




 eBOOK

# How to Simplify Monitoring for Complex Network Devices

by Destiny Bertucci

# Table of Contents

|                                                         |           |
|---------------------------------------------------------|-----------|
| <b>INTRODUCTION</b>                                     | <b>3</b>  |
| Is This eBook For You?                                  | 4         |
| <b>1 CHAPTER ONE</b>                                    | <b>5</b>  |
| Network Insight for Monitoring and Managing ASAs        | 7         |
| Failover Monitoring and VPN Tunnel Visibility           | 9         |
| Access Control Lists (ACLs)                             | 12        |
| Network Insight for Monitoring and Managing ACLs        | 13        |
| <b>2 CHAPTER TWO</b>                                    | <b>16</b> |
| Complexity of Monitoring F5 BIG-IP Load Balancers       | 16        |
| Network Insight for Monitoring F5 BIG-IP Load Balancers | 17        |
| Balancing Environment                                   | 18        |
| View Component Details                                  | 19        |
| DNS Resolution by Service                               | 20        |
| <b>3 CHAPTER THREE</b>                                  | <b>21</b> |
| Network Insight for Nexus VDC Capable Devices           | 23        |
| Virtual Device Contexts                                 | 24        |
| Additional Resources                                    | 26        |
| Dedication                                              | 27        |
| About The Author                                        | 28        |
| About SolarWinds                                        | 29        |

# Visualization of Complex Technologies

## INTRODUCTION

*"Once upon a time, there was a little network..."*

If you have ever been tasked with monitoring and managing Cisco® ASAs, F5® load balancers, or Cisco Nexus® devices, you understand this task could be challenging. Welcome to a much-needed safe zone where I plan on venting my frustrations and offering ways to overcome common issues of monitoring and managing these types of devices. Don't worry—this is not going to be a whiney or unhelpful manifesto about how virtual port channels (vPCs) or access control lists (ACLs) can be hard, and nothing can make it better. Instead, I am going to discuss things like the impact and insight within a visualization of connections with your load balancers, ACL management, and remote access VPN troubleshooting, then how you can quickly see the health of your vPCs to give you a solid understanding of the complications that tend to arise. I'll share ways you can fix (or at least improve) the situation with a little bit of insight from your monitoring and management software.

Writing this eBook has been long overdue, and I hope it presents new ideas for you as well. IT management teams have been multi-threading themselves on researching and solving issues after they arise. Chasing data and flow issues can be like finding a needle in the haystack. On top of this, load balancers have to stay healthy and aware of the services, global traffic managers, local traffic managers, virtual servers, pools, and pool members. That adds up to a lot of moving parts that need representation so you can instantly see what is not working in this tower of complexity!

Then, although we (meaning any network/security individual or team) deal with security, we do not have the opportunity to visually see what is secured. It can be incredibly frustrating running around opening this port or shutting down that port for a new application without being able to address security concerns or see the whole picture. In that situation, it is difficult to know if these new ACL network routing changes are creating other issues. It makes you feel like you are always putting out fires without ever discovering what started them.

We have entered an incredible era of security breakthroughs. We have a better handle on our networks because we are able to control them with adaptive devices based on rules of our own design. It has never been easier to acknowledge changes and verify that these changes are not causing problems. Monitoring baselines for multi-vendor devices allow us to see our environment's big picture, and how these devices and their configurations are affecting traffic or services.

Now bring to light the awesome power of vPCs from Cisco Nexus® devices, and we have quickly gone from a simple switch to connected switches with vPCs bundling multiple ports that appear as one. This is just the technology of today—can you imagine tomorrow? We need to be able to see if the vPC is fully up and functioning so we are ensuring our redundancy and not tracking down bandwidth issues. That's what it's designed for, right? Gone are the days of being left in the dark, only able to monitor up/down status.

The goal of this eBook is to provide you with a deeper understanding of today's monitoring needs and explain why we should be demanding more from our monitoring and managing software solutions. The power of Network Insight™ is the integration of modules to work together and provide insight. These modules are SolarWinds® Network Performance Monitor (NPM), Network Configurations Manager (NCM), NetFlow, and User Device Tracker (UDT).

## IS THIS EBOOK FOR YOU?

I do not want to waste your time reading this if it is above or below your level of understanding. To save you any further unnecessary reading, let me lay down some basic questions.

- » **If you do not** manage or monitor ASAs, Nexus, or F5® load balancers, this book is probably not for you. However, if you plan on being more involved with these types of devices, this could be the book you've been looking for.
- » **If you have** a CCIE® Security, CCDE® Design, or networking certification, this book can help you to look at monitoring in a whole new way. Then you can guide your junior colleagues and help them to understand these technologies and bridge gaps in concept.
- » **If you are** monitoring and managing ASAs, load balancers, and Nexus devices with third-party software, I believe this book is just right for you. You will be able to identify with (or at least understand) the frustrations I am describing and appreciate the concepts I share within.

# Chapter One

## Cisco Adaptive Security Appliance (ASA)

ASAs are Cisco next-generation, all-in-one security devices used by businesses all over the world. The need for enhanced security and data protection across the enterprise has only made these devices more critical to helping protect businesses from security breaches.

### Complexity of monitoring ASAs

To really dig into this topic, I have to break the device down to its component areas. Managing these devices is like a crash course in security. A lot of these devices are implemented by network engineers who have not fully jumped in to the security realm. That's not a huge problem, but being aware of the issues in this chapter will help you have a better grasp of the bigger picture.

### ASA firewall management

Packet filtering allows us to focus on protecting the inside network, demilitarized zones (DMZs), and the outside network. The device inspects all traffic that passes through it and uses access control lists (ACLs) to drop unwanted or unknown traffic. Each ACL contains an element that permits or denies statements, often known as an access control entry (ACE). The device can classify packets through inspection all the way up to Layer 4 headers.

