



# **Network Operations Analytics**

---

## **Solution Guide**

**Version 2.4.7 (Build 2.4.7.0.23)**

**January 2017**

# Contents

---

<b>Overview .....</b>	<b>2</b>
Announcement of product name change .....	2
Introduction.....	2
<b>Solution Components.....</b>	<b>3</b>
<b>Before You Begin .....</b>	<b>4</b>
Pre-Installation Checklist.....	4
Installation Steps.....	4
<b>NetFlow Optimizer.....</b>	<b>5</b>
Agentless Deployment .....	5
NFO Configuration .....	5
<b>Network Operations Analytics Module .....</b>	<b>6</b>
Reported Metrics.....	6
Network Path Reporting and Rendering .....	6
How sFlow and Sampling Affect Path Reporting .....	8
How VDS Configuration Affects Path Reporting.....	9
Module Configuration .....	9
<b>Network Operations Analytics for Splunk.....</b>	<b>13</b>
NetOps Splunk App Installation .....	13
NetOps Splunk App Dashboards .....	13
Path Analytics Dashboard.....	13
Top Tunnels (VTEPs) Dashboard.....	15
Top VMs by Traffic Dashboard .....	16
Network Health Dashboard.....	16
<b>Start Your Trial .....</b>	<b>17</b>

# Overview

---

## Announcement of product name change

Our core product NetFlow Integrator, or NFI, is now called NetFlow Optimizer (NFO). All references to NetFlow Integrator or NFI in this document apply to NetFlow Optimizer.

## Introduction

This solution brings new operational capabilities to network administrators or virtual infrastructure administrators by providing complete visibility into virtual and physical networks.

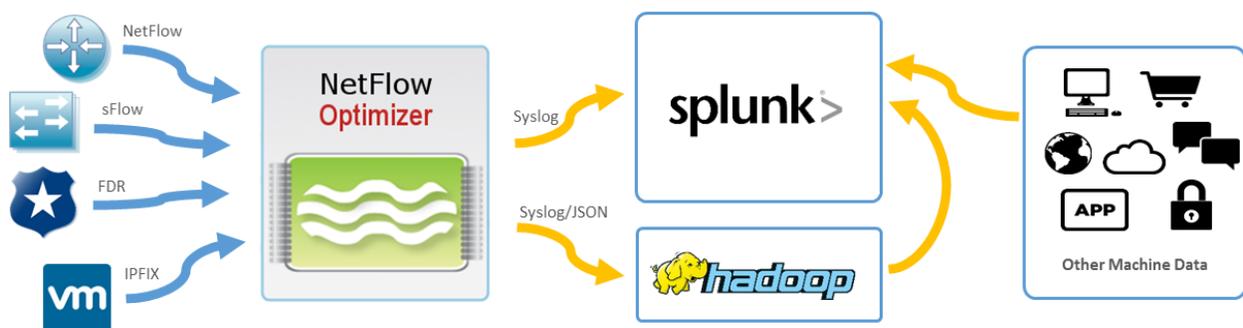
One of the biggest operational concerns for IT professionals and virtual infrastructure administrators is the lack of visibility between the virtual and physical networking layers -- how to trace and troubleshoot connectivity issues. Lacking traceability leads to longer time to resolution, and unacceptable outage time frames for many customers.

NetFlow Logic's Network Operations Analytics provides an easy and scalable way to analyze the massive volumes of network metadata generated by your routers, switches, next generation firewalls, load balancers (Cisco ASA, F5, Palo Alto Networks, etc.) – whether it is from a physical or virtual network, or both. This solution also provides end-to-end operational information about virtual and physical network traffic paths.

Monitor your end-to-end infrastructure to avoid service degradation or outages. Gain visibility into how network conversations traverse the virtual and physical network, including device interface health score, for efficient troubleshooting and planning.

The core of this solution is NetFlow Optimizer™ (NFO) and Network Operations Analytics NFO Module (NetOps).

**NFO** is a processing engine for network flow data (NetFlow, IPFIX, J-Flow, sFlow, etc.).



NetFlow Optimizer accepts network flow data from network devices (routers, switches, firewalls), applies map-reduce algorithms to the data to extract the information needed to address desired use cases, converts the processed data to syslog (or other formats such as JSON), then sends that useful information to your visualization platform such as Splunk Enterprise or other SIEM.

# Solution Components

Based on your environment you may choose appropriate components, listed below. You have to install three required NetFlow Logic's components -- NetFlow Optimizer, NFO Updater, and Network Operations Analytics Module, and one or more optional information management system components for visualization.

**Table 1 – Solution Components**

Component	Platform	Description
NetFlow Optimizer (NFO) RLS 2.4.7.0	Linux or Windows	This is core component for flow processing. Available for Windows, Linux, or as Virtual Appliance. (Downloadable from NetFlow Logic's web site – <a href="http://www.netflowlogic.com/download/">www.netflowlogic.com/download/</a> ).
NFO Updater RLS 2.4.7.0	Linux or Windows	This component is for integration with vCenter (and optionally with NSX). Available for Windows or Linux. (Downloadable from NetFlow Logic's web site – <a href="http://www.netflowlogic.com/download/">www.netflowlogic.com/download/</a> ).
Network Operations Analytics (NFO Module)	NetFlow Optimizer 2.4.7.0	NetFlow Optimizer Module. This component contains analytics for the solutions. Select Windows or Linux version to match your NFO platform. (Downloadable from NetFlow Logic's web site – <a href="http://www.netflowlogic.com/download/">www.netflowlogic.com/download/</a> ).
Network Operations Analytics for Splunk (Splunk App)	Splunk Enterprise 6.4 or above	Network Operations Analytics visualization in Splunk. It includes Network Health and Path Analytics dashboards. (Downloadable from <a href="https://splunkbase.splunk.com/app/2824/">https://splunkbase.splunk.com/app/2824/</a> ).

