53-1003839-01
May 2015

Intelligent Flow Management

Solution Design Guide

BROCADE

Contents

Intelligent Flow Management Solution Design Guide

Preface

The processing of large data flows across a Wide Area Network (WAN) can often take hours if not days depending on a number of factors including: the size of the data set being transferred, the speed of the external connection, the protocol being used, and intermediate congestion points. The data sets transferred across the WAN are often trusted. Trusted, defined as, between two known and intended endpoints, with the transfer of non-harmful data.

The Intelligent Flow Management Solution (IFM) uses an open, modular approach to speed the transfer of large trusted data sets between sites.

Overview

This solution guide will present a Software Defined, dynamic, open solution to addressing the need for transferring large, trusted data sets between endpoints across the WAN. This guide provides use case areas for which the solution can be used, and describes the solution components and requirements.

The IFM solution allows users to pre-program known good data transfer endpoints. With these transfer endpoints programmed, the solution will automatically recognize a large data transfer that is occurring and push down a policy to the border router that will speed the transfer of the data by routing the traffic directly to its intended destination. This will allow the firewall to focus on smaller data flows, freeing up resources, while utilizing the border router's high speed backplane for data forwarding.

Purpose of this Document

This document is intended to assist by example in the deployment of the Intelligent Flow Management Solution. Readers will be able to understand the components that make up the solution, the components' respective roles in the solution, and be able to utilize configuration examples as a guide for their own deployments.

Audience

This document is intended for Network Architects, Network Administrators, Development Operations and anyone interested in learning about how Software Defined Networking can be utilized in production environments to solve real world problems.

Objectives

Today, organizations large and small collaborate with peers inside and outside the walls of the office or the datacenter. That collaboration can take on different forms, such as e-mail, video calls, and data

analysis. For some organizations, the collaboration includes sharing large data sets with internal and external colleagues for review. These organizations include, but are not limited to: Research Universities, Government Labs, Biotechnology, Life Sciences, Energy Exploration and Refining, Film and Television Media, Banking and Financial, Insurance, Gaming and Software Development, Healthcare, Legal and Discovery.

These organizations are no stranger to transferring data sets that can be more than 1 terabyte files. Assuming no interruptions, a 1 terabyte file takes over 15 minutes to transfer over a 10Gbps connection and over 2.5 hours over a 1Gbps connection. When sending these files externally, organizations have other contention issues that will increase this transfer time. These issues include additional data traffic utilizing the same connection, firewalls that are inspecting inbound traffic and potential bottlenecks across the WAN.

Additionally, many organizations are looking for use cases to introduce Software Defined Networking into their production environments. The IFM Solution provides an opportunity to incrementally add SDN functions without a requirement to rip out a massive amount of infrastructure. By utilizing OpenFlow and the Open Daylight Controller, the Solution provides a non-proprietary introduction and operationalization opportunity of Software Defined Networking. SDN is one of the foundations of the New IP.

What is the New IP? It is the old IP re-imagined for our modern world, and designed to meet the needs of cloud, mobile, social and big data. The new IP is both hardware and software...and it has both business and technology benefits.

The old IP is based on closed, proprietary systems, innovation cycles are constrained by custom hardware, and provisioning network resource is difficult and manual. Security is bolted on, interoperation is achieved through standards, vendors are at the center of the ecosystem, costs are high, and innovation is slow.

The New IP is based on open source where provisioning network resource is automated and self service. Security is built in from the start, interoperation is achieved through open APIs, the end user is at the center of the ecosystem, capex and opex costs are lower and innovation happens at the speed of business.

Related Documents

- Intelligent Flow Management Solution Brief
- Brocade SDN Controller
- Brocade Flow Optimizer Application
- Brocade MLXe Series Routers

About Brocade

Brocade®networking solutions empower the world's leading organizations to transition smoothly to a world where applications and information reside anywhere. By delivering agility and innovation for cloud-based environments, Brocade helps organizations modernize their networks and accelerate their journey to the New IP.

In particular, Brocade solutions for storage networking, data center routing, Software-Defined Networking (SDN), and Network Functions Virtualization (NFV) give organizations the power to capitalize on the unique business opportunities driven by virtualization and the cloud.

To deliver a best-in-class solution, Brocade partners with world-class IT companies around the globe. www.brocade.com.

Benefits of the Solution

The US Department of Energy, ESNet, says that when elephant data flows have to traverse a firewall they are 20 times slower than smaller data sets. Packet loss and latency have detrimental effects to TCP-based large data transfers. One option is to implement the Science DMZ architecture as defined by ESNet.However, many institutions struggle with a security policy that requires research data nodes to be secured behind enterprise border firewalls.