The problem is that these statements can sometimes never be hit, preventing traffic from going to the right destination, or worse, not blocking traffic to certain destinations. Hit counts are when we know the rule has been represented and the traffic has been verified to the rule. If it is not hit, that means that the rules in front of the non-hit rule are not allowing traffic to continue.

Firewalls can have trillions of hits per second. Monitoring and managing these should be a priority since they are evaluating your network traffic at a high rate of speed. We will dive more into firewalls in the next chapter.

IP routing identifies which interface and gateway should be used to forward packets based on their destinations. These routing decisions can be made via dynamic routing protocols like routing information protocol (RIP), open shortest path first (OSPF), and static routes. Devices like Cisco ASAs have the ability to configure up to three default routes for traffic load-balancing, and should all point to the same interface.

Authentication, authorization, and accounting (AAA) gives access control to network devices and are vital for business infrastructures. They provide an extra layer of security support for managing and monitoring network devices. Yes, now we are bringing in local or external servers for authentication that use AAA authentication protocols like RADIUS and TACACS+.

High availability has moved from a “nice-to-have” to a “must-have” with today’s e-commerce reliance on LANs and WANs. Connectivity losses can have huge financial impacts. Large enterprises could potentially lose a million dollars in revenue per hour. Think about huge distribution centers for businesses and public sectors. Take the [Amazon® typo incident](#) for a quick “uh oh” moment. This mistake cost the company \$150 million for three hours of service downtime. If you would like to calculate your own downtime for critical services and outages, use the following formula:

Cost of downtime per day = lost revenue + lost productivity + cost to recover + cost of intangibles (loss in customers)

Then you have to take into consideration whether you have a 24/7 business or one that runs on a more traditional 8 a.m. - 5 p.m. schedule, because that will increase or decrease your hourly rate. Luckily, ASAs have failover capability, which allows you to have a disaster recovery device in place in case one fails or goes offline. This helps you stay online and continue business while you attend to the downed or dysfunctional device. Failover is vital. You do not want an open gateway to your network if the active device goes down.

### **ASA IDS/IPS management**

Intrusion Protection Systems (IPS) that include deep packet inspection (hence the “adaptive” portion of the ASA abbreviation) and IP auditing with basic signatures can be provided. This lets you decide if you want to use inline or promiscuous IPS mode. In inline mode, traffic passes through the firewall policies before it is sent to the Cisco FirePOWER® module. This helps by only inspecting permitted traffic based on the policies. Once in the FirePOWER module, if it is determined to be malicious based on defined security policies, it sends back the malicious verdict to the ASA, which will then block the traffic going forward. In promiscuous mode, copied packet traffic, which is defined in the service policy, is sent to the Cisco ASA FirePOWER module. However, if it is only in monitoring mode and if anything is deemed malicious, it can alert the administrator, but will not block the traffic. This is where monitoring health statistics comes into play. You can decide the amount of traffic and the effects of this traffic being inline or promiscuous for the increase in inspection.

## ASA VPN concentrator management

VPN site-to-site uses the IPSec protocol. Since managing this protocol setup is within the capabilities of an ASA, it is important to note that **IPSec can use Internet Key Exchange (IKE) for key management and tunnel negotiation**. Why is that in bold? It's simple: IKE uses different combinations of Phase 1 (IKE\_SA phase defined in the key exchange) and Phase 2 (CHILD\_SA phase defined within the data policy) attributes that will be negotiated between the peers. If **ANY** one of these attributes is misconfigured, the IPSec tunnel will fail. Because of the complexity of these tunnels and the required precision of configuration, debug commands like **"show crypto isakmp sa detail"** are used to verify these phases. We also have to know where to find these details quickly, to be able to better manage and troubleshoot if and when issues arise.

Remote-access VPNs open a whole new can of worms. Not only do we need to keep up with our network, now we have to verify the tunnels' stability via remote devices. A saving grace is that the steps and attention to detail required are similar to those of site-to-site.

Multiple security contexts are needed for both split tunneling and better traffic management. Contexts give you the ability to partition or separate an ASA into multiple virtual devices. This allows you to have multiple security contexts within your device, which means they will be included in the configuration file. These are separate security plans managed through one device. Split tunneling, which only sends corporate encrypted traffic to the ASA and sends things like email and messaging out through the internet, is a way of load-balancing traffic and preventing your ASA from being a complete bottleneck by having an intensive bandwidth load.

All of these capabilities can present a lack of monitoring situations. There are multiple features of ASA devices that we need to see and manage, but due to their complexity, it has been a bit of a mess to gather all of this information in one location. However, simple network management protocols and a show command to download a configuration file is not always the best or fastest way to approach these complex devices and their uses. The next section will show how we can untangle this mess of monitoring and management, and finally allow us to focus on performance and security.

## NETWORK INSIGHT FOR MONITORING AND MANAGING ASAs

Before managing and monitoring network devices, you should back up configuration files. These files normally contain the configuration for one device and how it should be used. However, there are devices that have what they call multi-context or security context configuration files. Just like it sounds, they have different configurations on one device that may tell the network traffic to behave or be directed to different locations based on the context that traffic adheres to. That being said, these multi-context configuration files have been an issue for recovery and comparisons. Recovery is when you have multiple contexts (segments) that have to be

accurately placed, which has been a challenge for network configuration managers. The network traffic needs to follow the correct context of the configuration so it is routed, secured, or dropped according to the context that has been created.

Comparing these configurations between like devices has always been a challenge when the order of placed ACEs have actual meaning, and security contexts can be seen as a configuration itself. As you read earlier, there can be multiple security contexts, often in companies like service providers, large enterprises, or governments that want to keep departments completely separate. The reason they exist seems obvious when you are managing traffic through a device to allow proper traffic in designated areas. Due to various complexities, the ability to download, store, compare, and restore specific context configs to a device for monitoring and management continues to be a challenge. The good news is that visibility into this ecosystem has never been clearer than it is now, with mature monitoring solutions providing the ability to monitor these multi-context configurations.