**Table 2 – Environment / Functionality Matrix**

Data Center Environment	NetOps Functionality
Any virtualized data center with heterogeneous network equipment capable of sending flows (NetFlow v5/v9, sFlow, IPFIX)	By processing flows from network devices, NetOps will report in real time: <ul style="list-style-type: none"> <li>- Health score and Failure risk of network physical interfaces</li> <li>- Path between communicating end points (VM to VM or VM to Host)</li> </ul>
VMware Virtual Distributed Switch (VDS)	All NetOps functionality is available. By enabling IPFIX on VDS, and flows on other network devices, NetOps will correlate virtual and physical traffic, resolving VM-to-VMhost adjacency, and ToR (Top of the Rack) switches
VMware Virtual Standard Switch (VSS)	NetOps will function with VSS. By processing flows from network devices, NetOps will report Health score and Failure risk of network devices and interfaces. Path between communicating peers (VMs and/or physical hosts) is reported, but VM-to-Vmhost adjacency is not resolved
VMWare NSX	NetOps relies on VDS for correlating virtual and physical network traffic. Other components, including VMware NSX virtualization platform itself, are optional

# Before You Begin

---

## Pre-Installation Checklist

Please be sure to have the following before you begin the installation of NetOps components:

- ✓ Installation Prerequisites – You have to login as **root** for Linux and **administrator** for Windows installations and updates of NetFlow Optimizer and NFO Updater.
- ✓ License – A license from NetFlow Logic sales is required before you can begin using NetFlow Optimizer software. Please contact [sales@netflowlogic.com](mailto:sales@netflowlogic.com) for a license.
- ✓ Physical and Virtual Network Devices - Please refer to the “Configuring NetFlow” section in your Cisco (or other) device or VMware VDS documentation.
- ✓ Visualization Platform – Splunk Enterprise.

## Installation Steps

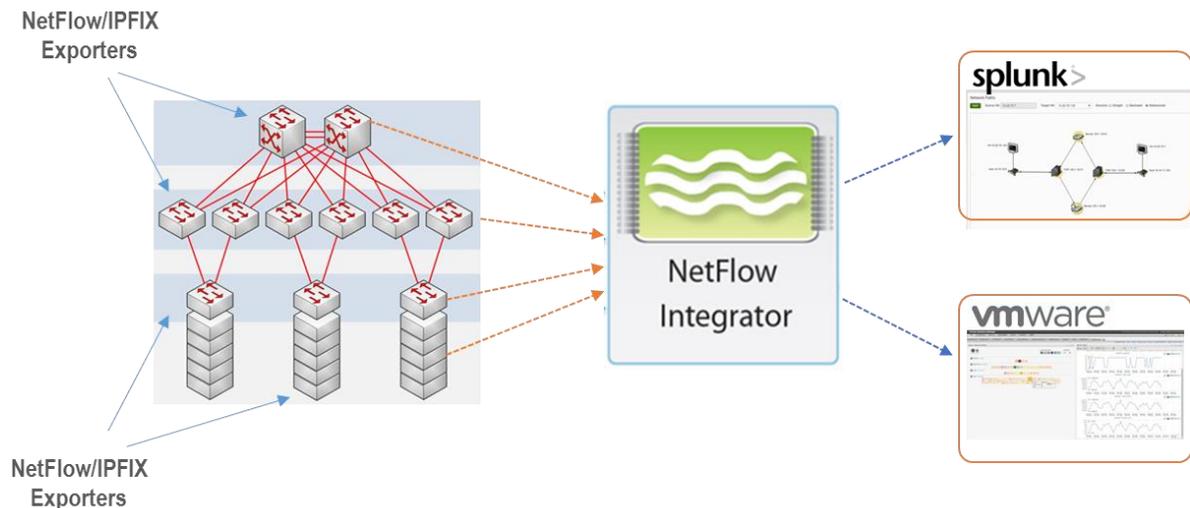
1. Install NetFlow Optimizer - see NetFlow Optimizer Installation and Administration Guide for details at <https://www.netflowlogic.com/resources/documentation/>
2. Configure NFO input (port numbers where flow information is sent to from network devices and VDSs)
3. Configure NFO output for Splunk (IP address and port number where NFO sends out syslogs)
4. Configure NFO SNMP data retrieval service
5. Install NFO Updater
6. Upload Network Operations Analytics Module into NetFlow Optimizer – see “Module Configuration” on page 9 of this document
7. Configure Network Operations Analytics Module parameters:
  - a. List of vCenter Virtual Machines – set automatic updates
  - b. List of ToRs management addresses – this list is optional, but it would help resolving ToRs
  - c. Install and configure your visualization component - Network Operations Analytics for Splunk (Splunk App) – see NetOps Splunk App Installation on page 13
8. Restart NetFlow Optimizer – the system is operational within approximately 15 minutes

# NetFlow Optimizer

---

## Agentless Deployment

Installing yet another agent on a large network to provide comprehensive network traffic information is costly, difficult to roll out and manage. NetFlow Optimizer was designed to avoid those issues. A single instance of NFO, deployed in a data center, is capable of processing and analyzing massive volumes of network metadata.



Configure your all flow-capable exporters, such as ToR switches, core and aggregation switches, routers, and virtual switches, such as VDS or Open vSwitch to send NetFlow/sFlow/IPFIX to NetFlow Optimizer.

## NFO Configuration

NetFlow Optimizer is available as Windows or Linux installers, or as a Virtual Appliance. When installed, NFO is automatically configured to listen to flow data on UDP port 9995. You can change this port and/or add additional ports to receive flows from all your exporters. You also need to configure NetFlow Optimizer output. You can configure up to 16 output destinations.

You also need to install NFO Updater on the same server with NFO or on a separate server. NFO Updater is required for integration of the solution with vCenter.

Please refer to **NetFlow Optimizer Installation and Administration Guide** for additional details.

# Network Operations Analytics Module

---

This Module reports network health by detecting anomalies in traffic volume and packet rate. It reports health and failure risk of each interface of all physical network devices configured to send flows to NFO. Health score is calculated using multivariate algorithms based on the actual network traffic reported by network devices in flow records. It also reports **virtual and physical** network traffic paths for all VM-to-VM, VM-to-Host, and Host-to-Host conversations.

## Reported Metrics

NetOps Module reports information about health of the network devices and their interfaces, as well as the actual path traversed through network devices by flows between communicating hosts (VMs or real hardware). This information is reported periodically, and the frequency is configurable (see Module Configuration section).

**Table 3 – Reported Metrics**

Metric	Description
Relative Traffic Load	Interface relative load is computed by dividing actual traffic going through the interface by the nominal speed of the interface. The nominal speed is obtained through SNMP polling of the network device (ifSpeed or ifHighSpeed), %
Relative Packet Rate	Relative rate is computed by dividing current packet rate of the interface by the maximal packet rate sustainable at a current average packet size observed during the data collection interval, %
Health Score	The health score is a single value metric calculated for each network interface over a short data observation time interval, e.g. 30 seconds. It takes into account relative traffic load and relative packet rate on the interface. Health score values range from 0 to 100. The health above 40 is normal, indicating no attention is required. Health score between 25 and 40 is the indication of medium to serious problem. The score below 25 indicates that the interface is not functioning properly
Failure Risk	Failure risk is a predictive indicator. Its calculation is based on Health Score, Load, and Relative Rate, plus dynamic changes of these metrics. Failure risk values range from 0 to 100. The risk below 60 indicates that the interface has no problem and no problem is expected in the future. 60 to 75 indicate a low to medium chance of a failure. The risk above 75 indicates a serious problem now or failure in the near future.

