the business or its suppliers to do things like scale an application or resources up or down, fail over between two network paths, adjust environmental cooling parameters of a part of a data center facility, and more.

Essentially, performance management becomes a way for businesses to measure their infrastructure for performance and cost. It makes reactive business decisions based on the data it collects and analyzes. A capable performance management solution will always provide out-of-the-box reports and graphs, but the real value ahead will be for businesses to encode it into their processes to react to changes in machine time instead of in human time.



SEVONE VISION FOR PERFORMANCE MANAGEMENT OF SDN & NFV.

In a software-defined infrastructure, you need to adjust your mindset when it comes to managing performance. You can't look for alarms or reports to dictate business decisions. Your performance management platform must be part of the loop, interacting with the digital infrastructure and orchestrating change based on intelligence gathered. It may trigger an action to fire off additional AWS or Azure resources to support an application's demand for more resources, or advise the scaling back of those resources when demand diminishes.

A performance management platform should adhere to the threestep process outlined in this whitepaper: collect and visualize, analyze and interact, and optimize and automate.

When considering an infrastructure performance management solution as part of your SDN architecture, look for one capable of:

- Providing full visibility into your dynamic infrastructure in hyperscale environments
- Understanding that services are now a collection of variables and not a fixed set of devices
- Presenting the effects of infrastructure on business applications and SLAs
- Making recommendations within a continuous feedback loop for optimized infrastructure

In an SDN environment, operational functions – including performance management – should require minimal configuration and respond to current workloads by provisioning their own infrastructure and services automatically. They should recover from faults in the environment without data loss or interruption to collection and reporting. And they must support elastic growth in the data storage layer.

Finally, it's important to be holistic in your approach to assurance. Even with the advent of SDN and NFV, there's still no such thing as a pure virtual infrastructure. At some point, you'll run into physical infrastructure, and you'll have to manage a hybrid environment for a while. You need to monitor and manage both physical and virtual aspects from a single platform, so you can understand dependencies and business impacts in these highly elastic and complex environments.

About SevOne.

SevOne provides the world's most scalable infrastructure performance monitoring platform to the world's most connected companies. The patented SevOne Cluster™ architecture leverages distributed computing to scale infinitely and collect millions of objects. It provides real-time reporting down to the second and provides the insight needed to prevent outages. SevOne customers include seven of today's 13 largest banks, enterprises, CSPs, MSPs and MSOs. SevOne is backed by Bain Capital Ventures. More information can be found at www.sevone.com. Follow SevOne on Twitter at @SevOneInc.