Combining monitoring elements like receiving real-time change notifications, compliance reports (based on your security needs), analyzing hit counts so you can simplify your ACLs, and being able to compare like devices for network standardization is a game changer. Think about the detection, prevention, analysis, and response (DPAR) cycle, and how the multiple abilities of the Cisco ASA tap into every aspect of it. When we are faced with the burdens of maintaining the information security goals of confidentiality, integrity, and availability, we cannot afford to monitor and manage these devices less than completely. Combining network performance metrics, such as overall platform health, load-balancing status, and failover, allows more time to troubleshoot. But you still need more to gain full visibility into the vast multitude of capabilities that ASA devices deliver. The following items help allow VPN tunnels with Cisco ASA monitoring and management to have a true baseline for better troubleshooting.

- » NetFlow and NSEL ingress and egress traffic with CBQoS information correlated to your interface bandwidth
- » General health statistics for capacity
- » VPN tunnel visibility (LAN-to-LAN)
- » Remote access VPN visibility
- » Failover monitoring
- » Connection counts
- » Hit counts
- » Interface monitoring with security levels
- » Reporting site-to-site and remote-to-site
- » Alerting for phases of failure with VPN tunnels



## FAILOVER MONITORING AND VPN TUNNEL VISIBILITY

SolarWinds® Network Configuration Manager (NCM) allows you to see how your configurations are working within your network devices. Validating your VPN tunnels is great, but adding the capability to store backups for network standardization, while also being able to alert on any changes made to these configurations, is security gold.

While troubleshooting VPN, firewall, or IDS/IPS issues, we should all strive to find one solution that provides a view into these details. For example, I want to quickly validate a failover's active and passive state within my ASAs while using a managing and monitoring device that lets me see a connection session's history.

**Node Details for ASA - EAST-FW-A - Platform**

**Details**

**EAST-FW-A**  
 ADMIN, Cisco ASA 5510SC  
 running 9.1(6)1  
 Routed mode  
 Polled at 10.1.39.51

Other details

**High Availability**

HA type: Active/Active; Last failover: 3 months ago

Diagram showing failover status:


- EAST-FW-B: Standby Secondary
- Standby ready
- Config synced
- Connection state synced
- this node: Active Primary

**Node Details for ASA - EAST-FW-A - Site-to-Site VPN**

**Overview**

**FILTER RESULTS**


| IP Address      | Status                                 | Up Status | MD5 hashing | In kbps        | Out kbps       | Time since last update |
|-----------------|----------------------------------------|-----------|-------------|----------------|----------------|------------------------|
| 10.131.71.33, 2 | Down Status, Failing at Phase 1        |           |             |                |                | 73 days since Feb 25   |
| 99.227.52.87    | Up Status, DES encryption, MD5 hashing |           |             | 645.87 kbps in | 34.57 kbps out | 1 day since May 9      |
| 96.51.134.55    | Down Status, Failing at Phase 1        |           |             |                |                | 1 day since May 9      |
| 68.146.88.84    | Down Status, Failing at Phase 1        |           |             |                |                | 1 day since May 9      |
| 66.222.188.97   | Down Status, Failing at Phase 1        |           |             |                |                | 1 day since May 9      |



I want the ability to drill down into site-to-site or even remote access VPNs to troubleshoot connection issues. I also want to see the overall performance of the device via CPU, memory, and traffic flow to any increases that may have occurred. This would let me forecast capacity planning for expansions to my ASA needs. Hit counts are vital to managing your ACL list, which can sometimes get out of control. Showing relationships; analyzing ACL configs; and showcasing redundant, shadowed, and unused rules is quick and easy, which is great because that means I don't have to personally go to each device and run show commands.

## NETFLOW AND HEALTH STATISTICS

When you combine NetFlow collected data, network performance fault monitoring, configuration management, and vulnerability detections with many other metrics, we finally have a bigger picture that provides complete ASA infrastructure monitoring and management. Full integration of these metrics gives you proactive awareness of your ASAs with fluid, active monitoring of traffic flow and the way your network, device, and users are affected.



Vulnerability reports automatically check for vulnerabilities against your platform's firmware. When you add firmware bulk upgrades, you save time and reduce your risk of exposure. Make sure to monitor the full realm of the device, from locations to labs to test, for potential issues that may arise after the upgrade. Forecasting time and impact during maintenance helps teams better prepare for and prevent issues that may cause service availability to falter.

## ACCESS CONTROL LISTS (ACLs)

For a more technical definition of access control lists, the image below provides a list of rules that specify which traffic should be granted or denied access to objects and locations. Each rule then provides an action for the traffic that matches it, to either drop or allow it. Traffic is checked in the order of the rule placement. If traffic matches a rule, then the traffic is either permitted or denied and will not pass to the next rule. If the traffic does not match the first rule, then the packets will continue to be processed against the subsequent rules until a match is found. ASAs have an implicit deny at the end of each ACL; any traffic that is not matched is dropped. This packet filtering allows you to determine the path a packet will take from then on. Rules can be based upon IP or MAC ACLs, so you have options to better secure and use filtering for your business needs.

## COMPLEXITY OF ACL MANAGEMENT

First, just as a baseline for practical usage, ACLs allow you to be the pass/no pass general of the IT department. If it takes me twice as long as it should to gather data about a problem, I am wasting valuable time that would be better spent resolving the issue. Meanwhile, service availability may be compromised. ACLs, and the access control entries (ACEs) that go along with them, can become a chore of managing new applications and access level changes happening weekly, if not daily. ACEs specify the decisions that the ACL must perform. I will focus mainly on Cisco ASA devices because I am more familiar with them and they generally have a bigger market share.

