



NetEqualizer User Guide



Latest NetEqualizer software version is now 10.5
this userguide is still valid. Addendums are available.

© Copyright 2003-2026 APconnections, Inc. All Rights Reserved.

No part of this publication including text, examples, diagrams or illustrations may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical or otherwise, for any purpose, without prior written permission of APconnections, Inc.



Table of Contents

Where to Install NetEqualizer	6
Setting up the NetEqualizer.....	7
Factory Default IP Settings	7
Cabling the NetEqualizer into your Network	7
Configuring the NetEqualizer.....	9
The NetEqualizer Dashboard.....	10
Navigation - How to Move Around the System.....	11
Setting Preferences	12
Configure Equalizing	13
Equalizing Defined	13
Configuring Equalizing Parameters	14
Set Key Equalizing Parameters	15
Using the RATIO Parameter to Influence Equalizing.....	15
Setting your Trunk Size	16
Key Parameter to Adjust Equalizing Sensitivity.....	17
Set Additional Equalizing Parameters.....	18
Additional Parameters to Adjust Equalizing Sensitivity	18
Parameters to Size Internal Tables	20
Viewing your Equalizing Parameter Settings.....	21
Equalizing in Action	21
Setting Bandwidth Limits.....	22
Configure Hard Limits by IP	23
Hard Limit Rules.....	24
Creating a Hard Limit.....	25
Modifying a Hard Limit	25
Deleting a Hard Limit	26
Adding Bursting to Hard Limits.....	27
Viewing your Hard Limits.....	29
Setting up Bandwidth Pools.....	30
Bandwidth Pool Rules.....	31
Creating a Pool	32
Adding Pool Members	33
Modifying a Pool or Pool Member	34
Deleting a Pool or Pool Member.....	34
Viewing Pools and Pool Members.....	35
Setting VLAN Limits	37
VLAN Limits Rules	38
Creating a VLAN Limit	38
Modifying a VLAN Limit.....	39
Deleting a VLAN Limit	39
Viewing VLAN Limits	40
View All Traffic Limits	41



Limit P2P Traffic	42
Connection Limits Defined	42
Viewing Connection Counts	43
Setting Connection Limits	43
<i>Connection Limit Rules</i>	<i>44</i>
<i>Creating a Connection Limit</i>	<i>45</i>
<i>Modifying a Connection Limit</i>	<i>46</i>
<i>Deleting a Connection Limit</i>	<i>47</i>
Viewing your Connection Limits	48
Consider Setting Bandwidth Priority	49
Defining Priority Traffic	49
<i>Creating a Priority Limit</i>	<i>50</i>
<i>Modifying a Priority Limit</i>	<i>51</i>
<i>Deleting a Priority Limit</i>	<i>51</i>
Masking Off Traffic	52
<i>Creating a Masked Host</i>	<i>53</i>
<i>Modifying a Masked Host</i>	<i>53</i>
<i>Deleting a Masked Host</i>	<i>54</i>
Viewing your Priority Limits and Masked Traffic	55
Restricting Bandwidth Usage	56
Establishing User Quotas	56
<i>User Quota Rules</i>	<i>57</i>
<i>Creating User Quota Rules</i>	<i>58</i>
<i>Resetting User Quota Rules</i>	<i>59</i>
<i>Modifying User Quota Rules</i>	<i>61</i>
<i>Deleting User Quota Rules</i>	<i>61</i>
<i>Starting the Quota System</i>	<i>62</i>
<i>Viewing User Quotas</i>	<i>63</i>
MAC Redirection	65
Perform Quick Edits	67
<i>Quick Edit - Deleting a Rule</i>	<i>68</i>
<i>Quick Edit - Adding a Rule</i>	<i>68</i>
Distributed Denial of Service Attack (DDoS) Tools	69
DDoS Monitor	69
DDoS Firewall	70
Monitoring and Reporting	71
Dynamic Real-Time Reporting (RTR)	73
RTR Dashboard	73
<i>Real-time General Traffic</i>	<i>74</i>
<i>Real-time Pool Data</i>	<i>74</i>
RTR Active Connections Reports	75
<i>View Active Connections</i>	<i>75</i>
<i>IP Lookups</i>	<i>76</i>
<i>View Active IPv6 Connections</i>	<i>79</i>
<i>View Connection Counts</i>	<i>80</i>
<i>View Active Penalties</i>	<i>81</i>
RTR Traffic History Reports	82
<i>Start RTR</i>	<i>83</i>
<i>Manage Tracked IPs</i>	<i>83</i>
<i>General Traffic History</i>	<i>84</i>
<i>Traffic History by IP/Pool/VLAN</i>	<i>84</i>
<i>Top Talkers</i>	<i>86</i>
<i>General Penalty Reports</i>	<i>86</i>
<i>Clear Reporting Data</i>	<i>87</i>



Export Data to a Reporting Data Warehouse	88
View NetEqualizer Log.....	88
Configuration.....	90
Running Processes	91
Start/Stop RTR	92
Autostart RTR	92
Email Notifications	93
Configure Email.....	93
Configure Alerts	94
Redundancy and Failover.....	95
Setting up Full Redundancy.....	95
Failover.....	96
Maintenance Tasks.....	97
Powering Off the NetEqualizer	97
Backing Up Your Configuration Settings	97
Getting Software Updates for the NetEqualizer	97
Troubleshooting.....	100
Frequently Asked Questions (FAQs)	103
Appendix 1 - Equalizing Parameters, Units, & Defaults.....	107
Appendix 2 - Setting/Forcing LAN Speeds and Duplex	108
Appendix 3 - Packet Capturing for taps such as CALEA	110
Appendix 4 - Tuning Parameters for a Large Number of subnet-ranged Limits, Pools, & Masks... ..	112
Appendix 5 - Syncing NetEqualizer Date/Time	113
Appendix 6 - Firewalling the NetEqualizer	115



Thank you for purchasing a NetEqualizer. You are now on your way to achieving "Faster Networks, with Zero Maintenance, at the Best Prices". Using NetEqualizer in default factory mode will take care of almost all network congestion and priority traffic flow requirements, and is the recommended operational mode for most customers. However, NetEqualizer also offers a wide range of bandwidth control options, while at the same time allowing you to keep it simple.

NetEqualizer Quick Start Guide

To perform your initial installation, you should reference the NetEqualizer Quick Start Guide. This contains the complete system settings configuration and minimal settings required to get you up and running. A hard copy is included in your shipping box. We also email a PDF copy with your shipping confirmation email. You can also find our [demo version](#) (without passwords) online.

Note: The NetEqualizer Quick Start Guide is a step-by-step instruction manual.

NetEqualizer User Guide

The NetEqualizer User Guide is intended to walk through NetEqualizer features in more detail than our NetEqualizer Quick Start Guide. It contains detailed descriptions of equalizing, limiting, reporting, DDoS, and our add-on modules. Once up and running, it is a good idea to review this entire NetEqualizer User Guide, to become familiar with all of the advanced features available to you.

Note: The NetEqualizer User Guide is not a step-by-step instruction manual. Select the feature you are interested in from the Table of Contents and go directly to that section.

For Additional Help

Should you need further assistance setting up your NetEqualizer, please call our Support Team at 303.997.1300 x102 or email support@apconnections.net. If you purchased through an authorized distributor or reseller, check with them first to determine if they support you directly.

Key to Reading the User Guide


Entity	Format	Example
GUI Parameter Name	Shown in bolded blue. Sometimes followed by SYSTEM PARAMETER NAME.	Bandwidth Up (TRUNK_UP)
SYSTEM PARAMETER NAME	System parameter names may follow the GUI Parameter Name. They are in CAPITAL LETTERS.	TRUNK_UP
Notes	Notes are shown in blue italics preceded by "Note:".	<i>Note: For a detailed list of the steps necessary to get up and running, please see the NetEqualizer Quick Start Guide.</i>
Type in: values to be entered	Values that you need to type in are shown in orange italics preceded by "Type in:".	Type in: <i>/bridge/bridge-utils/brctl/brctl rembrain my 99999</i>
Click on ->Menu Name or Tab Name	Menus are shown in orange italics preceded by "Click on->" or "->".	Click on -> <i>Manage Equalizing -> Configure Parameters.</i>
Click on -> [button_name]	Buttons are shown in orange italics surrounded by [square brackets] preceded by "Click on ->" or "->".	Click on -> <i>[Equalizing Status Indicator] -> [Start Equalizing].</i>



Where to Install NetEqualizer

NetEqualizer can be installed on any link whose traffic you would like to shape. For maximum effectiveness, most users should install NetEqualizer between the network users and the Internet trunk. Traffic running between your network and the Internet is generally a constriction point in traffic flow where many users compete for this limited resource. By placing your NetEqualizer at this junction you will automatically optimize your Internet speed.

The NetEqualizer operates as a Transparent Bridge on your network. There is typically no need to change anything in your network configuration to install the appliance. Simply install the NetEqualizer between your Router and Network Switch, or anywhere you can see the individual IP addresses you wish to shape. Set-up using the [Quick Start Guide](#) to modify any factory default settings, and then access it via a Web Graphical User Interface.





Setting up the NetEqualizer

For a detailed list of the steps necessary to get up and running, please see the [Quick Start Guide](#). If you do not have a copy of the Quick Start Guide, please request one by calling our Support Team at 303.997.1300 x102 or emailing support@apconnections.net.

We mention some of the key functions performed during set-up here. However, you will need to reference the Quick Start Guide to complete your set-up.

Factory Default IP Settings

The IP settings to access the NetEqualizer web Graphical User Interface (web GUI) have been set to:

IP Setting	Parameter Name	Factory Default
Web GUI IP	BRIDGEIP	192.168.1.143
Web GUI Netmask	BRIDGENETMASK	255.255.255.0
Web GUI Gateway	BRIDGEROUTE	192.168.1.1

Note: The IP address for the NetEqualizer is used to access the web GUI (for management purposes only). All factory default settings can be changed from the web GUI, the command line, or the API.