## Network Path Reporting and Rendering

Network path is a reconstruction of routes in a network taken by the traffic between a pair of communicating peers A and B (VM-to-VM, VM-to-Host, and Host-to-Host). Paths from A to B and from B to A are reported separately. It is important to keep in mind that flow-based path resolution is a best effort process, as device adjacency may not be certain. Path reporting relies on forwarding information (next hop), when available, and on traffic statistical methods of the Module. In certain cases additional path

resolution could be achieved by superimposing paths from A to B and from B to A. However, some segments of the network routes may remain unresolved, and are shown with dashed lines.

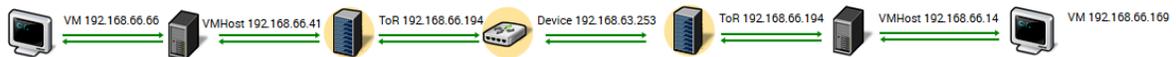
Let's consider the following scenarios.

### Scenario 1. Partially resolved network path



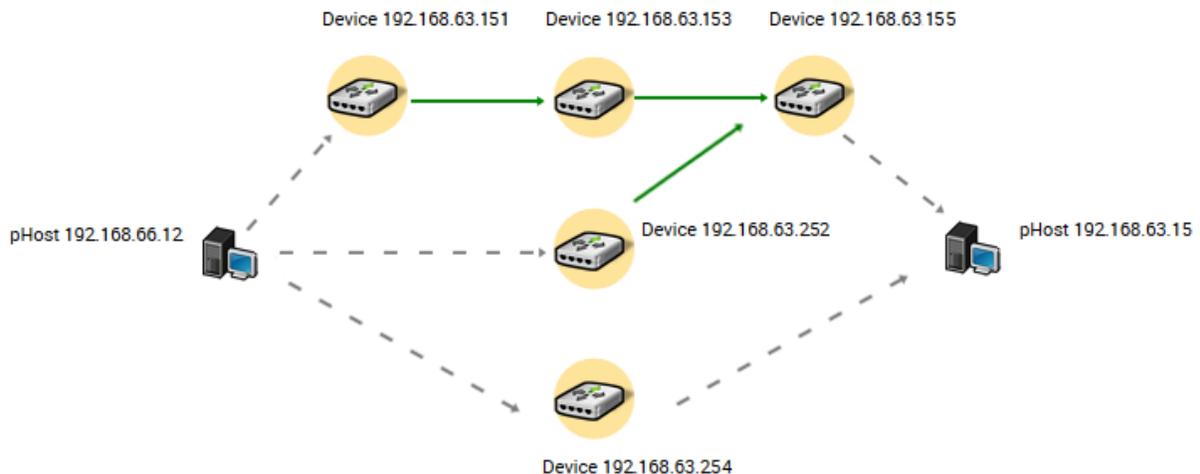
This example shows a partially resolved network path in the VM<sub>1</sub> -> VM<sub>2</sub> direction. Please note that the first hop (Host<sub>1</sub> -> ToR<sub>1</sub>) is not resolved: ToR<sub>1</sub> is a first network node in the segment which reports VM<sub>1</sub> -> VM<sub>2</sub> communication but we have no knowledge if there are other network nodes between Host<sub>1</sub> and ToR<sub>1</sub> which may be forwarding the traffic and not configured to send flows to NFO.

### Scenario 2. Fully resolved network path



A complete resolved network path is produced by superimposing VM<sub>1</sub> to VM<sub>2</sub> and VM<sub>2</sub> to VM<sub>1</sub> communications observations, as each direction complements the unresolved segment of the path in the opposite direction.

### Scenario 3. Multiple intersecting segments



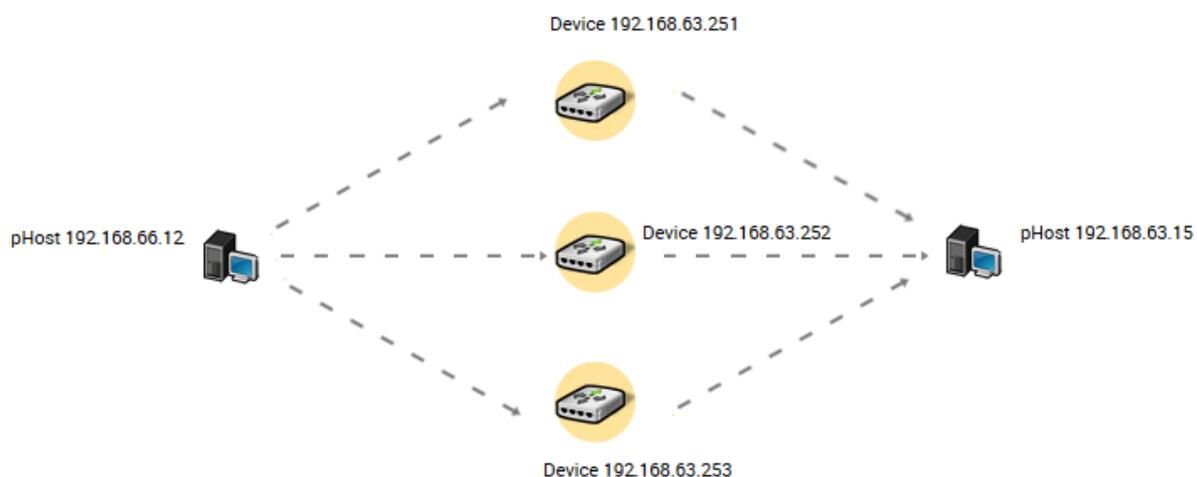
This example shows topology with five devices where network traffic from Host<sub>1</sub> (192.168.66.12) to Host<sub>2</sub> (192.168.63.15) may pass in three different ways: D<sub>1</sub> (192.168.63.151) to D<sub>3</sub> (192.168.63.153) to D<sub>5</sub>

(192.168.63.155), D<sub>4</sub> (192.168.63.252) to D<sub>5</sub> (192.168.63.155), and D<sub>2</sub> (192.168.63.254). In practice this is a common load balancing scenario.