As with most ASA types of devices, you will find they can have limitations. Cisco ASA limits include only 128 rules per ACL, or no more than 10,000 ACLs across all the ACLs in one virtual Ethernet module (VEM). You might think this is great because they limit the size of your configuration building. Not exactly—these rules have a hierarchical ordering, and they work by traversing the rules until each packet is either permitted or officially denied. To accurately configure these rules, you need to understand your traffic flow and business needs. For example, you may have a permit rule for your ACLs that grants access to a mail server from only one

subnet. Then you have all other traffic as a deny rule—all very basic. But then you are told there is a new application that only certain IPs need to access and they will be coming through this Cisco ASA. I now have to adjust my ACLs to allow specific traffic to access the application before the subnetted traffic. I could easily mistype and place it below the deny rule and have connectivity issues. So, there are multiple ways to cause issues with rules and the way they are arranged and managed. This presents the problem of users either having access to data they shouldn't, or the inability to access data they should be able to reach. Or both. Both is always a hoot.

Of course, you will always hear about the denied access that interferes with a user's work, but when it comes to accidental open access, that may be a sleeping giant, one that can be exposed with damaging repercussions. The upshot is that it is absolutely critical to get your ACLs correctly placed. You must constantly monitor the EFFECTIVE rules (how they actually work) to help ensure that you are getting what you think you designed. Troubleshooting issues can be a headache to resolve. You or your team will constantly gather informational metrics from the device and then manually correlate the problems. I literally had a security pocket book I used to keep with me that had my most-used commands for troubleshooting these devices. Here are a couple of entries.

#### **show access-list(name)**

```
hostname# show access-list outside_access_in
access-list outside_access_in; 3 elements; name hash: 0x6892a938
access-list outside_access_in line 1 extended permit ip 10.2.2.0
255.255.255.0 any
(hitcnt=0) 0xcc48b55c
access-list outside_access_in line 2 extended permit ip host
2001:DB8::0DB8:800:200C:417A any (hitcnt=0) 0x79797f94
access-list outside_access_in line 3 extended permit ip user-group
LOCAL\\usergroup any any (hitcnt=0) 0xb0f5b1e1
```

#### **Show Blocks**

```
testASA#show blocks
SIZE MAX LOW CNT
4 1600 1597 1600
80 400 399 400
```

As my familiarity with the commands grew, I found that being able to quickly pull massive amounts of information and make updates led to longer troubleshooting. I had to gather all of these complex rules and know which ones applied to the incoming traffic, and then correlate these to find the correct logical to physical name of the ACLs to their interfaces. If you find yourself in the same situation, you have probably thought—as I did—that it must be easier. Spoiler alert: It is.

Constantly shifting business and user needs create security holes and misalignment. These needs include things like users requiring access to different sites and areas within the internal network, setting up new internal applications that need to be accessed by the outside world, and demanding technical solutions for what are ultimately management issues (“I need you to keep Bob in accounting off Netflix®, but I still want to watch my show during lunch”). ACEs can stack up and be placed in numerical order so they can be added when needed, without being hit (the rules are not being tested since the previous ones are blocking them from getting to the new rule that’s been applied). Or worse, there could be an incorrect rule placed too high, allowing the wrong traffic through instead of having it filtered, denied, or sent elsewhere. Our ACLs can hinder our service availability and leave us scratching our heads if we do not have the whole picture of how they are being used.

## NETWORK INSIGHT FOR MONITORING AND MANAGING ACLs

Monitoring and managing firewalls and ASAs, up to this point, has been overly complex. However, now we finally have solutions that provide correlated data so we can truly take control. Have you ever wondered how your traffic is flowing and being directed? Do you ever want to see your logical and physical names of interfaces with their ACLs assigned, all within a single screen and from one monitoring solution? Yeah, me neither. (Note the sarcasm in my tone.) SolarWinds has the ability to take proven products and integrate their intelligence into a behemoth of ACL monitoring/management capabilities. Brace yourselves for this: backing up your ACLs while offering change notification is here! Now I can compare my ACLs to other like devices to validate and standardize the network. Also, if my ASA device goes down, I can now know if the standby has successfully changed to active and not worry because an alert let me know that a failover occurred with my ASAs. In case a hardware failure caused the failover and I need a new device replacement, I can restore the device configuration quickly and accurately. All of my ACLs and hard work make it appear like the failure never happened.

While ACLs on their own present a lot of challenges, security contexts—which present multiple, separate sets of ACLs that may or may not be working together—are like going from Tic Tac Toe to 3-D chess. Service providers who offer services to different customers use these security contexts. Multiple users in different companies need their own secure transactions that apply only to the services being provided. For example, say one is a large enterprise or college that keeps departments separate via different security contexts within one ASA. That said, managing ACLs within their specific context is crucial because they need precision to filter packets properly. If I can view the impact of an ACL change and troubleshoot complex ACLs within and across ASAs, I am in control. SolarWinds Network Automation Manager (NAM) delivers one single-pane-of-glass solution with the capabilities of NetFlow, configuration management, events, and platform health, allowing us to optimize ACLs by eliminating redundancy and shadow rules. Being able to snapshot and version ACL configs keeps these accurate and orderly for training, network standardization, and bare bones replacement.

## Differencing

Highlighting the differences between the same devices, such as Cisco ASAs, allows me to better manage my devices and protect the network from possible mistakes. Using compliance reports with my business’s security protocols allows users to quickly audit their devices and help ensure proper configuration in large environments. Auto-remediation allows for quick placement of dedicated ACLs that have to be on all like ASA devices.

The screenshot shows the 'Access List Rules Diff' page in SolarWinds. It compares two ACLs: 'CSEA110\_acl on EAST-FW-A' and '110-140\_in on Same Node'. The left pane shows the first ACL with four rules: permit icmp any4 any4, permit udp host 61.238.9.3 host 59.199.0.129 eq ntp, permit tcp 73.133.166.0 255.255.255.0 host 187.9.244.3 eq telnet, and permit icmp any6 any6. The right pane shows the second ACL with four rules: permit tcp 73.133.140.0 255.255.255.0 object-group ipat\_srvs object-group top\_17278, permit object-group DNS any4 object-group RNI INSIDE, permit object-group NTP\_Y1 any4 object-group NTP\_Y2, and permit object-group TWR\_e.s object-group PLLAstedge any4. A search bar is visible at the top right of the comparison area.