Cabling the NetEqualizer into your Network

First, make sure that you **power on the NetEqualizer**. Do this by pressing the red power button to the right of the LED panel.

Note: We recommend that you install your NetEqualizer on a UPS, to protect from power surges and outages.




All of the NetEqualizer models (series 3000/4000/5000) have two Ethernet interfaces and a Management Port. We are now using port plugs to help distinguish the various interface ports on the NetEqualizer. We use four colors: 1) blue (WAN), 2) orange (LAN), 3) clear (Management Port) and 4) black (unused).

The easiest way to figure out the ports is to look at the port plug color. The LAN port is orange, the WAN port is blue, and the Management Port is clear. If you have upgraded to fiber interfaces, the LAN and WAN ports will be white. For further clarification, please see the "Where to Install the NetEqualizer" diagram below.

Management Port (CLEAR)

First, find the Management Port (Eth2). As you face the back of the machine, Eth2 is on your top row LEFT (NE3000) or bottom row LEFT (NE4000/NE5000). Remove the clear port plug and then plug a *straight cable* into the Management Port (Eth2) and connect it to your Network Switch. If whatever you are plugging into does not have a built-in switch, then use a crossover cable. Once you have the Management Port connected, you can now configure your NetEqualizer.



WAN Port (BLUE or WHITE)

Remove the blue port plug (white if fiber) and then plug a *crossover cable* into the WAN Port (Eth0) and connect it to the Firewall/Router. As you face the back of the machine, Eth0 is on your bottom row LEFT (NE3000), or top row RIGHT (NE4000/ NE5000). If you have an auto-sensing Firewall or Router, you can use a straight cable or a crossover cable.

LAN Port (ORANGE or WHITE)

Remove the orange port plug (white if fiber) and then plug a *straight cable* into the LAN Port (Eth1) and connect it to your Network Switch. As you face the back of the machine, Eth1 is on your bottom row RIGHT (NE3000) or top row LEFT (NE4000/5000). If whatever you are plugging into does not have a built-in switch, then use a crossover cable.

Once your machine is on & connected, you should see green lights in the Power LED, Eth0, and Eth1 LEDs, as shown in the display panel picture on the previous page.

(optional): Access Point Configuration in a Wireless Network

Put your radios in bridging mode and set your Firewall/Router at your headend to do DHCP and NAT, instead of doing DHCP and NAT at your Access Points.

(optional): Setting LAN Port Speed and Duplex

Occasionally, customers need to manually set LAN Port Speed and Duplex as some Firewall/Routers do not auto-negotiate correctly with the NetEqualizer. If this is happening in your environment, you will see a large number of collisions and dropped packets as well as reduced network throughput. Although dropped packets are not a good thing, if you are seeing less than 1/10 of a percent (< 0.1%) of the total packets transmitted it will have no adverse effect on your network. If it starts to approach 1 percent (1%), you should follow [Appendix #2](#) to set this in your environment.

(optional): Firewalling off the NetEqualizer

If you do not install the NetEqualizer behind a firewall, you should use [Maintenance-> Manage Firewall](#) to firewall off the NetEqualizer. See [Appendix #6](#) for detailed instructions.



Configuring the NetEqualizer

Once the NetEqualizer is powered on and plugged in to your network, you need to access the web GUI to configure it. The web GUI is accessible out-of-the-box via the factory default IP address: <http://192.168.1.143>. If you have set up the Management Port as described above in "Cabling the NetEqualizer into your Network", you should now be able to access the web GUI. The NetEqualizer Dashboard is displayed (shown below). The NetEqualizer Dashboard is described in detail in the next section of the User Guide.




Once you login to the NetEqualizer, your next steps to configure the NetEqualizer will be to change the default passwords, install your license key, change your IP addresses, and set the date/time and time zone. All of these functions are found under the **[Setup]** button, circled in orange above on the NetEqualizer Dashboard.

When you *Click on -> [Setup]*, the Setup window at right come up. Note that all Setup functions are available by clicking on the links (orange text).

The Side Menu also can be used to navigate to functions. As Manage NetEqualizer (circled in orange) is followed by a "+", it is an expandable menu. Clicking on Manage NetEqualizer + expands to

contain the same functions as listed in orange under Manage NetEqualizer on the Setup window. All Manage NetEqualizer functions are described in detail in the [Quick Start Guide](#).



Note that clicking on a Side Menu Item or a Link (orange text) will bring up a new configuration window. This paradigm is repeated throughout the GUI.

Please see the Quick Start Guide, starting with Step #2: "Configuring the NetEqualizer" to complete NetEqualizer set-up & configuration.



The NetEqualizer Dashboard

At the heart of the NetEqualizer system is the NetEqualizer Dashboard. The NetEqualizer Dashboard provides an intuitive visual display of the status on critical data and settings within NetEqualizer. Think of the Dashboard as your command and control center for managing your NetEqualizer. On the picture below, the key elements that make up the Dashboard are labeled: Key Functions, Information Buttons, Common Tasks, Status Indicators, and the Current Activity Graph. The NetEqualizer Dashboard described below is the new redesign, available as of [Software Update 8.4](#).



Key Functions buttons are your main access point into configuring, monitoring, reporting, and maintaining your NetEqualizer. They include: 1) Setup, 2) Real-Time Reporting (RTR), 3) DDoS, and 4) Maintenance.

Common Tasks are shortcuts to areas within the system that you use frequently. From here you can Start/Stop Equalizing, Show your current NetEqualizer Configuration, View Active Connections (all traffic running through the NetEqualizer), or Run Diagnostics.

Information Buttons provide a quick overview of key settings. To help keep you up-to-date on if you need to upgrade, we display the current software version that you are running. You can also see the system date & time and time zone, as well as your license key setting and any key violations. Click on any of the green buttons to go the Install License Key screens, or on the blue Date & Time button to set your date, time or time zone.

The Dashboard contains **Status Indicators**, visually displaying on/off status on key functions: Equalizing, Real-Time Reporting (RTR), Quotas, and Packet Capture. Clicking on the Equalizing, RTR, or Quota indicators opens up the related start/stop screen. Clicking on Packet Capture brings up the associated set-up screen.

The **Current Activity Graph** is a view into Real-Time General Traffic, showing the total amount of upload and download data flowing through your NetEqualizer. Click on the graph to open up the graph in the full RTR Dashboard.


In this User Guide, we will discuss the features available via Key Functions in detail.



Navigation – How to Move Around the System

Each of the Key Function screens work in the same way. Here we will review one screen, so that you know how to navigate the NetEqualizer system.

We will use the Setup Function as our example. From the NetEqualizer Dashboard, *Click on -> [Setup]*. The following screen opens.



The Setup **Side Menu** (circled above), can be used to navigate to each Setup Function. If a Setup Function is followed by a "+", it is an expandable menu. On the Setup Side Menu, both Manage NetEqualizer+ and Manage Traffic Limits+ are expandable. Clicking on either one opens up a sub-menu that contains the same functions as listed in orange on the Setup window. For example, if you click on "Manage Traffic Limits +", you will see Configure Hard Limits, Configure Pools and VLANs, etc.

On the Setup screen, all Setup **Functions** are available by clicking on the links (orange text). For example, the Manage Traffic Limits functions are circled in blue above.

Note that clicking on a Side Menu Item or a Function (orange text) will bring up a new configuration window. This paradigm is repeated throughout the GUI.

The **Navigation Menu**, circled in orange above, contains shortcuts available from all windows. These help you quickly and easily get back to the NetEqualizer Dashboard (Home) or Key Functions (Setup, RTR, DDoS, and Maintenance) from anywhere in the system.

The **Minimize Button** can be used to shrink the Side Menu, so that the main screen is maximized. You can also use it to expand the Side Menu.


This paradigm is used throughout the screens in the NetEqualizer 8.4 GUI.



Setting Preferences

Before you start configuring the NetEqualizer, we recommend that you select the settings to be used for display in the GUI.

As of Software Update 8.4, we give you the opportunity to select the units you would like to see in the NetEqualizer GUI for both configuring and reporting on your traffic. We now offer a variety of settings to meet your needs. To get started, [Click on ->\[Setup\] -> Preferences](#). The Preferences window appears, as shown below.



Configuration Units

These units are used throughout the Setup function, specifically for configuring Equalizing and all Traffic Limits. We strongly recommend that you set this to megabits per second (Mbps), which best matches most customer's environments.

General Traffic/Real-Time Graph Units

These units are used anywhere the General Traffic Graph is displayed. This includes the NetEqualizer Dashboard, the RTR Dashboard, and in General Traffic History. We also recommend that you set this to megabits per second (Mbps).

Traffic by IP/Pool/VLAN Graph Units

These units are used anywhere Traffic is graphed by IP, Pool, or VLAN. This includes the RTR Dashboard, and in RTR Traffic History. We recommend that you set this to units that give you visibility to the granularity that is most useful in your environment.

Once you have selected your settings, [Click on -> \[Save Changes\]](#) to keep your settings, or click back to any other function or menu to ignore your changes.



Configure Equalizing

At the core of our traffic shaping capabilities is Equalizing. In this section we first define Equalizing, and then walk you through the concepts behind the Equalizing Parameters, and how you can tune them for your environment.

Equalizing Defined

Equalizing is a simple concept. It is the art form of looking at the usage patterns on the network, also known as traffic "behaviors". When things get congested on your network, equalizing applies fairness algorithms, based on the observed traffic behavior, to ensure that everyone gets a share of the available bandwidth.

This **behavior-based approach** usually mirrors what you would end up doing if you could see and identify all of the traffic on your network, but does not require the labor and cost of classifying everything. During congested periods on your network, applications such as VoIP, web browsing, web-based applications (SaaS, cloud applications, etc.), short downloads, and instant messaging (IM) all naturally receive *higher priority*, while large downloads, large videos, and live video streams receive *lower priority*. As priority is applied only during congested periods, traffic can flow freely when the network is not congested.