ESNet Firewall Performance Example

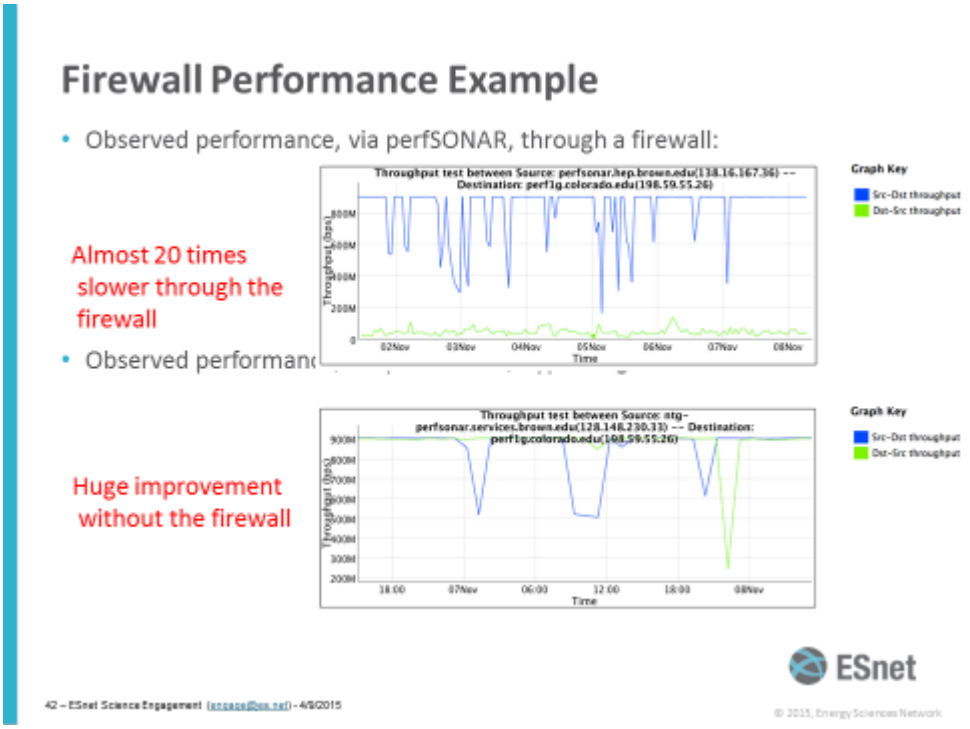

By utilizing the IFM solution, data nodes can sit behind firewalls, while trusted flows are identified in real-time and re-routed directly to its destination path in less than 10 seconds, vastly increasing the transfer speed. Now researchers can reduce their data transfer times from days to minutes.

Besides speeding large data transfers, the IFM solution brings additional value to Research Institutions. The Brocade Flow Optimizer application brings extra benefits for network administrators including: network traffic visibility, historical trending information, and identification/remediation of certain high volume traffic attacks. The open and software-driven solution is an excellent launching point for enabling the Software Defined Network. By enabling OpenFlow-capable infrastructure, an

[1] Eli Dart, Lauren Rotman, Brian Tierney, Mary Hester, and Jason Zurawski, “The Science DMZ: A Network Design Pattern for Data-Intensive Science”, SC13: The International Conference for High Performance Computing, Networking, Storage and Analysis. Denver CO, USA, ACM. DOI:10.1145/2503210.2503245, November 19, 2013, LBNL 6366E.

http://dl.acm.org/citation.cfm?id=2503245

[2] Slide from "Science DMZ and Architecture" - This webinar was presented on April 3rd 2015, in collaboration with the Great Plains Network, Network Nebraska, OneNet, KanREN, MOREnet and the South Dakota Board of Regents for the ENCITE (ENhancing CyberInfrastructure by Training and Education) project. Jason Zurawski from ESnet was the presenter for this webinar.

SDN controller based on OpenDaylight and utilizing the Brocade Flow Optimizer application in a high value production environment, the IFM solution takes SDN from theory to practice for Research Institutions. The IFM solution provides a solid infrastructure where additional SDN application capabilities and components can be built upon.