## Enumeration

The translation of interfaces from a physical name to a logical name (“outside,” for example) allows you to see the interfaces and the ACLs that are residing on them from one location. Combine this with NetFlow traffic for the interfaces, and you have a baseline that allows you to see possible intrusions or traffic changes within your network.

The screenshot shows the 'Node Details for ASA - EAST-FW-A - Interfaces' page in SolarWinds. It displays a table of all interfaces on this node. The table has columns for interface name, status, speed, errors in last hour, and traffic in/out. The interfaces listed are CPLANE, GIGABITETHERNET0/2, GIGABITETHERNET0/3, and GIGABITETHERNET0/4. The CPLANE interface is highlighted in green.

| Interface          | Status                           | Speed          | Errors in last hour | Traffic                         |
|--------------------|----------------------------------|----------------|---------------------|---------------------------------|
| CPLANE             | 10.1.39.57                       | 100 sec. level | 0                   | 0 bps in, 361.18 Mbps out       |
| GIGABITETHERNET0/2 | 10.1.39.51, 10.199.2.4 (standby) | 1 sec. level   | 0                   | 915.98 Mbps in, 2.31 Gbps out   |
| GIGABITETHERNET0/3 | 10.1.39.53, 10.199.2.4 (standby) | 100 sec. level | 0                   | 579.50 Mbps in, 650.74 Mbps out |
| GIGABITETHERNET0/4 |                                  | 1              | 0                   | 676.22 Mbps in, 351.73 Mbps out |

## Hit Counts

Analyzing hit counts can reduce the number of unnecessary rules. This can also be used to troubleshoot rules that are not being used because traffic is being denied or sent elsewhere.

Improve operational efficiency and optimize ACLs by automatically identifying shadow rules. When you have the ability to filter, search, and view ACLs through rule browsers, you can find and resolve issues quickly.

The screenshot shows the SolarWinds interface for managing ACL rules. The top navigation bar includes 'solarwinds', 'MY DASHBOARDS', 'ALERTS & ACTIVITY', 'REPORTS', and 'SETTINGS'. The main section is titled 'Rules Of This Access List'. On the left, there is a 'FILTER RESULTS' sidebar with options to 'EDIT FILTER PROPERTIES'. The filter criteria are: 'Rule Type (1)' with 'permit' selected (33 hits), 'Source (11)', and 'Destination (16)' with various object-groups selected. The main area displays a list of rules with their hit counts. A 'Hit count' dropdown is visible above the list. The rules are sorted by hit count in descending order.

| Hit count | Rule                                                  |
|-----------|-------------------------------------------------------|
| 35        | remark TTOW281698 PLL Ws c to PLL WsnsS01             |
| 47        | remark PLL Ws to LVOZ-Pt-As                           |
| 44        | remark TTOW281698 PLL-Ws p1 smtp to hvhc test1        |
| 42        | remark TTOW281698 PLL-WsnsEML01-p1 to AD svrs         |
| 38        | remark TTOW281698 PLL Ws c to LVO WsnsS01             |
| 36        | remark TTOW281698 PLL Ws c to LVO WsnsS01             |
| 8         | remark TTOW181559 CCN5618                             |
| 46        | remark TTOW302042 PLL-Ws p1 to LVO-Ws p1 6671 6643    |
| 15        | remark TTOW281698 PLL-Ws c smtp to hvhc test1         |
| 11        | remark TTOW253204 PLL-STEDGE smtp to smtp.ELTONvn.com |

Navigation: < 1 2 3 4 5 > 1-10 of 47

The bigger picture is integration for the win. When we integrate these metrics from different locations and fully understand devices like Cisco ASAs, we improve service availability, decrease human error, and optimize security. Visualizing the inner workings of the devices provides greater awareness in monitoring and managing ACLs, which, in turn, immediately increases our security. Managing and monitoring ACLs doesn't have to be a headache caused by multiple tools. Impact awareness and the ability to see traffic flow helps you make better decisions while cleaning up redundant rules.

We have crossed over into a new and exciting era of monitoring. Integration without barriers allows us to see technology and present it to us as a solid solution. I am impressed by that. I have always white-boarded these ideas to make them actually make sense, and now we have monitoring software that brings these ideas to life. It makes you wonder how we were ever able to breathe in the past.

## Chapter Two

### F5-BIG-IP Load Balancers

The concept of a load balancer seems to be self-explanatory. However, a lot of moving parts are required to make this happen successfully and to validate the reason you have these devices implemented within your environments. They allow you to increase the capacity of concurrent users, which increases your reliability of applications.

### COMPLEXITY OF MONITORING F5 BIG-IP LOAD BALANCERS

Network managers are challenged to maintain near 100% network uptime (I know it's stressful). Ensuring that business applications and services are fast, secure, and available is our primary responsibility and priority. There was a void here that the F5® BIG-IP® family was able to fill for network and system engineers. They offered new intelligence and more complexity to applications in these areas. The question I had was, how do I verify the load balancers are configured accurately and correct traffic is flowing with the appropriate balancing of user connections?

Now, there are ways of achieving this and, in retrospect, of upping our troubleshooting game. However, when you're logging in to individual components one by one, trying to draw out a picture mentally or physically on paper of how everything is related and checking each component's health, it feels like yet another fire drill to find that needle in the haystack.

Through all of the fire drills and the human multi-threading of hands to keyboards, we are still supposed to keep it business as usual. Many have tried, and again will try, to address these concerns by opting to buy more bandwidth, more servers, cloud installations, and even rewriting their applications. (To be fair, some applications need revision for performance, but that is a blog post and not an eBook topic—moving on!) Unfortunately, this can add up in cost, resulting in only partial fixes with the potential to create a more substantial complex infrastructure that you will have to maintain.