Once equalizing is in place, it automatically shapes your network when it is congested, using algorithms to implement "fairness". The concept of "fairness" enables your network to continue providing quick response times to the majority of your users while restricting the network hogs. Low bandwidth users do not have to share the pain of a slow, congested network with the network-hogging applications.

Equalizing does this by using our proprietary algorithms to implement fairness. First, equalizing tracks how much bandwidth is being used. If bandwidth used is over a predefined level, the network is considered congested. Once the network is considered congested, equalizing looks at every connection (IP address pair) and puts a PENALTY on those that are over a predefined level. Penalties are applied and removed as needed while traffic is congested. This process continues until network congestion eases.

Equalizing bases its decision to issue penalties based on built-in fairness rules:

- The persistence of the user's connections. We look at the length of time the connections have been live. The longer the time, the more likely a penalty.
- The amount of bandwidth used relative to the total size of the trunk.
- The number of users on the trunk. The more users active on the trunk, the less bandwidth NetEqualizer will allow per user before issuing a penalty.
- Is the overall trunk saturated? A trunk is saturated when it reaches the percentage defined by the RATIO parameter (default RATIO = 85%).

Equalizing is tunable, so that you can set the level of traffic that you wish to be given higher priority, based on your network environment. By setting equalizing parameters, we offer you the flexibility to control your level of equalizing as your network environment changes.

Rather than writing hundreds of rules to specify allocations to specific traffic as in traditional application shaping, equalizing gets out of the game of prioritizing applications altogether. We believe this to be a superior method to shape network traffic, particularly when the traffic types and amounts change frequently. You will find under our Equalizing approach



that behavior-based shaping does not need frequent maintenance & changes address new applications being introduced to your network.

Configuring Equalizing Parameters

Each Equalizing Parameter is discussed in detail below. For a summary of all Equalizing parameters, please see [Appendix 1](#), which contains a one-page cheat sheet with the default settings and recommendations.

It is important to ensure that Equalizing is "on" before you adjust Equalizing Parameters. NetEqualizer comes pre-configured to automatically start up with Equalizing turned on (Equalizing Rules = On and Equalizing Process started).

Equalizing Process Started

Click on -> [Setup] -> Configure Equalizing. The Configure Equalizing window appears, defaulted to the Key Equalizing Parameters tab, as shown below. As of Software Update 8.4, once you are in the Configure Equalizing screen, if Equalizing is OFF, you will see a warning message, seen here, that prompts you to start Equalizing. *Click on -> [Start Equalizing]* to start the Equalizing process.

Warning! Configuration parameters cannot be modified while equalizing is stopped. Start Equalizing if you want to modify configuration parameters.



Start Equalizing

You can also verify that the Equalizing process is started by going to the NetEqualizer Dashboard. *Click on -> Home* to return to the NetEqualizer Dashboard. The Equalizing Process is running if the Equalizing Status Indicator button is ON (GREEN).

Equalizing Rules On

Once the Equalizing Process is ON, you can easily determine if Equalizing Rules are ON. From the Configure Equalizing screen, scroll down until you see Equalizing Rules. Make sure that *Equalizing Rules = On*, as shown circled in blue below. This means that Equalizing Parameters are being applied to your network traffic.



Set Key Equalizing Parameters

Now you can configure **Equalizing Parameters**. There are two tabs: Key Equalizing Parameters, and Additional Equalizing Parameters. We will discuss both. After you change settings on each tab, scroll to the bottom of the window, and then [Click on -> \[Save Changes\]](#) to save or [Click on -> \[Reset\]](#) to discard.

As of Software Update 8.4, we offer configuration in Mbps, Kbps, or Bps. If you have not already set your Preferences, [Click on -> \[Setup\] -> Preferences](#) and set your Configuration Units. Then [Click on -> \[Setup\] -> Configure Equalizing](#) to return.

Using the RATIO Parameter to Influence Equalizing

Ratio (RATIO) (Percentage, Default = 85%)

NetEqualizer's **Ratio** parameter enables you to influence when equalizing is applied. The RATIO parameter refers to the network utilization on a percentage basis. RATIO can be set from 1 to 100. Our default value of 85 has the rules kick in when your network is at 85% utilization.

RATIO determines when Equalizing kicks in on your network trunk. This supplements any custom rules that you have set-up. When you *lower* RATIO, Equalizing will kick-in sooner, making equalizing more sensitive. When you *raise* RATIO, Equalizing kicks in later, making equalizing less sensitive.

To change the Ratio (RATIO) Parameter:

If you are not already on the Setup screen, from the NetEqualizer Dashboard, [Click on -> \[Setup\] -> Configure Equalizing](#). The Configure Equalizing window appears, defaulted to Key Equalizing Parameters Tab, as shown above. RATIO (circled in blue) defaults to 85. Type in: *your RATIO value* to change from the default. In our example above, we changed RATIO to 88%. For most cases, we recommend that you do not exceed 95% for RATIO.

Note: Throughout the GUI, as a visual reminder that you have unsaved changes, once you change a field it will be highlighted in yellow until you either Save Changes or Reset to cancel. You can either Save Changes now, or continue and save when you have finished all entries on this tab.

Why RATIO is Helpful

Sometimes the sheer volume of users on the network cannot be controlled by custom rules you have implemented. For example, setting a per-user hard limit of 2Mbps will prevent a user from going over the 2Mbps prescribed level. However, if 20 of your users get on at one time with large downloads, a DS3 trunk, for example, is quickly overwhelmed. To set custom rules, such as per-user hard limits, please see [Setting Bandwidth Limits](#)).



When Equalizing rules kick in at **RATIO** percent trunk utilization, they provide a unique safety valve for busy hours when your trunk gets full.

*Note: The **RATIO** parameter is applied to Equalizing your entire network trunk, and also to any [Bandwidth Pools](#) or [VLAN Limits](#) that you have established.*

Setting your Trunk Size

Bandwidth Up (TRUNK_UP) (Configuration Units set in [Preferences](#), Default = 15,360Mbps)
& Bandwidth Down (TRUNK_DOWN)

Set these parameters to the size of your network pipe for outbound traffic (**Bandwidth Up**) and inbound traffic (**Bandwidth Down**). NetEqualizer allows for different speeds for outbound and inbound links, as equalizing shapes bi-directionally. These parameters are set in whatever Configuration Units you selected in [Preferences](#).

Bandwidth Up and Bandwidth Down typically match your network capacity. These parameters are used by the NetEqualizer so it can react and take action when your trunk is nearing capacity, by starting [Equalizing](#) to shape your network traffic. Making either of these parameters *larger* than your actual trunk size will make the shaping rules less restrictive. Making them *smaller* than your actual trunk size will make them more restrictive. Alternatively, you can reduce [RATIO](#) to make shaping rules more restrictive.

The NetEqualizer defaults Bandwidth Up and Bandwidth Down to a 15,360 Mbps. However, you will see what bandwidth level you are licensed in the notes to the right of the parameter, as shown below. In our example, our license is for 750Mbps.

"Put in your uplink trunk size, in Mbps, here. A T1 is 1.5 Mbps. **The system is authorized to pass 750 Mbps.** Changes to this parameter will require equalizing to be stopped and restarted to take effect."

Click on -> [Setup] -> [Configure Equalizing](#) to bring up the Configure Equalizing window. Type in: *your **Bandwidth Up & Down values*** to set your trunk size, which should be less than or equal to ($\leq$) your license. In our example, we set both to 750Mbps.

Note: Bandwidth Up & Down do not enforce the link speed from your provider. We assume your provider has already enforced your contracted speed.

Whenever you change your Trunk Size, you will need to restart the Equalizing process for changes to take effect. You can either Save Changes now, or continue and save when you have finished all entries on this tab.

To Save Changes and restart now: *Click on -> [Save Changes]*. This message appears:



✔ **Success!** The following configuration parameters have been updated:

- Bandwidth Up (TRUNK_UP): 750
- Bandwidth Down (TRUNK_DOWN): 750
- Restart Equalizing to apply this update.



Restart Equalizing

Now **Click on** -> **[Restart Equalizing]** to restart the Equalizing process. Once complete, you will see the following message at the top of the screen. You can click on the "x" at the right of the message to dismiss it.

✔ **Success!** Equalizing has been restarted.



Note: We display messages at the top of the screen throughout the GUI. Error messages are red, warnings are yellow, and success messages are green.

Key Parameter to Adjust Equalizing Sensitivity

The last Key Equalizing Parameter that needs to be set is Hog Minimum. From the Setup menu, **Click on** -> **[Setup]** -> **Configure Equalizing**.

Hog Minimum (HOGMIN)

(Configuration Units set in [Preferences](#), Default = 1Mbps)

HOGMIN defines the minimum traffic level for which connections will not be penalized. In other words, a connection using less bandwidth in megabits per second than this number will never get penalized. The default value of 1 megabit per second will ensure that most business-critical traffic, such as VoIP, web browsing, and web applications, are never accidentally throttled back when NetEqualizer reaches a congestion threshold, as they will be below Hog Minimum.

With larger network pipes, you should raise Hog Minimum to allow more traffic types to pass without being penalized. Here are our recommended settings for HOGMIN, based on network size:

Network Size	HOGMIN
< 50Mbps	.5 Mbps
>= 50Mbps to < 200Mbps	.75 Mbps
>= 200Mbps to < 1Gbps	1 Mbps
>= 1Gbps	2 Mbps

Type in: **your preferred Hog Minimum value** to set Hog Minimum. In our example, we kept Hog Minimum at 1Mbps, which is the recommended value for our 750Mbps pipe.

As we have now completed all entries on the Key Parameters Tab, **Click on** -> **[Save Changes]** to save or **Click on** -> **[Reset]** to discard.

Configure Equalizing

Home / Setup / **Configure Equalizing**

KEY EQUALIZING PARAMETERS
ADDITIONAL EQUALIZING PARAMETERS

Use this page to configure key parameters that control equalizing.