Please note a dashed line between D<sub>5</sub> and Host<sub>2</sub>. This means that none of the nodes in this path was determined to be adjacent to the Host<sub>2</sub> and therefore this segment is shown as unresolved.

A dashed line between D<sub>2</sub> and Host<sub>2</sub> implies that traffic flowing from Host<sub>1</sub> to Host<sub>2</sub> was reported by device D<sub>2</sub> which did not provide any forwarding information and was not reported as next hop by any of the other devices, so both segments Host<sub>1</sub> to D<sub>2</sub> and D<sub>2</sub> to Host<sub>2</sub> are shown as unresolved.

#### Scenario 4. “A bag of routers”



In this example all network devices between Host<sub>1</sub> (192.168.66.12) and Host<sub>2</sub> (192.168.63.15) did not report any forwarding information, therefore all segments of three paths are shown as unresolved.

#### Scenario 5. Multiple paths to a device



In this example device D<sub>1</sub> reports traffic from Host<sub>1</sub> (192.168.66.12) to Host<sub>2</sub> (192.168.63.15) via a plurality of the interfaces. Such situation is possible when there is one or more intermediate network devices situated between Host<sub>1</sub> and D<sub>1</sub> (192.168.63.252) which do not report the traffic flows.

## How sFlow and Sampling Affect Path Reporting

NetFlow Optimizer NetOps Module determines path only from information carried in NetFlow. It will use next-hop IP address to build the corresponding path segment when this information is available. When next-hop is not reported in NetFlow, the Module will use statistical methods to build paths, thus the sampling rate of NetFlow or sFlow may have a direct effect on the accuracy of path building. In addition, as path building process is dynamic, sampling may result in showing path segments that are no longer active.

# How VDS Configuration Affects Path Reporting

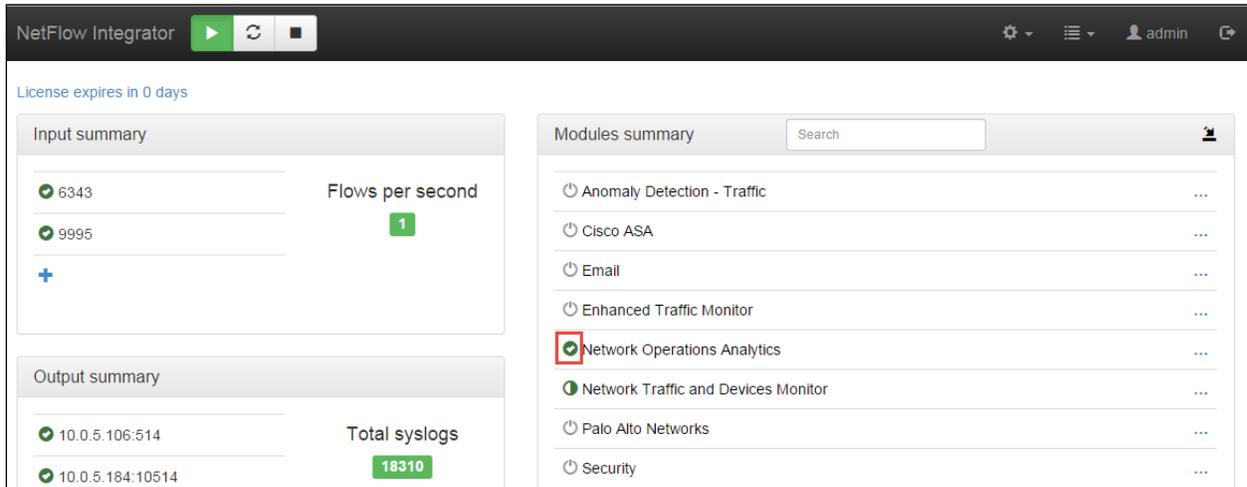
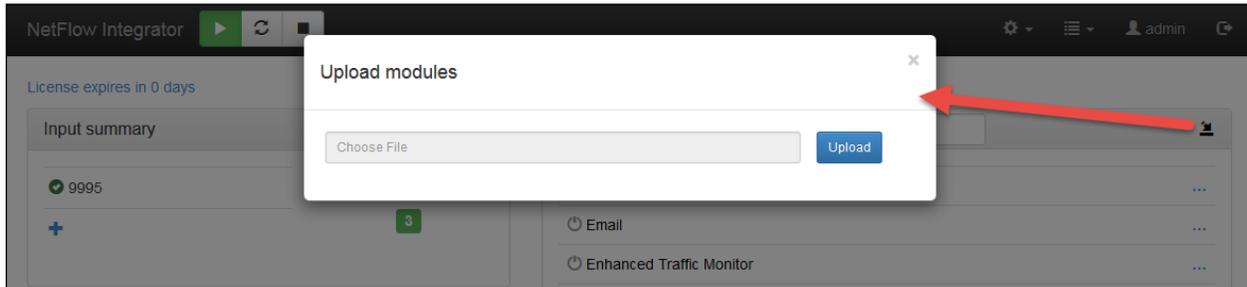
When configuring VMware vSphere Networking (for details go to <https://pubs.vmware.com/vsphere-65/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-55FCEC92-74B9-4E5F-ACC0-4EA1C36F397A.html>) make sure:

1. The Switch IP address is specified (Step 5)
2. The option **Process internal flows only** is disabled (Step 8)

## Module Configuration

The Module should be uploaded into NFO and enabled. In NetFlow Optimizer Home page click on upload button, and select the package to upload (e.g. `network_operations_analytics-2.4.7.0.23-linux-x86_64.zip`).

Once uploaded, click on  to enable the Module. You don't have to restart the server – the Module is operational when enabled.



NetOps Module is highly configurable. Click on Network Operations Analytics Module and you will be presented with the following screen.