Furthermore, changing business requirements means that the network is rarely stable. There are always changes, upgrades, and expansions to perform. Therefore, accurate real-time visibility into the stability and performance of the network is necessary and critical to maintaining business continuity. Network managers employ a broad range of management tools, from home-grown solutions to multi-million dollar enterprise management frameworks, to gain visibility and control over their network.



Typically, home-grown solutions do not scale well and are easily outgrown by the companies that use them. Consequently, many enterprise management frameworks have been the focal point of larger customers for over 10 years, but these products can take months to implement and dedicated consulting staff to maintain. With their considerable expense to purchase and maintain, network managers are typically disappointed with the return these investments provide.

The next section will show how comprehensive views that access all components of the application delivery help to understand relationships and dependencies between component systems. Identify problems before your users begin calling, and address them to keep your time where it is needed most and the users away from idle hands.

## NETWORK INSIGHT FOR MONITORING F5 BIG-IP LOAD BALANCERS

SolarWinds® Network Performance Monitor's Network Insight feature provides comprehensive monitoring of F5 BIG-IP DNS (formerly BIG-IP Global Traffic Manager™) and F5 BIG-IP Local Traffic Manager™ to give you the insight you need to keep your most important services running smoothly. I know you are probably wondering why this is a big deal. How is this going to help me in my everyday IT life? Keep reading and prepare to understand once you see these overview pages of awesome!

The monitoring of health, performance, and availability of application delivery has been complex in the past. The managing of individual components requires you to individually check their health, which is time consuming. Understanding these dependencies can be difficult with the different relationships between component systems. Then you are left with identifying what the actual problem is and where the possible slowness and service outages are occurring.


SolarWinds leaned in and wanted to have a better monitoring environment that would allow you to solve these issues and get your weekends back. Here is how we addressed the issues above.

- » Visualize your entire balancing environment
- » Graphically display relationships and component status
- » View component details in a single page
- » DNS resolution by service

## BALANCING ENVIRONMENT


With the Network Insight feature, you can monitor the health and performance of all components of application delivery, including WideIPs, virtual servers, pool members, and more.

Simply click on one of the status indicators to show relationships or dig into additional details about that component.



## DISPLAY STATUS AND GRAPH RELATIONSHIPS

Easily view the relationships from service to traffic managers, virtual servers, pools, and pool members, along with the detailed status of each component. Use the search function to find specific host nodes or domains. This allows you to quickly segment to each component needed to provide a balanced environment. You can check all health by hovering over each component.




## VIEW COMPONENT DETAILS

Network Insight provides everything you need in a single console.


With F5 BIG-IP Global Traffic Manager, you get a summary of supported services, and F5 BIG-IP Local Traffic Manager, high availability status, DNS resolution by service, and the supported sites and services.

With F5 BIG-IP Local Traffic Manager, you will see a summary of virtual servers, pools, and pool members, and virtual server details that include concurrent connections by virtual server, port, default pool, balancing algorithm, and connections.



## DNS RESOLUTION BY SERVICE

Verify your DNS hit count per service within your F5 environment.



Managing your application deliverance doesn't have to be a complex situation. Using the correct monitoring tools to help you to visualize your environment while actively monitoring is vital in today's times. The network will be consistently changing and adopting new technology, which means you should be demanding more from your monitoring and management tools, period.

Earlier I mentioned how I'd whiteboard out solutions for ASA issues and setup. You don't even want to see how I used to do this just to round robin four additional SolarWinds web servers. I'm already getting a headache just thinking about those days.

# Chapter Three

## Nexus Data Center Switches

These devices allow you to bundle interfaces from different switches and create virtual port channels (vPCs). It sounds like a super awesome vLAN-type environment, but this concept is way more than just security features with a larger playground. Their ability to use multiple interfaces as vPCs allow you to increase bandwidth and have redundancy. These are modular and fixed port network switches designed for the data center. NX-OS has some high-availability features optimized for high-density 10 Gigabit Ethernet, like running a modular NX-OS firmware/operating system on the fabric.

## COMPLEXITY OF MONITORING NEXUS

When we are discussing devices that run anywhere from \$250,000 to over \$500,000 each, by sheer cost alone we know they are important. In the past, we relied on network engineers to map out their technologies in Visio®, Excel®, or even Notepad to keep track of their intricate designs and configuration details. A junior engineer coming into this environment with that type of documentation would find it hard to understand and very time consuming. Using your team as multi-threaded command line interface drones to validate connections and proper access can be both overwhelming and a costly waste of human time and effort.

In describing these devices, I want to give you a high-level overview of what their capabilities are and the reason for the complexity in design. The Nexus 7000 NX-OS software supports three distinct functional areas that I want to delve into.

1. Virtual device contexts (VDCs)
2. Virtual port channels (vPCs)
3. Port-level access control lists (PCLs) and VLAN-level access control lists (VACLs)

### Virtual Device Contexts (VDCs)

VDCs allow the partitioning of a single physical Nexus 7000 device into multiple logical devices. This logical separation provides the following benefits.

- » Administrative and management separation
- » Change and failure domain isolation from other VDCs
- » Address, VLAN, VRF, and vPC isolation

Once you create your VDC, you would then assign it physical interfaces to downstream devices like Nexus 5000 series. These VDC configurations allow switches to be virtualized at the device level. Each configured VDC presents itself as a unique device to connected users within the framework of that physical switch. The VDC runs as a separate logical entity within


the switch maintaining its own unique set of running software processes, with its own configuration, and managed by a separate administrator.

### Virtual Portal Channels (vPCs)

You can assign these physical interfaces to virtual port channels, which will allow you to increase your bandwidth, redundancy, apply security, and add features depending on your sub-modules. Creating these bonded channels from one Nexus device and pairing them to bonded channels on another physical Nexus device gives you a single circuit, and physical and logical redundancy. As you can see, there are a lot of moving parts to configure these correctly and to create a design that will work completely.