Ratio (RATIO)

88

% utilization

Bandwidth Up (TRUNK_UP)

750

Mbps

Bandwidth Down (TRUNK_DOWN)

750

Mbps

Hog Minimum (HOGMIN)

1

Mbps

Equalizing Rules (DEFAULT_RULES)

☒ On
 ☐ Off

Save Changes

Reset



Set Additional Equalizing Parameters

You may wish to set **Additional Equalizing Parameters**, which are shown at right.

From the NetEqualizer Dashboard, *Click on -> [Setup] -> Configure Equalizing* to bring up the Configure Equalizing window. *Click on -> Additional Equalizing Parameters* (circled in blue at right) to open the Additional Equalizing Parameters tab.

There are four (4) Additional Equalizing Parameters that you can configure. You may need to tune your Penalty Unit and Connection Tracking Table Size to optimize for your network pipe. Maximum Penalty and Inactive Tics are rarely changed. We will review each below, and offer tips and recommendations on what you should set these parameters to, based on your network size.

After you change settings on this tab, scroll to the bottom of the window, and then *Click on -> [Save Changes]* to save or *Click on -> [Reset]* to discard.

Additional Parameters to Adjust Equalizing Sensitivity

In some instances, Equalizing defaults may need to be custom tuned for sensitivity. For example, if streaming music feeds break midstream at times when the total usage on the trunk is light, it might be because Equalizing is tuned to be too sensitive.

Penalty Unit (PENALTY_UNIT)

(100ths of seconds, Default = 5)

PENALTY_UNIT is the minimum penalty that will be inflicted on a packet when a penalty is set up on an IP address. Values for this variable are integers in the range 1 -100, with 1 being the least restrictive.

PENALTY_UNIT is the unit of time that NetEqualizer will start with when delaying a packet of Internet data. It iteratively increases penalties by this value should a "hog" not respond to the initial penalty. By increasing the size of this parameter, the NetEqualizer will scale back hogs more quickly. Note that the higher your network speed, the more sensitive it is to PENALTY_UNIT. The default value of 5 will work fine on any network, but if you see the NetEqualizer slowing streams too severely, you may want to reduce this value.

Here are the recommended settings for PENALTY_UNIT, based on network size:

Network Size	PENALTY_UNIT
>= 10Mbps to < 50Mbps	2 or 3
* >= 50Mbps	1



* Networks much larger than 50Mbps may require a PENALTY UNIT resolution smaller than 100ths of seconds. In the NetEqualizer Web GUI, the smallest penalty that can be applied to an IP Packet is 1/100 of a second. If you are finding that a default PENALTY of 1 is putting too much latency on your connections, then you can adjust the PENALTY unit to 1/1000 of second with the following command:



From the NetEqualizer Dashboard, *Click on ->[Maintenance] -> Troubleshooting Tools -> Run a Command.*

Type in: */bridge/bridge-utils/brctl/brctl rembrain my 99999*

Maximum Penalty (MAX_PENALTY)

(100ths of seconds, Default = 140)

This parameter is rarely changed from the Default. If it is changed, it needs to be set to a value that is greater than Penalty Unit.

This is the maximum delay that NetEqualizer will allow. NetEqualizer increments a delay by the value of PENALTY_UNIT every few seconds in the event a connection continues to use excessive bandwidth, until MAX_PENALTY is reached. A MAX_PENALTY of 200 (2 seconds) usually kills the connection altogether, as most servers on the Internet give up communicating when communications lag for more than two seconds.

Moving Average (MOVING_AVG)

(no longer on Web GUI)

No longer visible on the NetEqualizer web GUI, as this is rarely changed. MOVING_AVG keeps NetEqualizer from penalizing short bursts of activity. For example, if this variable is set to 8 and the network is hit with a burst of 8000 bytes over a second from an IP address, the moving average for the second would be 8000/8 or 1000 bytes. If the burst persisted for four seconds, the average would be 32000/8 or 4000 bytes.

The larger this number, the longer a burst can be before it gets penalized. Note that if this parameter is set too high, nothing will ever get penalized. The preset value for MOVING_AVG from our factory-delivered NetEqualizer is designed to handle any size network and need not be changed.

Note: If you manually edit the NetEqualizer configuration file, you will see MOVING_AVG in the configuration. Please keep it set to its default value.

ANCIENT

(no longer on Web GUI)

Ancient is no longer visible on the NetEqualizer GUI. The Ancient parameter is how long to keep a penalty in effect, in seconds. The preset value for ANCIENT from our factory-delivered NetEqualizer is designed to handle any size network and need not be changed.

Note: If you manually edit the NetEqualizer configuration file, you will see ANCIENT in the configuration. Please keep it set to its default value of 20.

HOGMAX

(no longer on Web GUI)

Legacy variable, no longer visible on the NetEqualizer Web GUI, and no longer used.

Note: If you manually edit the NetEqualizer configuration file, you will see HOGMAX in the configuration. Please keep it set to its default value.



Parameters to Size Internal Tables

If you are not already on the Configure Equalizing screen, from the NetEqualizer Dashboard, [Click on -> \[Setup\] -> Configure Equalizing -> Additional Equalizing Parameters](#). You should now be on the Additional Equalizing Parameters Tab. Type in a new value to modify the following parameters to size internal tables:

Connection Tracking Table Size (# of connections to track in 1 second, Default=10,000)
(BRAIN_SIZE)

Connection Tracking Table Size determines how many connections (IP pairs) the NetEqualizer watches at one time during any given second. NetEqualizer keeps a mini-history of the activity of all users on a trunk. It uses this database to make decisions on who is using too much bandwidth.

Here are recommended settings for Connection Tracking Table Size, based on network size:

Network Size	Connection Tracking Table Size (BRAIN_SIZE)
< 100Mbps	20,000
>= 100Mbps to < 1Gbps	30,000
>= 1Gbps	40,000

Type in: [a new value for BRAIN_SIZE](#) to change this parameter.

Note: NetEqualizer can handle up to 5 million or more connections every minute. We point this out as many customers compare our connection ability with that of their Router, which uses a timeframe of minutes.

Inactive Tics (INACTIVE_TICS) (units are hundredths of seconds, Default = 200)

This parameter is rarely changed from the Default. This is how long an entry in the Connection Tracking Table (BRAIN_SIZE from above) will live before being removed if no activity is detected. Generally, we are not interested in connections that are idle. For example, a value of 200 for this parameter instructs the NetEqualizer to "cancel" tracking a connection after 2 seconds.

A recommended maximum value of 800 is 8 seconds. The minimum is 100 (or 1 second).

Type in: [value for INACTIVE_TICS](#) to change this value.

BUFFERS

(no longer on Web GUI)

No longer visible on the latest NetEqualizer Web GUI, as it is rarely changed. BUFFERS control the number of connections that can simultaneously be penalized (slowed down). When NetEqualizer sets a penalty on a connection, it assigns a "delay" buffer to this connection to slow it down. NetEqualizer reserves a finite number of "delay" buffers when it powers up. The preset value for BUFFERS from our factory-delivered NetEqualizer is designed to handle any size network and need not be changed.

Note: If you manually edit the NetEqualizer configuration file, you will see BUFFERS in the configuration. Do not change this value; changes will be ignored.

As we have now completed all entries on the Additional Equalizing Parameters Tab, [Click on -> \[Save Changes\]](#) to save or [Click on -> \[Reset\]](#) to discard.



Viewing your Equalizing Parameter Settings



Once you have set all your Equalizing Parameters, you should view them in your NetEqualizer Configuration file to make sure that they are set as you expect.

On the Navigation Menu, [Click on -> \[RTR\] -> Configuration](#).

Current parameter settings are listed using the [SYSTEM PARAMETER NAME]. For example, you can see **RATIO** is now set to 88% (circled in blue) on the screen below. You can scroll through the Configuration screen to view all of your parameter settings.

The screenshot shows the 'Configuration' page in the NetEqualizer interface. It features a table with columns 'Line #', 'Parameter', and 'Value'. The 'RATIO' parameter is highlighted with a blue circle. The table lists various system parameters and their current values.

Line #	Parameter	Value
1	STATE	ON
2	ETHA	eth0
3	ETHB	eth1
4	NAME	NetEqualizer
5	NAME PREFIX	10
6	RATIO	88
7	PENALTY LIMIT	5
8	MAX PENALTY	140
9	QUEUE LIMIT	15
10	BUFFERS	900
Line #	Parameter	Value

Showing 1 to 10 of 35 entries (processed in 0.00040197372436523 s)

Equalizing in Action

Once NetEqualizer is installed and running, there are several reports available in RTR to help you monitor and analyze how NetEqualizer is responding to your network's traffic.

To see the current penalties being applied to your network, from the Dashboard:

[Click on -> \[RTR\] -> Active Connections -> View Active Penalties](#).

To see a history of the penalties being applied, from the Dashboard:

[Click on -> \[RTR\] -> Traffic History -> General Penalty Reports](#).

When your network is experiencing moderate to heavy use, you will see entries containing the word **PENALTY** followed by two IP addresses in the NetEqualizer Log. To view the NetEqualizer Log, [Click on -> \[RTR\] -> NetEqualizer Log](#).

PENALTY indicates that NetEqualizer's built-in fairness rules have determined that the communication link between these two IP addresses (a connection) is using too much bandwidth, so NetEqualizer has issued a penalty against this connection. The penalty causes all data on that connection to slow down. At periodic intervals, if NetEqualizer determines that this connection is still using too much bandwidth, it will increase the delay on the connection. The **PENALTY** will be removed in a few seconds should the congestion on your Network subside. Log entries are discussed in more detail in the [NetEqualizer Log](#) section.



Setting Bandwidth Limits

Bandwidth Limits can be set if you want to restrict the amount of bandwidth a specific IP address or set of IP addresses can use. This is done through the use of **Bandwidth Limiting Rules** to carve out maximum bandwidth usages for a particular subscriber base.