Design

Topology

FIGURE 1 Intelligent Flow Management High-Level Components

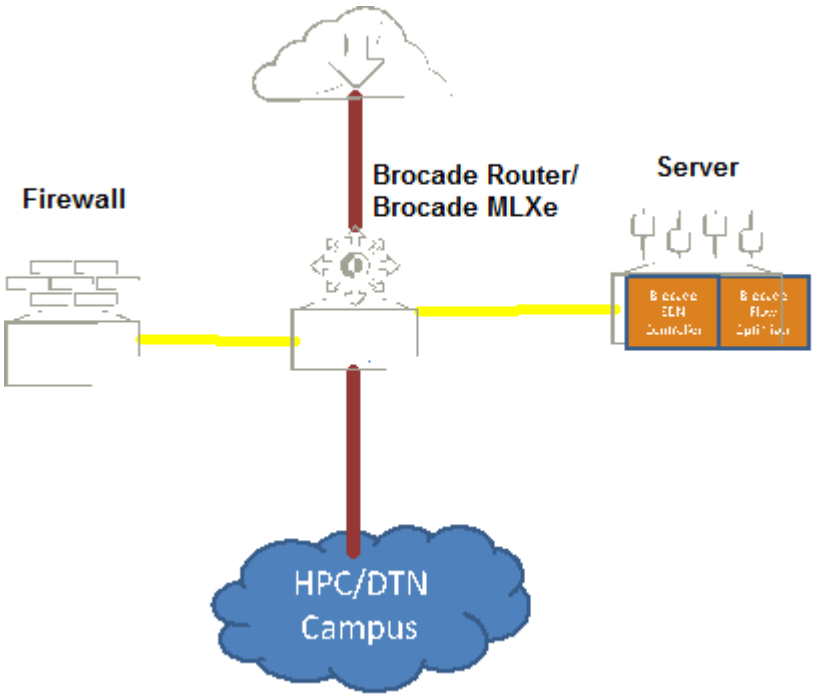

Border Router – Brocade MLXe

The Brocade MLX series core router delivers unprecedented scale and performance, high reliability, and operational efficiency for the most demanding service provider and enterprise networks. Built on a programmable architecture with high-density 100 Gigabit Ethernet (GbE), 40 GbE, and 10 GbE routing, these routers meet massive bandwidth demands, while maximizing ROI. Leading OpenFlow 1.3 scale in hybrid port mode provides a seamless transition to SDN for increased network agility and programmatic control.

The Brocade MLXe router is used as a border router by many of the worlds largest research universities. As more research institutions continue to upgrade their northbound connection to 100Gig high speed networks, like Internet 2, ESNet, GEANT, Regional Research and Education Network and commodity Internet Service Providers, they are implementing border routers that are open, scalable platforms.

Minimum system requirements:

- High Speed Connection to Upstream Provider (e.g. Internet Service Provider, Research and Education Network, Fiber Connection to Remote Data Center.
- OpenFlow v1.3
- OpenFlow Hybrid Port Mode Support
- sFlow

Firewall

There is no requirement for a specific vendor or list of supported features for the firewall. It is expected that most research institutions will have a 10Gbps connection or greater from the firewall to the border router.

Brocade SDN Controller

The Brocade SDN Controller is a quality-assured edition of the OpenDaylight controller - the industry's leading open source SDN controller. Brocade provides the tools and services to quickly implement software-defined networks and a completely open platform for application developers, with the backing of an established technology provider and its leaders within the OpenDaylight developer community.

Minimum System Requirements

- Operating System Options:
  - RedHat RHEL - 6.5
  - Ubuntu - 14.04
  - CentOS - 7
  - Fedora - 20
- Server System Specification Recommendations:
  - 3.0 GHz Intel® Xeon® or Intel® Core™ - 4 Cores or equivalent
  - RAM: 8GB
  - Storage: 64GB
  - Network: at least 1Gbps Ethernet

Brocade Flow Optimizer Application

The Brocade Flow Optimizer is a Software Defined Networking application that works in conjunction with the Open Daylight Controller. The application provides policy-based large flow detection and management with fine-grained control. With the assistance of an on-board sFlow collector within the application, it provides a number of dashboards and reports for visibility into traffic flow. Reports include real-time graphs, historical and trending graphs, and events list through a browser-based user interface. The Brocade Flow Optimizer also supports REST APIs that provide the capability for integration into a third party cloud orchestration system.

In addition to reporting and traffic visibility capabilities, the application can take pre-defined actions for traffic flows matching specific profiles. There are standard profiles as well as customer profiles that can be created within the application using any combination of Layer 2, Layer 3, and Layer 4 information.

Pre-defined actions to take on matched traffic flows include: re-direct, re-mark, meter or discard. The Brocade Flow Optimizer directs these actions through the OpenDaylight Controller via OpenFlow commands to network devices. Note that all these actions occur without any disruption to the Forwarding plane.