I've talked to customers and their concern was not the setup, as they have had senior engineers set the design up and do not need to change these very often. The problem comes from troubleshooting the interfaces that are bundled to create the vPCs and making sure you know if there was a change from the primary to the secondary.

Failures of vPCs can be hard to figure out based on the sheer amount of components involved in the setup. If we break these down, you have to look at how the setup begins. This will allow you to understand why it's so frustrating to see this technology as a whole. I have included a design to help visualize the capability and use of Nexus and its virtual device context or designed vPCs. The vPCs are showing with the dotted circle wrapping the connections.



From the diagram, we can see the creation of redundancy, which is critical to applications and services to make sure they are able to pass their traffic. Now you are able to see how many parts need to be monitored. Then we have the VDC domains themselves that will include the keepalive heartbeat and the peer-link. The physical diagram shows the connections and how they are bundled to create what appears to devices passing their traffic as the logical diagram.

The logical diagram represents how devices would see the Nexus to send their traffic to their destination. On top of the logical design, you would still have orphan interfaces from the original devices that would need to be monitored. These orphan devices are interfaces that are not bundled or involved with the VDC or vPCs. They can go to their destinations and are not restricted or in use with this feature set.

All of these sometimes thousands of interfaces, we cannot forget, are individually configured. That represents long configuration files and a lot of searching through show commands to validate their use and to manually map their connections. This is where we find IT departments multi-threading individuals in the heat of the moment to track down, troubleshoot, map, and find issues that may be presented.

### Port Level ACLs

- » Security in the form of ACLs becoming more distributed and specific (i.e., closer to the endpoint)
- » How Nexus implements this idea
- » The impact of having to manage thousands of ACLs
- » The impact of having mismatched ACLs in virtual port channels

In the next portion, we will cover how SolarWinds has responded with their visualization and troubleshooting needs for a user with this type of setups within your environment.

## NETWORK INSIGHT FOR NEXUS VDC CAPABLE DEVICES

SolarWinds Network Performance Monitor and SolarWinds Network Configuration Manager's Network Insight feature provides comprehensive monitoring of Cisco Nexus devices and will give you the insight you need to keep your VDC environments healthy, monitored, and managed. Increased visibility allows you to focus on solving common issues with vPCs, VDC access lists, and configuration errors instead of searching for them in high-stress scenarios.


Device complexity is growing across the board. Shouldn't your monitoring of these devices change as well? We should no longer be multi-threading our human IT counterparts for show commands and CLI engagement to create logical mapping of bundled physical ports. We need to know where the issue is immediately when looking at our monitoring so we can then quickly access the configuration, adjust, and resolve the issue before users start contacting us.

SolarWinds once again broke down the barriers of traditional monitoring and jumped right into interacting with the device, which allows for a better monitoring environment. I would rather be proactively warned and have visual mapping to pinpoint where I can focus my time and energy to solve issues, instead of acting like a squirrel gathering up information for a coming together of minds. Let's talk about how we addressed the issues I described earlier in the complexity of Nexus VDC feature-enabled devices.

- » Virtual device contexts (VDCs)
- » Virtual port channels (vPCs)
- » Port-level access control lists (PACLs) and VLAN-level access control lists (VACLs)
- » DNS resolution by service


## VIRTUAL DEVICE CONTEXTS (VDCs)

VDC domains on large environments can be as little as one or as vast as your imagination and device limitations. Before, show commands and Visio diagrams would be used to show which switch had what VDC domains while quickly being able to see the health of the domain and any possible issues to troubleshoot.



## VIRTUAL PORT CHANNELS (VPCs)

Bundled physical ports make up a vPC, as we have discussed previously, and the general mapping of which ones to primary and secondary has usually been a very tedious and manual process. Now we are able to see the connected physical interfaces, which device they are connected to, and their general health.



This is great information right out the gate, but there is more to the view above. Now that I can see which interfaces are applied to each VDC and vPC, I can further investigate. If there is an interface that may be having an issue, we can drill into its context and see just the interface's config snippet. We are then able to view the configs and compare primary to secondary config snippets of these physical port channels.



**vPC Configuration Details**

**Nexus-VPC-2**  
Download time 5/13/2018 5:20 PM (Current)

|                                         |
|-----------------------------------------|
| interface port-channel30                |
| description ADFCHV03                    |
| switchport mode trunk                   |
| switchport trunk allowed vlan 10,16,300 |
| vpc 30                                  |

**Nexus-1**  
Download time 5/13/2018 5:20 PM (Current)

|                                         |
|-----------------------------------------|
| interface port-channel30                |
| description ADFCHV03                    |
| switchport mode trunk                   |
| switchport trunk allowed vlan 10,16,300 |
| vpc 30                                  |

**Member Interfaces**

**interface Ethernet1/35**

|                                         |
|-----------------------------------------|
| description ADFCHV03                    |
| switchport mode trunk                   |
| switchport trunk allowed vlan 10,16,300 |
| channel-group 30 mode active            |

**interface Ethernet1/41**

Now you can see only the interface portion of a config per physical device that you are managing within NCM.

**Interface Config** [HELP](#)

```
interface Ethernet1/11
description
switchport mode private-vlan trunk promiscuous
switchport private-vlan trunk native vlan 888
switchport private-vlan trunk allowed vlan 1-3967,4048-4093
```

## Port-level access control lists (PACLs) AND VLAN-LEVEL access control lists (VACLs)

You no longer have to worry about PACLs that can cause issues but not be fully understood or visible. It's the same for the allowance of the Layer 2 VACLs that need to be in place and work together for seamless traffic flow. From below, we can now see if the PACLs and VACLs have verification and show if they are being used by the hit count. All of this previously would have been a combination of time-consuming commands and trial and error to find and verify if a PACL or VACL was shadowing a rule you may have recently set up.