For example, a college network administrator may want to set up separate bandwidth usage categories with bandwidth limits for each of three subnets: 1) students, 2) faculty, and 3) administrators. This would be accomplished by using the "Pools" Bandwidth Limiting Rule, which can be used to set up three separate Pools, each with their own bandwidth limit:

Pool1 = students 100Mbps up/120Mbps down

Pool2 = faculty 50Mbps up/60Mbps down

Pool3 = administrators 25Mbps up/30Mbps down

Bandwidth Limiting Rules define and restrict the amount of bandwidth a specific IP address or set of IP addresses can use. Bandwidth Limits do NOT physically reserve bandwidth on your network. They are used to set a virtual ceiling or limit for an IP or group of IPs. There are multiple ways to configure this in the NetEqualizer, to best meet your needs. The Bandwidth Limiting Rules are as follows, and will be discussed in detail below:

Bandwidth Limiting Rules

- | | |
|--|--------------------------------------|
| 1. Hard Limits by IP | - Individual limits by IP or subnet. |
| 2. Adding Bursting by IP | - Burst a Hard Limit by IP. |
| 3. Bandwidth Pools | - Shared limits by IP or subnet. |
| 4. VLAN Shared Limits | - Shared limits by VLAN. |

Note: As Bandwidth Limiting Rules are allocations, not reservations, bandwidth is not set aside for them. Rather, network bandwidth is available for use at all times.



Configure Hard Limits by IP

[\(back\)](#)

Use Hard Limits to set a fixed amount of individual bandwidth to a single IP address or an entire set of IP addresses specified by a subnet mask.

Subnetted Hard Limits are not shared bandwidth. EACH IP address in the hard-limited subnet range will receive the specified hard limit. For example, if you set up a 3Mbps up/5Mbps down for a subnet, each IP address will get 3Mbps/5Mbps to use. Subnetted Hard Limits can be set up for a Class B subnet, Class C subnet, or any legal subnet value 1-32.

NetEqualizer allows up to 60 thousand (60,000) unique active Hard Limits.

From the NetEqualizer Dashboard or Navigation Menu [Click on ->\[Setup\] -> Manage Traffic Limits -> Configure Hard Limits](#). The Configure Hard Limits screen opens, shown here.

Configure Hard Limits

Home / Setup / Manage Traffic Limits / **Configure Hard Limits**

Use this page to configure hard limits by IP address.

- Use [Quick Edit Hard Limits](#) to add and delete hard limits by IP address without having to restart equalizing.
- IP addresses are described using [CIDR notation](#) -- a CIDR of 32 means "this IP address only".
- An IP address cannot be included in a hard limit range and also have an individual hard limit rule.
- An IP address cannot have a hard limit and also be a member of a pool.
- Changes to hard limits by IP address will require stopping and restarting the equalizing process to take effect.
- No changes are made to the hard limits until they are explicitly saved.

	Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor
1	192.168.1.112	/	32	10	5	1
2	210.1.20.1	/	24	5	3	1

[Save Changes](#) [Reset](#)

Configure Hard Limits is a Batch Entry Screen, where you can add or edit many rules at once and then save your changes, saving you time in the setup process. Click on the dark blue "+" buttons to add entries, red "x" to delete entries (circled in blue above), and the blue "up/down arrows" (circled in orange) to reorder Hard Limits. Once you have entered, modified, or deleted Hard Limits, [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes.

In the light blue box above the entry fields, you will find notes to help you in creating Hard Limits. Links are in orange, which you can click on to either get more information, or to move to another screen.

Once you save changes, you will be prompted to Restart Equalizing. Your Hard Limits will NOT take effect until you restart equalizing, even if you see them in your NetEqualizer Configuration.

We have also maintained your ability to add or delete a rule without restarting Equalizing. In Software Update 8.4, this maintained in our Quick Edit screens. From the



Configure Hard Limits batch entry screen, you can move to the [Quick Edit Hard Limits](#) screen, to add or delete Hard Limit Rules without having to restart equalizing. Quick Edit cannot be used to modify rules.

Hard Limit Rules

Before creating Hard Limits, it is important for you to understand several rules that apply.

A Hard Limited IP cannot also be in a Bandwidth Pool

Once you put a Hard Limit on an IP address, you cannot also have that IP address exist within a Bandwidth Pool. An IP address either has a fixed amount of bandwidth (Hard Limit) or shared bandwidth (Bandwidth Pools), but not both. See [Bandwidth Pools](#) to determine which option is right for your needs.

Override Hard Limit for an IP in a subnetted Hard Limit Range

You can set a Subnet Hard Limit, which applies to an entire set of IPs, such as a /24 or /16, and also set an *override hard limit for an individual IP in that range*. It is important that the Subnet Hard Limit base address does NOT match the override individual IP.

For example, this is allowed:

Subnet Hard Limit = 10.1.1.0 /24

Override Hard Limit = 10.1.1.143 /32

This is NOT allowed:

Subnet Hard Limit = 10.1.1.143 /24

Override Hard Limit = 10.1.1.143 /32

When setting up Subnet Hard Limits, we recommend using .0 for the base address. This means that you can have override Hard Limits for any IP except 10.1.1.0 (using our example). This should not be a problem, as DHCP does not usually assign out .0 in common practice.

Hard Limiting an External IP address

As of Software Update 8.4, you can assign a Hard Limit to an IP address that physically sits on the WAN side of the NetEqualizer. For example, you can put a Hard Limit on an Internet site that is used to host video. This is useful in conjunction with setting Priority Traffic, for cases where you want your hosted video to not be equalized, but do not want it to run completely unrestricted. For example, you could set a hard limit of 20Mbps Up/Down, so that your hosted video could only consume a maximum of 20Mbps of your network pipe.

Tuning for a Large Number of Subnetted Hard Limits

If you plan to set up a large number of subnet-ranged Hard Limits (≥ 32 subnet ranges), you will need to set several tuning parameters. Please see [Appendix #4](#) for detailed instructions.



Creating a Hard Limit

	Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor
1	192.168.1.112	/	32	10	5	1
2	210.1.20.1	/	24	5	3	1

Save Changes Reset

To create a Hard Limit, *click on your computer's TAB to put the focus into the Host IP address field*. Type in: *Host IP address* in 11.22.33.44 format.

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B use /16, for a Class C use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will get the Hard Limit.

TAB to Download Mbps Hard Limit field. Type in: *a Positive Number*. TAB off the field to complete your entry.

TAB to Upload Mbps Hard Limit field. Type in: *a Positive Number*. TAB off the field to complete your entry.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. This will be fixed in a future release.

Click on any of the "+" icons (circled in blue) to continue adding Hard Limits. Any unsaved fields are shown in yellow, as a visual reminder that you have not yet saved your changes. If you were to click off this screen before saving changes, your data would be cleared.

Once you have entered all of your limits, *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard changes. We are now prompted to Restart Equalizing, for the new Hard Limits to take effect. Once we *Click on -> [Restart Equalizing]*, our new Hard Limits will be available to the NetEqualizer process.

In our example, we have created two Hard Limit (HL) rules. The first is a HL of 10Mbps Down/5Mbps Up for individual IP address 192.168.1.112 (/32). The second is a subnet HL (/24) of 5Mbps Down/3Mbps Up for all IP addresses within 210.1.20.

Modifying a Hard Limit

As of Software Update 8.4, we provide you with the ability to modify your Hard Limits, using the Configure Hard Limit screen. Additionally, you can make multiple edits at once, and then *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard.

From the NetEqualizer Dashboard or Navigation Menu *Click on -> [Setup] -> Manage Traffic Limits -> Configure Hard Limits*. The Configure Hard Limits screen opens, shown below.

In our example, we first changed our subnetted HL CIDR from /24 to /16. We also changed the Upload Mbps from 3 to 4. These fields turn yellow, showing that we have updated the Hard Limit but not saved it.



			Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor		
		1			192.168.1.112	/	32	10	5	1
		2			210.1.20.0	/	16	5	4	1

Save Changes Reset

As we are done with our changes, we will now [Click on -> \[Save Changes\]](#) to save them. As shown below, our HL rows show no remaining yellow fields, as our changes are saved. We are now prompted to Restart Equalizing, for the updated Hard Limit rules to take effect. Once we [Click on -> \[Restart Equalizing\]](#), our new HL settings will be available to the NetEqualizer process.

Success! Hard limits by IP address have been updated. Restart Equalizing to complete the hard limits by IP address update.

			Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor		
		1			192.168.1.112	/	32	10	5	1
		2			210.1.20.0	/	16	5	4	1

Save Changes Reset

Deleting a Hard Limit

In order to remove a Hard Limit, click on the red "x" button on the row that you wish to delete. The HL row will disappear. To completely remove the HL from the configuration file, you need to [Click on -> \[Save Changes\]](#). We are now prompted to Restart Equalizing, for the Hard Limit rule to be completely removed from the NetEqualizer Configuration. Once we [Click on -> \[Restart Equalizing\]](#), our HL will be permanently deleted.

In our example, we have clicked on the red "x" next to the 2nd row, to delete the subnetted HL for 210.1.20.0 /16. You can see on the screen below that only the CL for 192.168.1.112 remains.

			Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor		
		1			192.168.1.112	/	32	10	5	1

Save Changes Reset

Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a Hard Limit has been removed, see the following [View Traffic Limits](#) section of this User Guide.



Adding Bursting to Hard Limits



([back](#))

In addition to setting a Hard Limit by IP address, as of [software update 4.7](#), we have enabled "bursting" above the Hard Limit. Prior to the bursting feature, the top speed allowed for each user was fixed at the set Hard Limit.