**NetFlow Integrator** 10180: Network Operations Analytics

Status: ● Enabled (Last enabled: 2015-11-17 17:14:07)

**Parameters:**

Default sampler rate	<input type="text" value="1"/>
Network path information reporting interval, sec	<input type="text" value="300"/>
Output method: 0 - syslog (push), 1 - REST (pull), 2 - both	<input type="text" value="2"/>
Inactivity timeout for network paths, sec	<input type="text" value="300"/>
N - number of reported tunnels	<input type="text" value="50"/>

**Data collection interval, sec**

Network Operations Analytics	<input type="text" value="23"/>
Network Paths Analytics	<input type="text" value="37"/>
Network Paths Reporter	<input type="text" value="203"/>
VM - VM Hosts Adjacency	<input type="text" value="41"/>
ToR Switches Discovery	<input type="text" value="29"/>
Top Tunnels	<input type="text" value="30"/>

**Data sets**

List of vCenter Virtual Machines	Auto created by "VMWARE vCenter Virtual Machine monitor" agent	
List of ToRs management addresses	10180: List of ToRs management addresses	

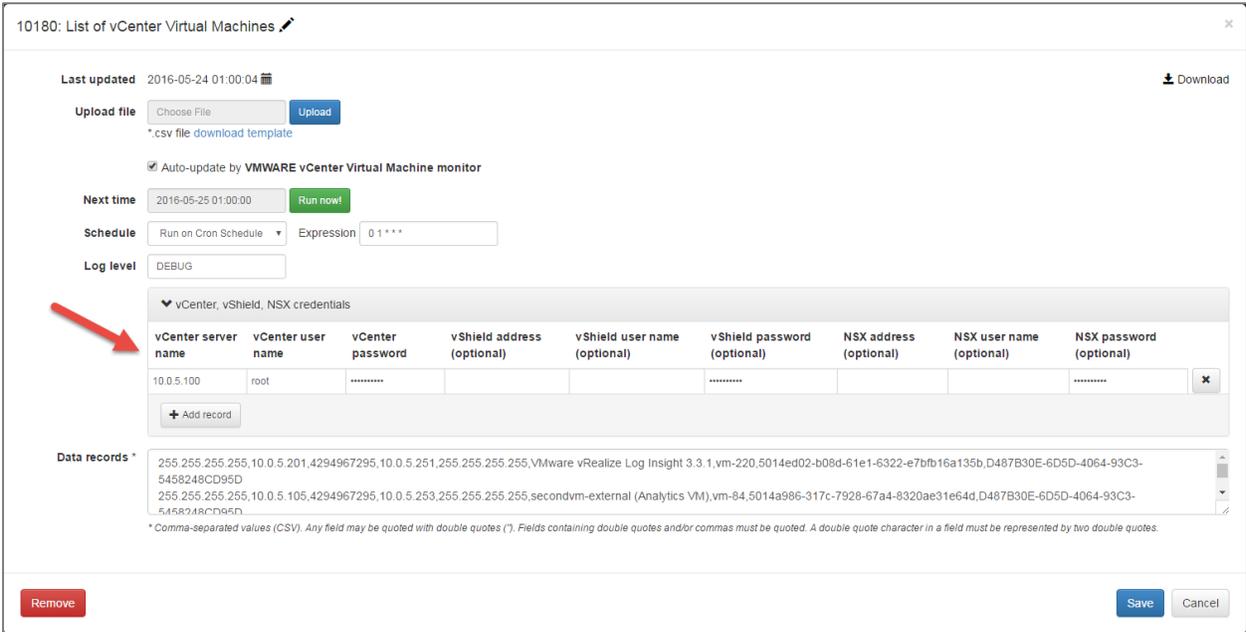
Contact NetFlow Logic support before changing any parameters.

**Table 4 - Configuration Parameters**

Parameter Name	Description
Default sampler rate	A default sampling rate for sampled flows for which exact sampling information is missing. min = 1, max = 100000, default = 1
Network path information reporting interval, sec	This parameter controls Network path periodic updates. Once discovered path is detected and reported, if it is still active, it will be reported again after this period of time. min = 120 sec, max = 14400 sec, default = 300 sec
Output method: 0 - syslog (push), 1 - REST (pull), 2 - both	This parameter controls how you want the output from this Package to be available. values {0, 1, 2}, default = 0
Inactivity timeout for network paths, sec	This parameter controls inactivity period after which path is no longer reported. min = 60 sec, max = 14400 sec, default = 300 sec

Network Operations Analytics	This parameter controls network health score reporting interval in seconds. Network health is reported in syslogs with nfc_id = 20180 and nfc_id = 20181. min = 5, max = 600, default = 23
Network Paths Analytics	This parameter controls the internal interval used in detecting paths between pairs of peers. min = 5, max = 600, default = 37
Network Paths Reporter	This parameter controls network path reporting interval in seconds. Network path is reported in syslogs with nfc_id = 20183. min = 5, max = 600, default = 203
VM - VM Hosts Adjacency	This parameter controls VM – VM Host-ToR Adjacency reporting interval in seconds. This information is reported in syslogs with nfc_id = 20184. min = 5, max = 600, default = 41
ToR Switches Discovery	This parameter controls the internal interval in seconds used in discovering ToR switches. min = 5, max = 600, default = 29
Top Tunnels	This parameter controls Top Tunnels reporting interval in seconds. This information is reported in syslogs with nfc_id = 20187. min = 15, max = 600, default = 30
List of vCenter Virtual Machines	This watch list is populated by NFO Updater Agent by connecting to one or several vCenters
List of ToRs management addresses	This watch list helps the Module to determine which devices are Top of the Rack Switches. ToRs detected automatically if list is empty

Click on List of vCenter Virtual Machines to setup integration with vCenters, vShield, and NSX as shown below.



You don't need to restart NetFlow Optimizer or enable/disable the Module in order for the new parameters to take effect.

**Table 5 - Configuration Parameters - List of vCenter Virtual Machines**

Parameter	Description
Auto-update by <b>VMWARE vCenter Virtual Machine monitor</b>	This check box enables NFO Updater agent which queries vCenter and populates the list of Virtual Machines with their VDS and VTEP.
Next time	Next scheduled update time based on the schedule setting below.
Schedule/Expression	Unix cron expression
vCenter server name	IP address or name of vCenter server
vCenter user name/password	vCenter login credentials
Log level	Leave blank or set to DEBUG for troubleshooting
Data Records	This watchlist is  csv list in the following format: VDS IPv4 address, VM IPv4 address, VM host , VTEP IPv4 address, VM name, VM UUID

**Table 6 - Configuration Parameters - List of ToRs management addresses**

Parameter	Description
Data Records	Management IP addresses of network devices which are Top of the Rack Switches.