System requirements:

- To be used in conjunction with Brocade SDN Controller 1.3 or Open Daylight Controller Helium release
- Brocade Flow Optimizer and OpenDaylight Controller can be installed on the same server, separate, or on different VMs
- Server Needs to have In Band Connectivity to Border Router to Receive sFlow Packets
- Server specs:
  - Linux - Ubuntu 14.0.4 (64 bit), Centos 7
  - 8GB RAM
  - 64 GB of free HDD space

Traffic Flow Process

FIGURE 2 *Traffic Flow Process*

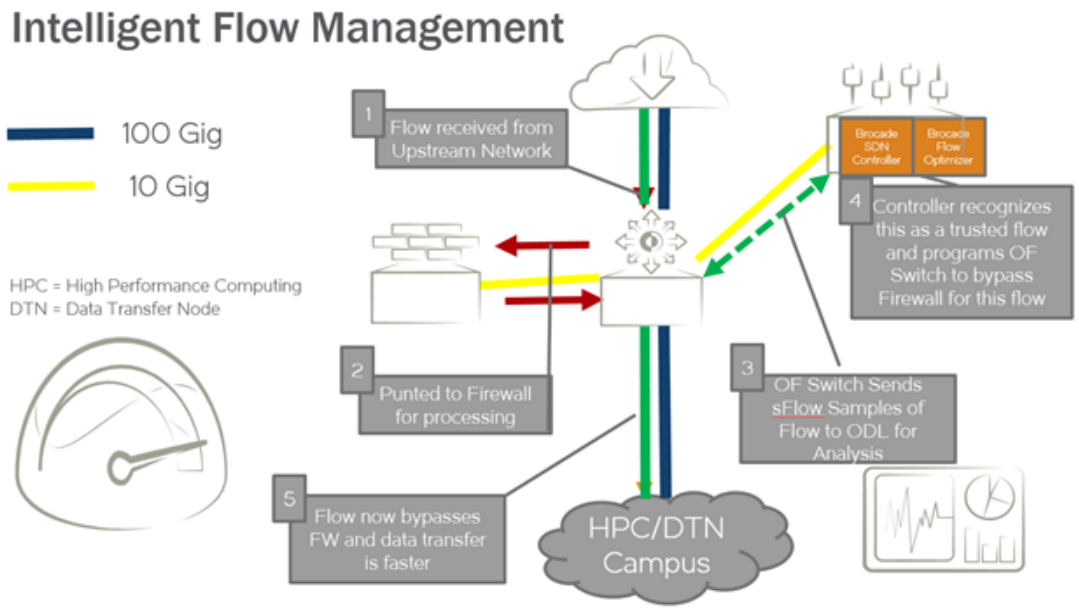

Process detail:

1. The targeted, trusted large data flow is received from the upstream network. The upstream network can be a connection to a Northbound Research Network or general ISP. These connections are typically at 100Gig speeds.
2. The traffic flow is received by the border router and sent to the Enterprise Firewall traffic inspection. After firewall inspection, the flow is sent back to the border router for routing to its intended destination. Typically the intended destination is the High Performance Computing cluster, Data Transfer Node or a Node on Campus.
3. As the border router receives the incoming packets from the upstream network, sFlow packets are sent to the Brocade Flow Optimizer Application. The Brocade Flow Optimizer application collects the sFlow packets and puts them into an on-board sFlow database. This database looks for pre-programmed matches of traffic profiles.
4. Once a traffic profile match of the trusted data flow has occurred. The Brocade Flow Optimizer application employs the Brocade SDN Controller to issue a preset OpenFlow command to the Border Router.
5. The preset action to take on this flow is for the Border Router to re-direct this trusted flow directly to its destination port, bypassing the firewall. Once the Openflow command is issued by the Brocade

SDN Controller to the border router, the flow is re-directed to its pre-defined port and data transfer is sped up.

The entire process from initial flow receipt to re-direction is accomplished in ten seconds or less. The same process can be applied to large data flows being transferred FROM the Campus/HPC or DTN TO the Upstream Network, speeding outbound transfers.

Component Configuration Examples

The following sections provide configuration examples of the components that are used in the Intelligent Flow Management Solution. These are meant to be used as a guide for the respective requirements that need to be installed and configured on each of the components. The example configurations should be modified to meet the respective environments for which they are to be installed.