Now with bursting, a user can be allowed a burst of bandwidth for up to 10 seconds at two, three, four, or any multiple of their base Hard Limit. For example, if a user has an incoming base Hard Limit of 2 megabits a second, and a burst factor of 4, then their inbound connection will be allowed to burst all the way up to 8 megabits for 10 seconds (2Mbps HARD LIMIT x 4 BURST FACTOR = 8Mbps inbound BURST LIMIT), at which time it will revert back to the original 2 megabits per second. If the outgoing base Hard Limit was set to 1 megabit per second, with the same burst factor, the outbound BURST LIMIT would be 4Mbps. This type of burst will be noticed when loading large web pages loaded with graphics. From a user's perspective, they will essentially fly up in the browser at warp speed.

In order to make bursting a "special" feature, it obviously can't be on all the time. For this reason, by default the NetEqualizer will force a user to wait 80 seconds before they can burst again.

Setting up Bursting on an IP Address

From the NetEqualizer Dashboard or Navigation Menu [Click on ->\[Setup\] -> Manage Traffic Limits -> Configure Hard Limits](#). The Configure Hard Limits screen opens, shown below.

			Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor		
		1			192.168.1.112	/	32	10	5	4
		2			210.1.20.0	/	16	5	4	1

[Save Changes](#) [Reset](#)

The last field on each Hard Limit row is the Burst Factor. When you create a Hard Limit, by default the Burst Factor is set to 1 (no bursting). Leave this field set to 1 for no bursting, or set to a multiple greater than 1 for bursting. BURST FACTOR is multiplied times the incoming and outgoing HARD LIMITs to arrive at the BURST LIMITs (default speed you wish to burst up to). BURST FACTOR must be a positive integer. If you try to set BURST FACTOR to zero, a fraction, or a negative number you will get an error message when you Save Changes, as shown here.

⚠ **Oops! The following errors were found:**

- limit 1: burst factor must be a positive integer (0)

For our example above, we have changed the Burst Factor to 4 on IP address 192.168.1.112/32. With bursting, the Hard Limits will look as follows when bursting:

10Mbps download HARD LIMIT x 4 BURST FACTOR = 40Mbps download BURST LIMIT
5Mbps upload HARD LIMIT x 4 BURST FACTOR = 20Mbps upload BURST LIMIT

Once you have changed your Burst Factors, [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes.

Note: Once bursting has been set-up, bursting on an IP address will start when that IP exceeds its



rate limit across all connections for the IP. The burst applies to all connections across the IP address.

Global Burst Parameters

There are two global burst parameters that apply to all bursting that you have setup. These are BURST DELAY, time between bursts, and BURST DURATION, how long a burst lasts. BURST DELAY defaults to 80 seconds. BURST DURATION defaults to 10 seconds.



To change the global burst parameter defaults:

From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[Maintenance\] -> Troubleshooting Tools -> Run a Command](#). Type in: `/usr/sbin/brctl setburstparams my BURSTDELAY BURSTDURATION`

For example, assuming the BURST DELAY to be 40 and BURST DURATION to be 30, our command would be `/usr/sbin/brctl setburstparams my 40 30`

The first parameter, BURST DELAY, is the time an IP must wait (in seconds) before it can burst again. If an IP has done a burst cycle it will be forced to wait this long in seconds before it can burst again. BURST DELAY is defaulted to 80 seconds. This means that an IP address will wait 80 seconds after its last burst duration completes before bursting again.

The second parameter, BURST DURATION, is the time an IP will be allowed to burst (in seconds) before being relegated back to its base Hard Limit.

Note: At this time, the global burst parameters are not persistent. You will need to put a command in the NetEqualizer Autostart File, as shown below, if you want them to stick between reboots.




To edit the Autostart File:

From the Maintenance and Reference Menu, [Click on -> Maintenance -> Manage Autostart File](#). The Manage Autostart File window appears, defaulted to the Edit Autostart File tab, as shown below.

Add the SAME line that you created above to the bottom of the Autostart File to change your global burst parameter settings. In our example, circled in blue below, we typed in: `/usr/sbin/brctl/setburstparams my 40 30`

Then [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes.





Bursting and Speed Tests

With the default settings of 10 second bursts and an 80 second time out before the next burst, it is unlikely a user will be able to see their full burst speed accurately with a speed test site. The easiest way would be to extend the burst time to minutes, instead of the default 10 seconds, and then run the speed test.

With the default set at 10 seconds, the best way to see a burst in action is to take a [continuous snap shot](#) of an IP's consumption during an extended download.

Note: Before you implement bursting, you may want to consider the downside of bursting. See our [bursting blog article](#) on this subject.

Viewing your Hard Limits

There are two places in the NetEqualizer where you can view your Hard Limits. We will discuss both here.



View All Traffic Limits

From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[Setup\] -> Manage Traffic Limits -> View All Traffic Limits](#). Hard Limits are the first section of the report. You will see a list of your Hard Limits, as shown in the following excerpt from the report. From this report, you can also navigate back to the Hard Limits batch entry, by clicking on the dark blue "Edit" button.

HARD LIMITS							Edit
	Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor	
1	192.168.1.112	/	32	10	5	1	
2	210.1.20.0	/	16	5	4	1	



Configuration File

You can also see your Hard Limits in the NetEqualizer Configuration File. From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[RTR\] -> Configuration](#). If your configuration file is more than 25 lines, you may need to click on the "Last" button to go to the end of your file, as we did below.

38	HARD	192.168.1.112 1250000 625000 32 1
39	HARD	210.1.20.0 625000 500000 16 1
Line #	Parameter	Value
Showing 26 to 39 of 39 entries (processed in 0.00053000450134277 s)		
First Previous 1 2 Next Last		

There will be one row for each Hard Limit. The row will be classified as "HARD", followed by the IP address or subnet, and then the Download Limit, Upload Limit, CIDR, and Burst Factor. The Download and Upload limits are shown in bytes per second. In our example above, you can read the first line as follows: a Hard Limit on 192.168.1.112, 1,250,000Bps (10Mbps) download limit, 625,000Bps (5Mbps) Upload Limit, for an individual IP (/32) and no bursting (1).

Note: Pools look similar to Hard Limits in the Configuration File. You can tell apart by the Pool IP address, which is made up of the Pool# repeating. For example, Pool 1 would be shown as HARD 1.1.1.1/32 2250000 2250000 100001 1.



Setting up Bandwidth Pools

[\(back\)](#)

A Bandwidth Pool (Pool) is a collection of IP addresses that *share* a bandwidth allocation. The sum total of bandwidth for all of the IP addresses will not be allowed to exceed more than the total bandwidth allocated to the Pool. For example, if four IP addresses are assigned to a pool, and the pool download bandwidth limit is set at 10Mbps, then the total download bandwidth for all four IPs together is 10Mbps (the total, not per IP). *Pools are a bandwidth restriction, not a reservation.*

Think of a Pool as a "**virtual NetEqualizer**". You can group users into logical trunks by IP address and apply [equalizing technology](#) to each logical group (Pool). Equalizing is performed in the same fashion as across your entire network trunk; but in this case it equalizes *within* the Pool. When the total bandwidth threshold for that Pool is reached, determined by the RATIO parameter, then any large connections (over HOGMIN) associated with IP addresses within the Pool will be penalized. For our example, if RATIO was 85%, then equalizing would occur on the 10Mbps Pool when download bandwidth hits 8.5Mbps.

The bandwidth restriction on a Pool may fluctuate a bit depending on the type of traffic. Heavy use of UDP traffic tends to run over the limit, and heavy TCP/IP traffic (FTP for example) will tend to be held below the limit.

Pools were added for network topologies where bandwidth congestion is occurring at nodes in the network, not necessarily at the WAN/LAN connection. For example, this could be occurring in a wireless network where bandwidth congestion occurs at the wireless hotspots or in the backhaul connections. Individual bandwidth pools can be defined with the IPs of users at each hotspot and equalizing applied per hotspot.

One example of using Pools is to accommodate cases where bandwidth is advertised and sold by Internet Providers as "you are one of n customers sharing x bandwidth". Another case to use bandwidth pools is to set up equalizing at the subnet level. For example, a university may split their network into faculty, administrators, and student subnets. Each of these subnets could be defined as a bandwidth pool, with separate upload/download speeds that are shared by all users in the pool.

From the NetEqualizer Dashboard or Navigation Menu [Click on ->\[Setup\] -> Manage Traffic Limits -> Configure Pools and VLANs ->By Pool tab](#). The following screen opens.

Configure Pools and VLANs

Home / Setup / Manage Traffic Limits / **Configure Pools and VLANs**

BY POOL **BY VLAN**

Use this page to configure limits by pool.

- Use [Quick Edit Pools](#) to add pools and to add and delete pool members without having to restart equalizing.
- Pool numbers are integers in the range 1-250 and must be unique.
- Click [+](#) below to expand a pool to view and/or modify its members.
- An IP address cannot be a member of more than one pool.
- An IP address cannot be a member of a pool and also have a hard limit.
- Changes to limits by pool will require stopping and restarting the equalizing process to take effect.
- No changes are made to the limits until they are explicitly saved.

Pool Number	Members	Pool Name	Download Mbps	Upload Mbps
1	3	corporate	18	18
2	1	wireless	12.04	12
3	1	bldg 1	10.24	2.048

[Save Changes](#) [Reset](#)



Configure Pools and VLANs is a Batch Entry Screen, where you can add or edit many rules at once and then save your changes, saving you time in the setup process. Click on the dark blue "+" button at the top to add a Pool, or the red "x" next to each row to delete a Pool (both circled in blue above).

As order is not important with Pools, the blue "up/down arrows" (circled in orange) are NOT used to reorder Pools. Rather, they are used to *add members (IP addresses and/or subnets)* to Pools. Use the up/down arrows at the top of the page to expand ALL Pools, or the up/down arrows on each line to expand ONE Pool individually.

Once you have entered, modified, or deleted Pools or Pool Members, *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard changes.

In the light blue box above the entry fields, you will find notes to help you in creating Pools. Links are in orange, which you can click on to either get more information, or to move to another screen.

Once you save changes, you will be prompted to Restart Equalizing. Your new Pools or additional Pool members will NOT take effect until you restart equalizing, even if you see them in your NetEqualizer Configuration.