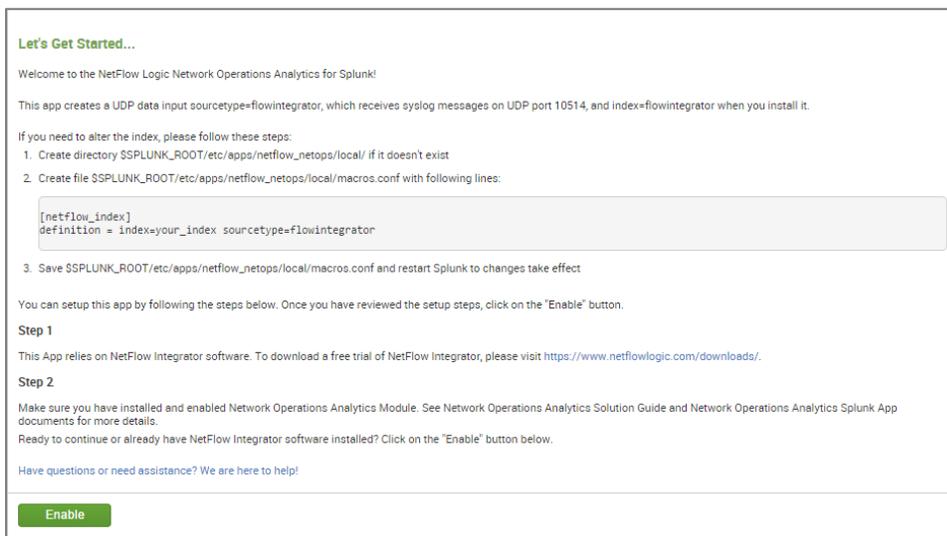
# Network Operations Analytics for Splunk

This section describes installation, configuration, and user instructions of NetFlow Logic Network Operations Analytics for Splunk App (NetOps Splunk App).

## NetOps Splunk App Installation

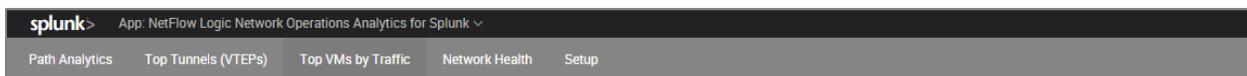
When you install this app it automatically creates the index flowintegrator, and UDP Input listening on port 10514 (NetFlow Optimizer output should be configured to send data there).

Once you install the App you will be presented with this screen:



Press Enable button and launch the App.

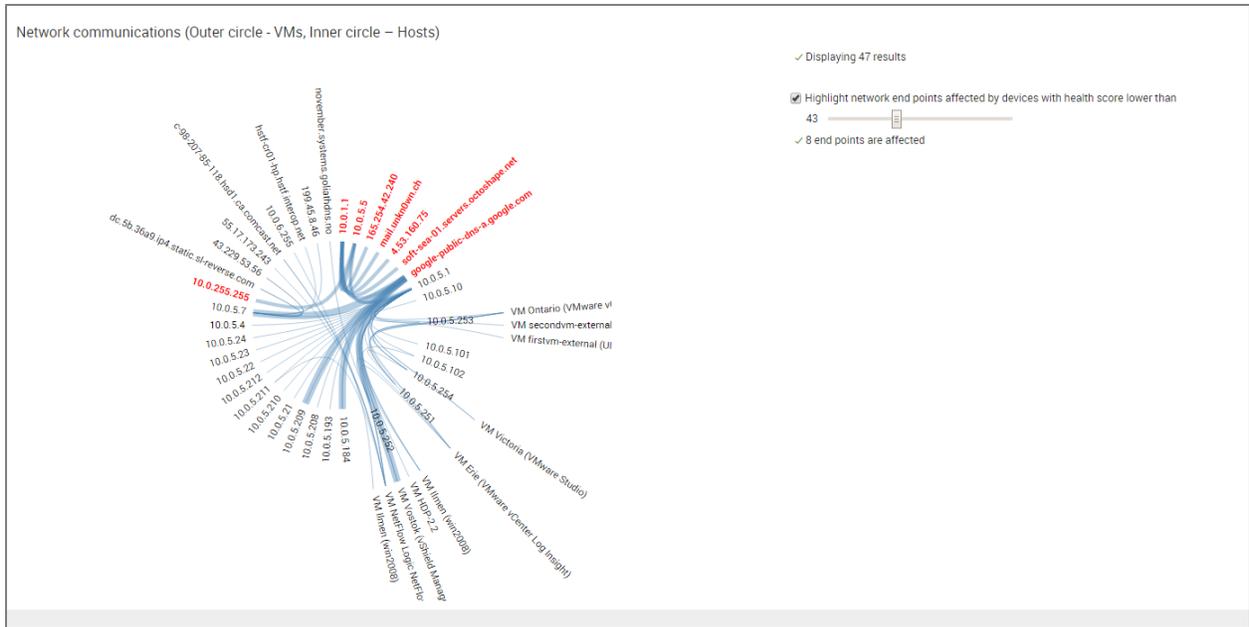
## NetOps Splunk App Dashboards



### Path Analytics Dashboard

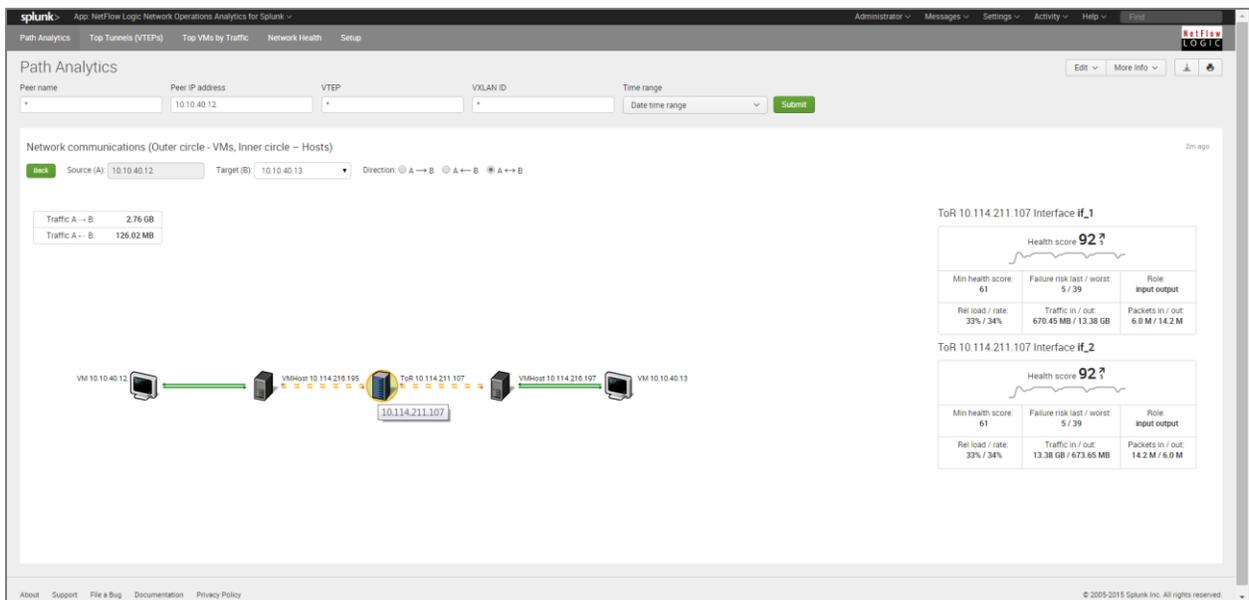
This dashboard enables you to explore network conversations between communicating network peers.