Brocade MLXe Configuration Examples

Configuring OpenFlow on the MLXe router

1. Telnet or SSH into the router and get to the Configure Terminal mode
```
NetIron MLX-4 Router>enable
NetIron MLX-4 Router#configure terminal
NetIron MLX-4 Router(config)#
```
2. Enable OpenFlow Version 1.3 and configure the OpenFlow controller IP address

NOTE
The controller IP address used in this example is 10.1.2.11.

```
NetIron MLX-4 Router(config)#openflow enable  ofv130
NetIron MLX-4 Router(config)#openflow controller ip-address
10.1.2.11 no-ssl port 6633
```
3. Enable OpenFlow hybrid port mode on the desired interfaces
```
NetIron MLX-4 Router(config)#interface ethernet 1/1
NetIron MLX-4 Router(config-if-e10000-1/1)#openflow enable
layer23 hybrid-mode
```

NOTE
You can specify Layer 2, Layer 3, or both layers in hybrid mode as Layer23 matching mode to supported on the interface. By default, interfaces on these devices support Layer 2 matching mode. If you enable Layer 2 matching mode on the specified interface, only Layer 2 matching fields are supported on that interface.

4. Set the system maximum. (System reload is required once you change the system maximum values.) The system maximum values for OpenFlow entries:
```
NetIron MLX-4 Router(config)#system-max  openflow-flow-
entries <Valid Decimal Entry>
```

NOTE
DECIMAL: Valid range is 0 to 65536. Default: 0.

5. OpenFlow protected VLAN entires
```
NetIron MLX-4 Router(config)#system-max  openflow-pvlan-
entries <Valid Decimal Entry>
```

NOTE
DECIMAL: Valid range is 0 to 2048. Default: 0.

6. OpenFlow unprotected VLAN entries
```
NetIron MLX-4 Router(config)#system-max  openflow-
unprotectedvlan-entries  <Valid Decimal Entry
```

NOTE
DECIMAL: Valid range is 0 to 4096. Default: 0.

7. MAX NP OpenFlow Entires
```
NetIron MLX-4 Router(config)#system-max np-openflow-
entries layer2or3 | layer23IPv4 value slot [ i j k |
i to z | all].
```

NOTE
The slot number value can be any of the valid slot numbers in the device. For slots, you can provide "all", "slot 1 to 2" or individual slot options. The following parameters are available for this command:

- layer2or3 - layer 2 flow or layer 3 flow entries
- layer23IPv4- Layer 2 and 3, including L2 and IPv4 flow entries

8. Enable sFlow services and configure sFlow destination (collector) IP address

NOTE
For example, 10.1.1.10 is an In-Band IP address (NIC) for BFO where sFlow packets are received on default port 6343

```
NetIron MLX-4 Router(config)#sflow enable
NetIron MLX-4 Router(config)#sflow destination
<Inband IP Address of the VM where BFO is installed> 6343
```

9. Enable sFlow null0-sampling on the global level
```
NetIron MLX-4 Router(config)#sflow null0-sampling
```
10. Set the sFlow sampling rate at the recommended rate: 8192
```
NetIron MLX-4 Router(config)#sflow sample  8192
```
11. Enable sFlow forwarding on all the Ingress ports being monitored for Ingress Large Data Traffic Flows:
```
NetIron MLX-4 Router(config)#interface ethernet 1/1
NetIron MLX-4 Router(config-if-e10000-1/1)#sflow
forwarding
```
12. The "show running-configuration" should display the sFlow configuration on the router
```
NetIron MLX-4 Router(config)#show running-
configuration
```
13. Perform a write memory operation to save the running-configuration to the startup-configuration and retain the changes
```
NetIron MLX-4 Router(config)#write memory
   Write startup-config done.
```

Brocade SDN Controller Configuration Examples

In addition to the system requirements outlined earlier, the server running the Brocade Vyatta Controller needs to have the following applications installed: JAVA, Zip, Unzip, cURL, Node.js, and OpenSSH. More information about installing these applications can be found here.

Installing the Controller and GUI

1. Download the following installation directories of the controller.

a. bvc-<version>.zip
b. bvc-dependencies-<version>.zip

2. Enter the following command to create the /opt/bvc directory:
   `sudo mkdir /opt/bvc`
3. Enter the following command to change the ownership of the directory:
   `sudo chown $USER /opt/bvc`
4. Enter the following commands to unzip the installation directories:
   ```
   unzip -o bvc-<version>.zip -d /opt
   6.    unzip -o bvc-dependencies-<version>.zip -d /opt
   ```

   These commands create files in the /opt/bvc/ directory.
5. Enter the following commands to install the controller:
   ```
   cd /opt/bvc
   ./install
   ```

   The controller and the GUI automatically start at the end of the installation.

   The installer displays the URL that starts the controller GUI. For example: `Server@http://<controller_ip>:9000`

Enabling HTTPS on the Controller

HTTPS provides authentication of the web site and the associated web server and also provides a bidirectional encryption of communication between the client and the server. This process provides a reasonable guarantee that the contents of the communication between you, the controller, and the controller GUI cannot be read or forged by any third party.

You have already installed the controller and the controller GUI and verified that both are running successfully.