We have also maintained your ability to add a Pool, or add/delete Pool Members without restarting Equalizing. In Software Update 8.4, this maintained in our Quick Edit screens. From the Configure Pools and VLANs batch entry screen, you can move to the [Quick Edit](#) Pools screen, to add or delete Pool Members or Add a Pool without having to restart equalizing. Quick Edit cannot be used to modify rules.

Bandwidth Pool Rules

Before creating Pools, it is important for you to understand several rules that apply.

Bandwidth Pools cannot overlap with Hard Limits by IP

Once you add an IP address to a Bandwidth Pool, you cannot also have a Hard Limit on that IP address. An IP address either has a fixed amount of bandwidth (Hard Limit) or shared bandwidth (Bandwidth Pools), but not both. See [Configure Hard Limits by IP](#) to determine which option is right for your needs.

An IP address can only exist in one Pool

You cannot put an IP address or IP subnet into more than one Pool at a time.

Priority Traffic and Pools

If you create a Priority Traffic rule for an IP address, and the IP address exists within a Bandwidth Pool, it will receive priority over other IP addresses within the pool.

Tuning for a Large Number of Subnet-ranged Pools

If you plan to set up a large number of subnet-ranged Pools (≥ 32 subnet ranges), you will need to set several tuning parameters. Please see [Appendix #4](#) for detailed instructions.

You can define up to 250 Pools

In [Software Update 8.4](#) and above, Bandwidth Pools can number from 1 to 250; up to 250 different bandwidth pools are allowed per NetEqualizer.

You can add individual IP addresses or entire subnets to Bandwidth Pools

IP addresses within a Bandwidth Pool need not be contiguous. You can add members to a Bandwidth Pool in any order.



Creating a Pool

	Pool Number	Members	Pool Name	Download Mbps	Upload Mbps
1	1	3	corporate	18	18
2	2	1	wireless	12.04	12
3	3	2	bidg 1	10.24	2.848
4	4	1	another pool	20	10

	IP Address	CIDR
1	192.168.1.1	24

To create a Pool, *Click on -> dark blue "+" button at the top of the screen* (circled in blue above) to add a new blank row. In our example above, a blank Row 4 and a blank Pool Member row tied to Row 4 (large orange circle above) are added.

Then *click on your computer's TAB to put the focus into the Pool Number field*. Type in: *A Positive Integer between 1-250* to assign a number to your Pool. We typed in "4". The window displays the current Pool #'s in use, so that you do not reuse a Pool #.

If your Pool # is not in the range, you will see the following error when you click on Save Changes.

Oops! The following errors were found:

- limit 4: pool number must be an integer in the range 1-250 ()

TAB to Pool Name field. This field is OPTIONAL. If you define a Pool Name, it will be used in Real-Time Reporting on your Pool reports. Type in: *letters, numbers, and spaces* to assign a Pool Name. No special characters (##@!) are allowed. We typed in "another pool". You will see the following error when you click on Save Changes if your Pool Name is not valid.

Oops! The following errors were found:

- limit 4: poolname must be letters, digits and spaces only (!! special characters #)

TAB to Download Mbps field. Type in: *a Positive Number*. You can use fractional values, such as 12.04, as shown above for Pool 2. We typed in 20.

TAB to Upload Mbps field. Type in: *a Positive Number*. You can use fractional values, such as 2.848, as shown above for Pool 3. We typed in 10. TAB to complete your entry.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. For this reason, if you are adding a lot of Pools, you may want to save changes after creating each Pool with its associated Pool Members. This will be fixed in a future release.

You can *Click on -> [Save Changes]* to save your Pool and add Pool Members later, or you can add your Pool Members now. *Click on -> [Reset]* to discard changes.

In our example above, we now have Pool4, also known as "another pool", with a shared allocation of 20Mbps Download and 10Mbps Upload. Pool 4 will be equalized at RATIO, to share bandwidth fairly among pool members. In our next step, we will add Pool Members to Pool 4.



Adding Pool Members

TAB to IP Address of the Pool Member row (large orange circle above). Type in: *IP address* in 11.22.33.44 format.

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B, use /16, for a Class C, use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will share the Pool Limits.

Click on any of the "+" icon (small orange circle above) to continue adding Pool Members. In our example show below, a second blank Pool Member row is now added for us to add another Pool Member to Pool 4. Any unsaved fields are shown in yellow, as a visual reminder that you have not yet saved your changes. If you were to click off this screen before saving changes, your data would be cleared.

Pool Number	Members	Pool Name	Download Mbps	Upload Mbps
4	1	another pool	20	10

IP Address	CIDR
192.168.1.1	24
	32

Save Changes Reset

Once you have entered all of your Pool Members, *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard changes.

In our example, Pool 4 now has one Pool Member, a /24 subnet on 192.168.1.1. All IPs in the /24 range will share the 20Mbps Down /10Mbps Up assigned to Pool 4.

We are now prompted to Restart Equalizing, as show below, for the new Pools and Pool Members to take effect. If you are NOT yet done adding Pools, dismiss the message by clicking on the "x" (circled in blue below) and continue adding your Pools and Pool Members. If you are with your additions, *Click on -> [Restart Equalizing]*, and your new Pools & Pool Members will be available to the NetEqualizer process.

Success! Limits by pool have been updated. Restart Equalizing to complete the limits by pool update.

Restart Equalizing

Pool Number	Members	Pool Name	Download Mbps	Upload Mbps
1	1	corporate	18	18
2	1	wireless	12.04	12
3	2	bldg 1	10.24	2.848
4	1	another pool	20	10

Save Changes Reset



Modifying a Pool or Pool Member

As of Software Update 8.4, we provide you with the ability to modify your Pools and Pool Members, using the Configure Pools and VLANs screen. Additionally, you can make multiple edits at once, and then *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard.

From the NetEqualizer Dashboard or Navigation Menu *Click on -> [Setup] -> Manage Traffic Limits -> Configure Pools and VLANs*. The Configure Pools and VLANs screen opens.

In our example, shown below, we first changed the Pool 4 Pool Name to "bldg. 3". We also changed the Upload Mbps from 10 to 15. And finally, we changed the Pool Member subnet from 192.168.1.1 to 192.200.1.1. These fields turn yellow, showing that we have updated the Pool and Pool Member but not saved them.

As we are done with our changes, we will now *Click on -> [Save Changes]* to save them. We are prompted to Restart Equalizing, for the updated Pool Limit rules to take effect. Once we *Click on -> [Restart Equalizing]*, our new Pools settings will be available to the NetEqualizer process.

Deleting a Pool or Pool Member

In order to remove a Pool or Pool Member, click on the red "x" button on the row that you wish to delete. The Pool or Pool Member row will disappear. To completely remove the Pool or Pool Member from the configuration file, you need to *Click on -> [Save Changes]*. We are now prompted to Restart Equalizing, for the Pool or Pool Member rule to be completely removed from the NetEqualizer Configuration. Once we *Click on -> [Restart Equalizing]*, our deletion will be permanent.

In our example, we have clicked on the red "x" next to the 4th row (Pool 4), to delete Pool 4 and its Pool Members. You can see on the screen below that Pool 1, 2, and 3 remain.

Note: You do not need to remove all members from a Bandwidth Pool to remove the Pool.

Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a Pool or Pool Member has been removed, see the following [View Traffic Limits](#) section of this User Guide.



Viewing Pools and Pool Members

There are several places in the NetEqualizer where you can view and validate your configuration for Pools and Pool Members. These reports are useful to confirm that your adds/changes/deletes have been saved and are now part of your configuration. We will discuss below.



View All Traffic Limits

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Manage Traffic Limits -> View All Traffic Limits*. Limits By Pool is the 2nd section of the report. You will see a list of your Pools, with all Pool Members shown directly below. In our excerpt, we show two Pools, corporate and wireless, and their pool members. From this report, you can also navigate back to the Configure Pools and VLANs batch entry, by clicking on the dark blue "Edit" button.

LIMITS BY POOL

Pool Number	Members	Pool Name	Download Mbps	Upload Mbps	Edit
1	1	corporate	18	18	
			IP Address	/	CIDR
			1 204.48.96.51	/	32
			2 204.48.111.0	/	24
2	2	wireless	12.04	12	
			IP Address	/	CIDR
			1 192.168.3.1	/	24



Configuration File

You can also see your Pools in the NetEqualizer Configuration File. From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [RTR] -> Configuration*. If your configuration file is more than 25 lines, you may need to click on the "Last" button to go to the end of your file, as we did below.

31	HARD	210.1.20.0 625000 500000 16 1
32	HARD	1.1.1.1/32 2250000 2250000 100001 1
33	HARD	204.48.96.51 32 0 200001 1
34	HARD	204.48.111.0 24 0 200001 1
35	HARD	2.2.2.2/32 1500000 1500000 100002 1
36	HARD	192.168.3.1 24 0 200002 1

Each Pool row will start with the Pool #, displayed as "#.#.#.#/32". For example, in the report above, you can see that row 32 is Pool 1 (1.1.1.1), and row 35 is Pool 2 (2.2.2.2). The row will be classified as "HARD", followed by the Pool #/32, and then the Download Limit, Upload Limit, and Pool # displayed as 100###, and Burst Factor, which always defaults to 1 for Pools. The Download and Upload limits are shown in bytes per second.

In our example above, you can read row 32 as follows: Pool 1, 2,250,000Bps (18Mbps) download limit, 2,250,000Bps (18Mbps) Upload Limit, and no bursting (1).

Pool Member rows are always shown directly below the Pool row. Pool Member rows contain IP address, CIDR, a non-used value defaulted to zero (0), 200Pool#, and Burst Factor, which



always defaults to 1. In our example above, Pool 1 has two Pool Member rows, row 33 and row 34. You can read row 33 as follows: an individual IP 204.48.96.61 (/32) is a member of Pool 1 (200001), with no bursting (1). Row 34 reads as: the subnet 204.48.111 and all its IP addresses (/24) are members of Pool 1 (200001), with no bursting (1).