First select the time interval and press <Submit> button. You will see a graph representing all network conversations during the selected time. Physical hosts are grouped in the inner circle, and VMs are grouped in the outer circle.



Narrow down the results by specifying full or partial IP address of communicating peers. Apply additional filters by specifying VTEP IP and/or VXLAN ID. Use the slider to highlight communicating endpoints affected by devices and interfaces with low health score.

Select the first object (A) by clicking on the IP address. Select the second object (B) from a drop down. The Path panel will dynamically populate with the path between the selected objects.



Explore paths by selecting direction: A->B, B->A, or bidirectional A<->B. This panel is a great way to see where the problem is. The color of network nodes indicates the health score of the interfaces in the path of communication between the selected objects.

The table on the right hand side provides the following information about the selected device:

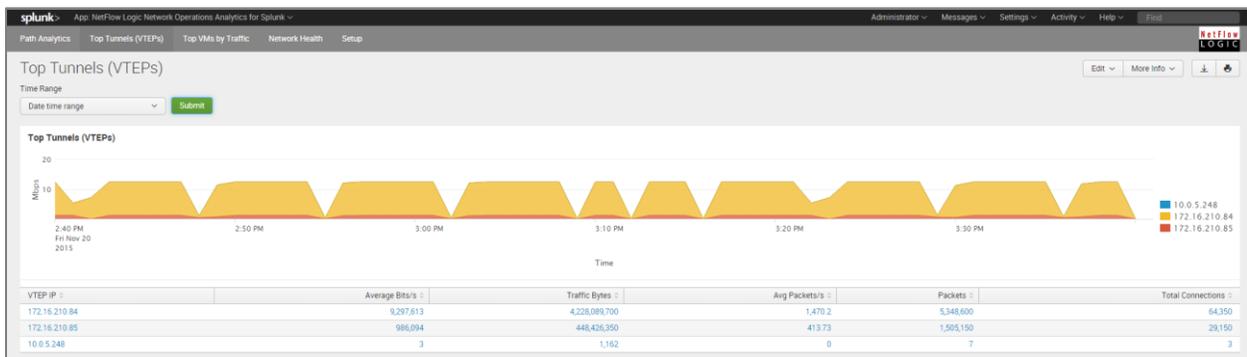
- Device health score – a number between 0 and 100 with the lower scores indicating a device with one or more problematic interfaces
- Device failure risk score - a number between 0 and 100 with the higher scores indicating an interface with a high probability of failure
- Relative traffic load on a most active interface as a percent of its nominal capacity
- Relative packet rate on a most active interface as a percent of a maximal packet rate sustainable at a current average packet size

For interfaces relaying a traced communication the following information is presented:

- Interface health score
- Interface failure risk score
- Relative traffic load on this interface as a percent of its nominal capacity
- Relative packet rate on this interface as a percent of a maximal packet rate sustainable at a current average packet size
- A total number of bytes passed in each direction through this interface over a selected time interval
- A total number of packets passed in each direction through this interface over a selected time interval

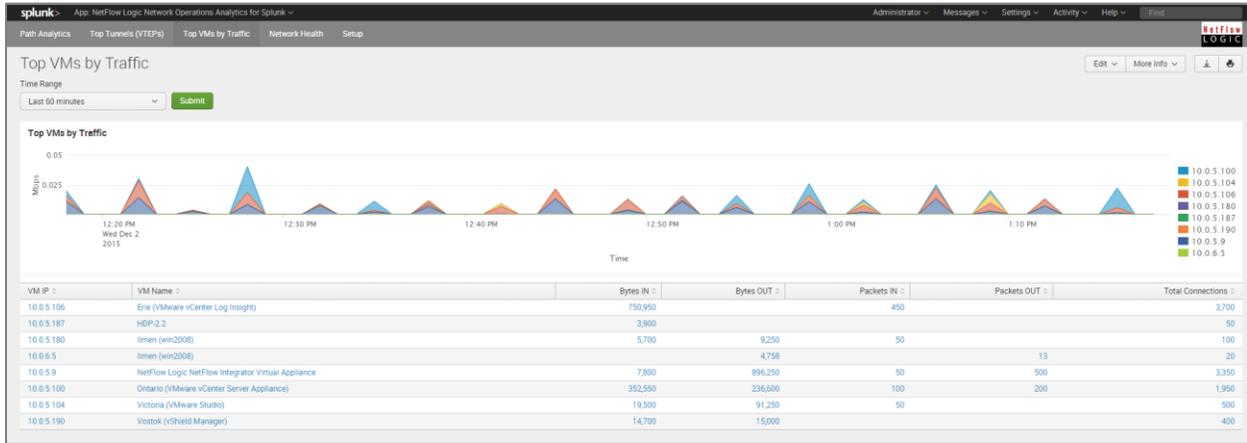
## Top Tunnels (VTEPs) Dashboard

This dashboard shows top VTEPs by traffic over selected period of time. Drill down functionality allows to view individual VM to VM conversations going through selected VTEP. You can drill down further by selecting VM to VM conversation to see the path of communication between peers, which shows VM hosts where virtual traffic is encapsulated, and how it traverses the physical network, including device interface health score, for efficient troubleshooting and planning.