Follow the steps in the next section to enable HTTPS and to disable HTTP communication on the controller and the controller GUI.

1. Stop the controller by entering the following command:
   `/opt/bvc/bin/stop`
2. Run the setup_https script to enable HTTPS by entering the following command:
   `/opt/bvc/bin/setup_https on`
3. Start the controller by entering the following command:
   `/opt/bvc/bin/start`

   The controller responds to HTTPS on the port 8443 instead of the port 8181.

   The controller GUI responds to HTTPS on the port 9443 instead of the port 9000.

NOTE
When you run the setup_https script, the script generates a self-signed certificate and installs the certificate on the web and the controller processes,

Starting and Shutting Down the Controller and GUI

The controller is automatically started after it is installed. These instructions, to start and shut down the controller, are useful only for troubleshooting.

1. Enter the following command to start the controller and the GUI:
   `/opt/bvc/bin/start`
2. Enter the following command to stop the controller and the GUI:
   `/opt/bvc/bin/stop`

Brocade Flow Optimizer Configuration Examples

Installing the Application, this can be done on a separate server from the Brocade SDN Controller, on a separate VM or on the same server under a different service port:

1. Install all of the required software listed under System Requirement
2. Download the Brocade Flow Optimizer software (BFO1.0_bld <X> distribution.tar)
3. Create a directory for the Brocade Flow Optimizer software distributable archive
4. Copy the software (*BFO1.0_bld <X> distribution.tar*) to the directory created
5. Extract the tar archive to the directory

In the configuration properties file (*config.properties*), set the values of the following parameters based on your Brocade SDN Controller installation:

1. Go to the home directory of the Brocade Flow Optimizer (where it was installed)
2. Open the configuration folder, and input the parameters for the following entries:
   - odlcontroller.ip
   - odlcontroller.username
   - odlcontroller.password
3. Save the changes.

Discover MLX Devices Using the Brocade Flow Optimizer

1. Go to the home directory of the Brocade Flow Optimizer (where it was installed)
2. Open the configuration folder
3. In the configuration properties file (*config.properties*), set the values of the following parameters based on your Brocade Vyatta Controller installation: `node.list`

   (For example: node.list= 00:24:38:ae:cb:00,00:24:38:88:54:00)
4. In this example, there are two nodes. The MAC address of each device is included in the value.

NOTE
If you do not know the MAC address of the device (or devices), use the following steps to get the MAC address information.

---

   a. Telnet to the device
   b. Execute the s`h int management 1` command
   c. Search for the string below in the output and add the MAC address as `"00:24:38:ae:cb:00 Hardware is Ethernet, address is 0024.38ae.cb00 (bia 0024.38ae.cb00)"`
5. Save the changes.

Starting the Brocade Flow Optimizer Application

1. Go to the home directory of the Brocade Flow Optimizer (where it was installed)
2. Open the configuration folder
3. In the configuration properties file (*config.properties*), specify the valid MAC address of the node and the IP address of the Brocade SDN Controller
4. Change the directory to the bin directory
5. Run the *service.sh* script to start the application. This script automatically initiates and starts the database server
6. Open the Brocade Flow Optimizer client in your Google Chrome browser using one of the following URLs:
   - https://localhost:8089
   - https://<IP-address of Brocade Flow Optimizer server>:8089
7. Login to the Brocade Flow Optimizer:

8. The Brocade Flow Optimizer dashboard appears. You will notice no flows appear in the dashboard:

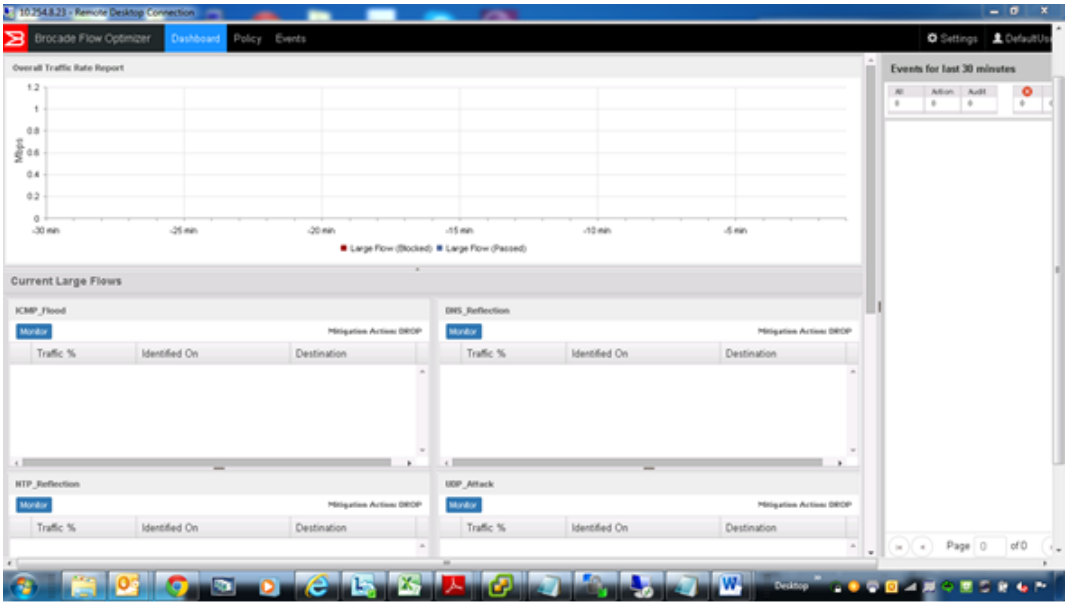

9. By clicking the "Policy" tab at the top of the Brocade Flow Optimizer dashboard, you are then able to "Add Custom Profile":