**Access List Rules Diff**

[CHANGE ACLs COMPARED](#)

Search...

**acl-01-ipv6 from mgmt0 on Nexus-2**  
ACL changed May 13, 5:20 PM (Current)

|    |        |      |                   |                         |  |
|----|--------|------|-------------------|-------------------------|--|
| 10 | permit | ipv6 | 2001:db8:85a3::48 | 2001:db8:85a3::2112::64 |  |
| 20 | deny   | ipv6 | 2001:db8:85a3::48 | 2001:db8:85a3::2112::64 |  |

**ADMIN-VTY-ACCESS from port-channel1 on Same Node**  
ACL changed May 13, 5:20 PM (Current)

|     |        |                |               |             |         |
|-----|--------|----------------|---------------|-------------|---------|
| 10  | remark | Router         |               |             |         |
| 20  | permit | ip             | 172.88.0.0/23 | any         | log     |
| 30  | remark | SERVER Subnets |               |             |         |
| 40  | permit | ip             | 10.0.88.0     | 0.255.0.255 | any log |
| 70  | remark | MGMT VLANs     |               |             |         |
| 80  | permit | ip             | 10.0.33.0     | 0.255.0.255 | any log |
| 90  | remark | IT VPN Subnet  |               |             |         |
| 100 | permit | ip             | 10.77.66.0/24 | any         | log     |

## CONCLUSION

Monitoring has changed for the better in my eyes. We are engaging with devices and showing you the logical views that you are CLI creating and designing. We should be demanding more of these abilities from our monitoring software, as the time saved can increase your productivity. However, learning and fully understanding the technology is vital to any growing IT department.

I hope you have enjoyed this short book on the new mindset of Network Insight from SolarWinds. I appreciate your time and I hope you are able to increase your infrastructure awareness through your tools of today.

## TROUBLESHOOT LIKE AN EXPERT WITH SOLARWINDS NETWORK INSIGHTS™

[DOWNLOAD NOW](#)

Get deeper visibility into critical network devices from SolarWinds. With Network Insight features, SolarWinds is dedicated to continually bringing you better information about key network gear so you can manage your network, not your network monitoring. Download a free 30-day trial today with NPM and NCM.

## Additional Resources

### Network Insight for Cisco Nexus

Help ensure service availability with health and performance monitoring of critical data center switches.

### Network Insight for Cisco ASA

Automate the monitoring and management of your firewall infrastructure in a single unified platform.

### Network Insight for F5 BIG-IP

Monitor the health and performance of all components of application today

# Dedications

## TO THE THWACK MVPs

Let me start by stating that if you do not know of our THWACK® community, then you should immediately go to THWACK and sign up. There is no cost to you, but I assure you the knowledge you'll gain from being a part of this community is worth gold.

Now why I'm dedicating this to the MVPs of THWACK is simple. They are truly like family to myself and SolarWinds. Their dedication to beta testing and making sure that the information is solving their everyday issues is like nothing I have ever seen before. We all have a bond that stems from lifting each other up in our thoughts and understanding of technology and its challenges.

There is nothing for me to jump into a session with one of them and quickly figure out a new technology advancement and how we can monitor it effectively. Then on top of all the dedication to helping others within THWACK, they also are what I consider the best type of family. True, they're the heart of my inspiration and the bloodflow of my troubleshooting mind.

So, to all of the THWACK community and the MVPs, thank you from the bottom of my heart.  
~Dez~

## About The Author

Destiny Bertucci is a Head Geek™ at SolarWinds with a broad array of certifications and degrees, such as Cisco® Certified Network Associate (CCNA), (ISC)² Methodologies, CompTIA® IT Operations Specialist (CIOS™), CompTIA Secure Infrastructure Specialist (CSIS™), INFOSEC, database development degree, and SolarWinds Certified Professional®.



**Destiny Bertucci**  
SolarWinds Head Geek

In her 16 years as a network manager, she has worked in healthcare, federal, and application engineering, which allowed her to be a successful SolarWinds Senior Application Engineer for over nine years.

She started her networking career in 2001 by earning CCNA/Security+ certification and launching a networking consultant business. After using SolarWinds tools for many years, she joined the company and continued earning certifications and degrees to expand her professional reach into database development and (ISC)² methodologies. Customizing SolarWinds products while working on setups and performance deepened her knowledge of the complete SolarWinds product line. She is now skilled and experienced in network, security, application, server, virtualization, cloud, and database management.

# About SolarWinds

SolarWinds is Geek Built.™ That means that geeks, including SysAdmins, engineers, and other IT professionals, produce solutions for other geeks. SolarWinds addresses real problems that geeks face every day at work. We're not designing solutions based on which buzzwords are getting the most play on social media. Instead, we spend a lot of time talking to people in the trenches to find out not only what they are thinking about in terms of problems, but also how they would like to see those problems addressed. That feedback becomes the list of features we build into the next version.

Second, it's modular. You don't need to get the whole suite in one monolithic installation. You can determine which functionality you need, and then get the modules that meet those needs. The modules will snap together under a common framework, and also integrate well with solutions from other vendors. Because real geeks know that you don't get to pick every single piece of software the company uses, and that, like a good mutt, heterogeneous solutions are often the most robust and faithful allies you can have in the data center. The flipside of this is that each module is flexible. Each tool has a variety of "outside the box" actions you can take to get almost any job done.

Finally, and there's no way to dress this up, SolarWinds solutions are affordably-priced. Especially when you consider the features in each module, and function profile. Is it free? Of course not. But you are getting enterprise-class solutions at SMB prices.

Head over to [solarwinds.com](https://solarwinds.com) for more detailed information on products and pricing. You can also download a free, unlimited (meaning you can load up as many devices as you want), 30-day demo of any (or all) of the SolarWinds modules.

Pro Tip: There are also about two dozen free tools you can download over at: [solarwinds.com/free-tools/](https://solarwinds.com/free-tools/).