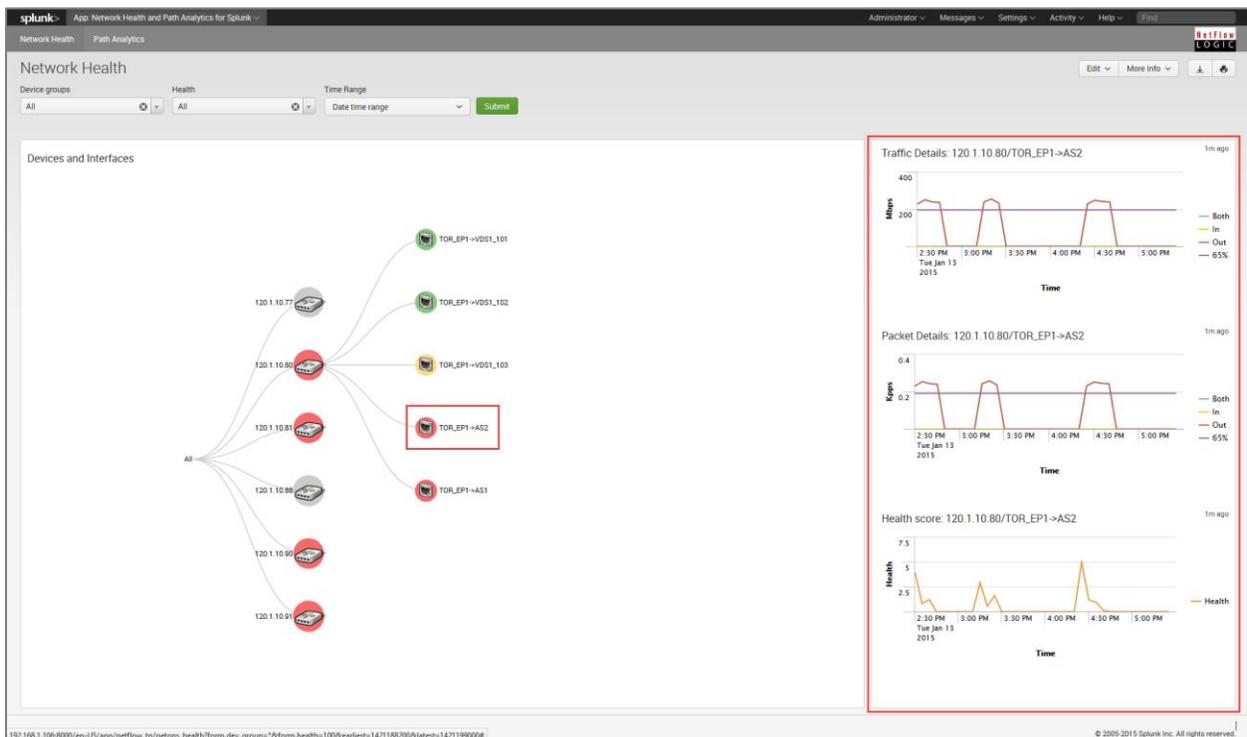
## Top VMs by Traffic Dashboard

This dashboard shows top VMs by traffic over selected period of time. Drill down to path to see VM to VM conversations over physical network.

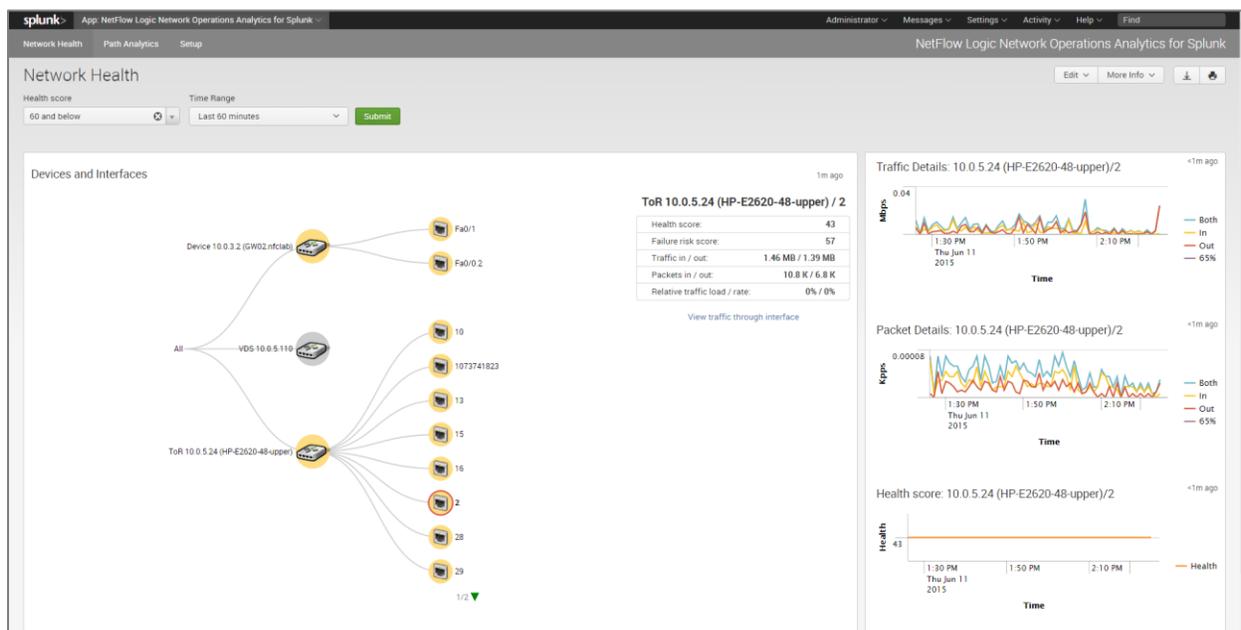


## Network Health Dashboard

In this dashboard the network devices comprising the physical and the virtual parts of the network are presented in a tree view. Icon colors indicate NFO health score of each device based on its interface with the lowest score. NFO health score is a number between 0 and 100 with lower scores corresponding to a problematic device. On this dashboard green icons indicate healthy devices while yellow icons identify devices reporting some abnormal conditions. Red icons indicate that one or more interfaces on a device are not functioning properly. Grey icons represent virtual network devices, such as VMware VDS. Their health score is not calculated.



When you click on a network device node, the tree expands showing all interfaces of the selected device, again color-coding them according to their health score. When you click on an interface, the drill-down panels on the right show the traffic details for the selected interface.



The table on the right hand side provides the following information about the selected interface:

- Interface health score – a number between 0 and 100 with the lower scores indicating a problematic interface
- Interface failure risk score - a number between 0 and 100 with the higher scores indicating an interface with a high probability of failure
- Ingress and egress traffic rate
- Ingress and egress packet rate
- Relative traffic load on the interface as a percent of its nominal capacity
- Relative packet rate as a percent of a maximal packet rate sustainable at a current average packet size

Three drilldown panels on far right show network traffic characteristics and interface wellbeing scores over a selected time interval. A 65% threshold mark displayed on the historic charts helps identifying time intervals when network paths served by this interface are in a high risk of failure territory.

From this dashboard you can jump to **Path Analytics** dashboard to view all network conversations going through the selected interface.

## Start Your Trial

Download NetFlow Optimizer and NetOps components and register to receive your 60-day evaluation license (<https://www.netflowlogic.com/download/>).