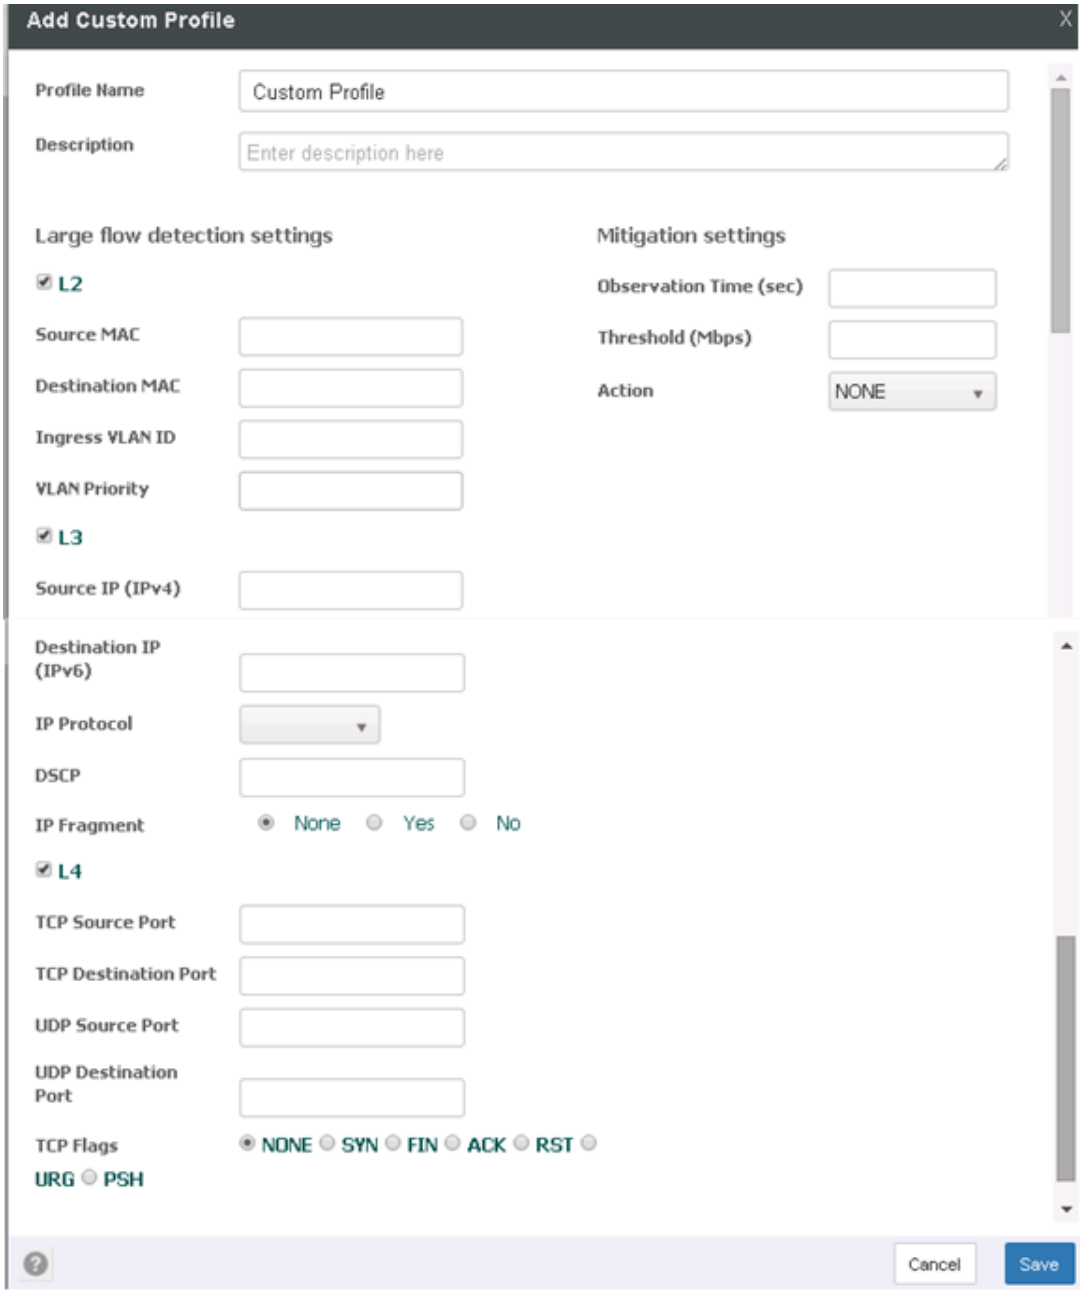

Your flow detection options include: Source/Destination MAC address, Source/Destination IP address (v4 or v6), Source/Destination TCP/UDP port number. You also have options for VLAN ID, VLAN Priority, and DSCP matching

Flow Mitigation Options Include: Drop, redirect, and Meter

Configuration Example:

Suppose we want to redirect a flow going from 10.10.10.10/24 to 20.20.20.20/24. We want to redirect this traffic on router 10.254.6.31 to port 50. Here is how this would be configured in the Brocade Flow Optimizer:

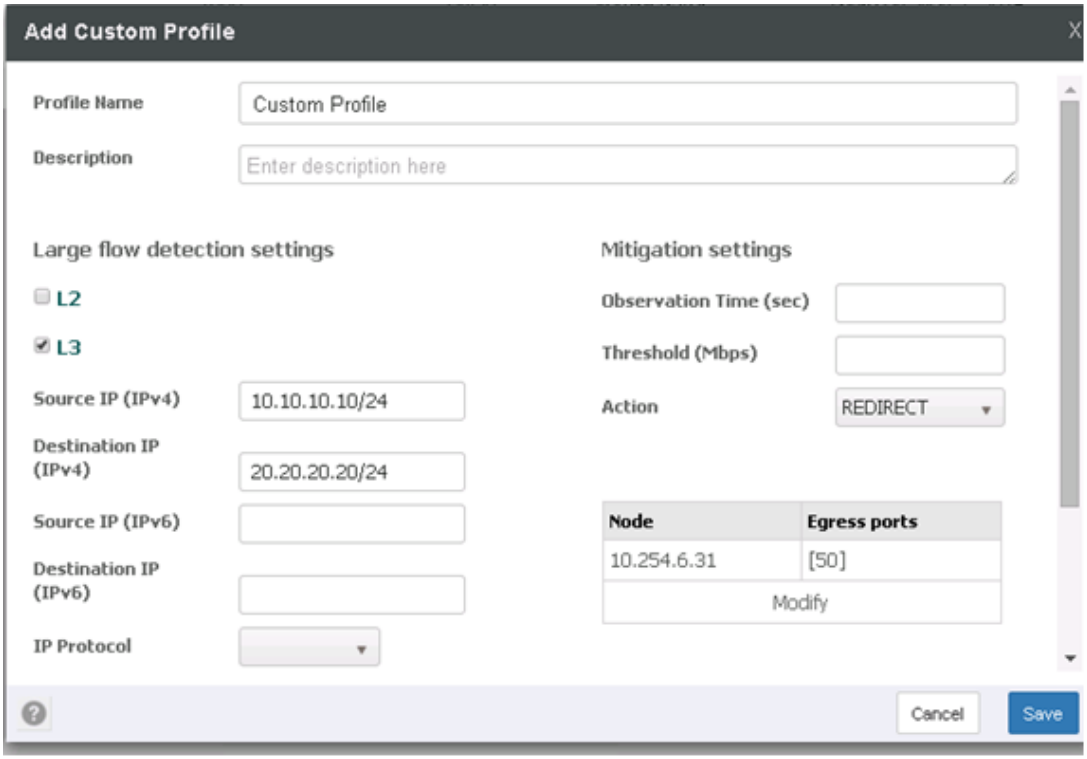

Stopping the Brocade Flow Optimizer Application

1. Go to the home directory of the Brocade Flow Optimizer (where it was installed)
2. Change the directory to the bin directory)
3. Run the *stopService.sh* script to stop the application. This script automatically stops the application and the database server