*Note: Pools look similar to Hard Limits in the Configuration File. You can tell apart by the Pool IP address, which is made up of the Pool# repeating. For example, Pool 1 would be shown as HARD **1.1.1.1/32** 2250000 2250000 100001 1.*



Setting VLAN Limits

[\(back\)](#)

If you utilize VLANs on your network, you can set up your bandwidth limit rules to utilize your predefined VLANs. VLAN Limits work in the same fashion as Bandwidth Pools – they are *shared* bandwidth allocations.

The sum total of bandwidth for all of the IP addresses will not be allowed to exceed more than the total bandwidth allocated to the VLAN Limit. For example, if 200 IP addresses are assigned to a VLAN, and the VLAN Download Limit is 100Mbps, then the total downloadable bandwidth for all two hundred IPs together is 100Mbps (the total, not per IP). *VLAN Limits are a bandwidth restriction, not a reservation.*

In addition to enforcing the VLAN rate limits, the NetEqualizer will perform [equalizing](#) across all users on the VLAN when equalizing is ON. This works like [Bandwidth Pools](#), in that "[virtual equalizing](#)" is applied across all users on a VLAN.

In our example, shown below, if you set the download limit on VLAN #200 to 100Mbps, and the VLAN usage level reaches RATIO (default value is 85%), the NetEqualizer will begin to penalize any connection exceeding the value of [HOGMIN](#) within the VLAN.

From the NetEqualizer Dashboard or Navigation Menu [Click on ->\[Setup\] -> Manage Traffic Limits -> Configure Pools and VLANs ->By VLAN tab](#). The following screen opens.

VLAN ID	Download Mbps	Upload Mbps
1	100	100

Configure Pools and VLANs is a Batch Entry Screen, where you can add or edit many rules at once and then save your changes, saving you time in the setup process. Click on the dark blue "+" button at the top to add a VLAN Limit, or the red "x" next to each row to delete a VLAN Limit (both circled in blue above).

Once you have entered, modified, or deleted VLAN Limits, [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes.

In the light blue box above the entry fields, you will find notes to help you in creating VLAN Limits. Links are in orange, which you can click on to either get more information, or to move to another screen.

Once you save changes, you will be prompted to Restart Equalizing. Your new VLAN Limits will NOT take effect until you restart equalizing, even if you see them in your NetEqualizer Configuration.



We have also maintained your ability to add or delete VLAN Limits without restarting Equalizing. In Software Update 8.4, this maintained in our Quick Edit screens. From the Configure Pools and VLANs batch entry screen, you can move to the Quick Edit VLAN Limits screen (circled in orange on the screen above), to add or delete individual VLAN Limits without having to restart equalizing. Quick Edit cannot be used to modify rules.

VLAN Limits Rules

Before creating VLAN Limits, it is important for you to understand several rules that apply.

VLAN Limits cannot overlap with Hard Limits by IP

Once you add a VLAN Limit, you cannot also have a Hard Limit on an IP address in that VLAN range. An IP address either has a fixed amount of bandwidth (Hard Limit) or shared bandwidth (VLAN Limit), but not both. See [Configure Hard Limits by IP](#) to determine which option is right for your needs.

Priority Traffic and VLAN Limits

If you create a Priority Traffic rule for an IP address, and the IP address exists within a VLAN Limit, it will receive priority over other IP addresses within the VLAN Limit.

Creating a VLAN Limit

	VLAN ID	Download Mbps	Upload Mbps
1	200	100	100
2	205	50	25

A VLAN Limit is a shaping rule that causes the NetEqualizer to enforce your rate limit such that the aggregate bandwidth usage of all current VLAN users will not exceed the values selected for incoming and outgoing bytes per second.

To create a VLAN Limit, *Click on -> dark blue "+" button at the top of the screen* (circled in blue above) to add a new blank row. In our example above, a blank Row 2 was added.

Then *click on the VLAN id field*. Type in your VLAN id: *A Positive Integer between 1-4094*. We typed in "205". The window displays the current VLAN #s in use, so that you do not reuse a VLAN #. If your VLAN # is not in the range, you will see the following error when you click on Save Changes.

Oops! The following errors were found:
• limit 2: VLAN ID must be an integer in the range 1-4094 (-888)

TAB to Download Mbps field. Type in: *a Positive Number*. You can use fractional values, such as 50.5, or whole numbers. We typed in 50.

TAB to Upload Mbps field. Type in: *a Positive Number*. You can use fractional values, such as 50.5, or whole numbers. We typed in 25. TAB to complete your entry.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. For this reason, if you are adding a lot of VLAN Limits, you may want to save changes after adding each VLAN Limit. This will be fixed in a future release.



You can [Click on -> \[Save Changes\]](#) to save your VLAN Limit, or you can continue adding VLAN Limits. [Click on -> \[Reset\]](#) to discard changes.

In our example above, we now have a VLAN Limit for VLAN ID = 205, with a shared allocation of 50Mbps Download and 25Mbps Upload. VLAN 205 will be equalized at RATIO, to share bandwidth fairly among all IP address in the VLAN.

Modifying a VLAN Limit

As of Software Update 8.4, we provide you with the ability to modify your VLAN Limits, using the Configure Pools and VLANs screen. Additionally, you can make multiple edits at once, and then [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard.

From the NetEqualizer Dashboard or Navigation Menu [Click on -> \[Setup\] -> Manage Traffic Limits -> Configure Pools and VLANs -> By VLAN tab](#). The Configure Pools and VLANs screen opens.

In our example, shown below, we changed the Download Limit from 50 to 75Mbps, and the Upload Limit from 25 to 50Mbps for VLAN 205. These fields turn yellow, showing that we have updated the VLAN Limit but not saved our changes.

		VLAN ID	Download Mbps	Upload Mbps
	1	200	100	100
	2	205	75	50

Save Changes Reset

As we are done with our changes, we will now [Click on -> \[Save Changes\]](#) to save them. We are prompted to Restart Equalizing, for the updated Pool Limit rules to take effect. Once we [Click on -> \[Restart Equalizing\]](#), our new Pools settings will be available to the NetEqualizer process.

Deleting a VLAN Limit

In order to remove a VLAN Limit, click on the red "x" button on the row that you wish to delete. The VLAN Limit row will disappear. To completely remove the VLAN Limit from the configuration file, you need to [Click on -> \[Save Changes\]](#). We are now prompted to Restart Equalizing, for the VLAN Limit rule to be completely removed from the NetEqualizer Configuration. Once we [Click on -> \[Restart Equalizing\]](#), our deletion will be permanent.

In our example, we have clicked on the red "x" next to the 2nd row (VLAN id 205), to delete it. You can see on the screen below that only VLAN id 200 remains.

		VLAN ID	Download Mbps	Upload Mbps
	1	200	100	100

Save Changes Reset



Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a VLAN Limit has been removed, see the following [View Traffic Limits](#) section of this User Guide.

Viewing VLAN Limits

There are several places in the NetEqualizer where you can view and validate your configuration for VLAN Limits. These reports are useful to confirm that your adds/changes/deletes have been saved and are now part of your configuration. We will discuss below.



View All Traffic Limits

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Manage Traffic Limits -> View All Traffic Limits*. Limits By VLAN is the 3rd section of the report. Click on any of the dropdown arrows, like the one circled in blue below, to minimize or maximize any section of the report.

In our excerpt, we show two VLAN Limits, 200 and 205. From this report, you can navigate back to the Configure Pools and VLANs batch entry, by clicking on the dark blue "Edit" button.

LIMITS BY VLAN				
	VLAN ID	Download Mbps	Upload Mbps	Edit
1	200	100	100	
2	205	75	50	



Configuration File

You can also see your VLAN Limits in the NetEqualizer Configuration File. From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [RTR] -> Configuration*. If your configuration file is more than 25 lines, you may need to click on the "Last" button to go to the end of your file, as we did below.

38	VLAN	200 12500000 12500000
39	VLAN	205 9375000 6250000

Each VLAN Limit row will be classified as "VLAN", followed by the VLAN id, and then the Download Limit and Upload Limit. The Download and Upload limits are shown in bytes per second.

In our example above, you can read row 38 as follows: VLAN #200, 12,500,000Bps (100Mbps) download limit, and 12,500,000Bps (100Mbps) Upload Limit.



View All Traffic Limits



From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Manage Traffic Limits -> View All Traffic Limits*. This report will show all the Traffic Limits that you have set for the NetEqualizer, in the following order:

- 1) [Hard Limits](#)
- 2) [Limits by Pool](#)
- 3) [Limits by VLAN](#)
- 4) [Masked Hosts](#)
- 5) [User Quotas](#)
- 6) [P2P Traffic Limits](#)
- 7) [Priority Traffic](#)

We show an excerpt of the View All Traffic Limits Report here, containing two sections: 1) P2P Traffic Limits, and 2) Priority Traffic. We describe each section of the report in more detail under the appropriate section of the User Guide. If you want more detail, please refer to the section you are interested in by clicking on the link above in our 1-7 list.

From this report, you can also navigate back to any of the Traffic Limits' batch entry screens, by clicking on the dark blue "Edit" button in the section that you wish to edit. You can also minimize or maximize any section by clicking on the gray dropdown arrow in any section.

P2P TRAFFIC LIMITS					
	Host IP	/	CIDR	Connection Limit	Edit
1	192.168.1.0	/	24	60	
2	88.91.77.233	/	32	2	
PRIORITY TRAFFIC					
	Host IP	/	CIDR		Edit
1	192.168.1.1	/	32		
2	192.168.1.188	/	32		



Limit P2P Traffic

Peer-to-Peer (P2P) Traffic attempts to open 100's to 1000's of simultaneous connections on your network, and can absorb a significant portion of your available network bandwidth if left unchecked. As P2P traffic may be short, bursty-type traffic, it may not be stopped by Equalizing. Another mechanism is needed to control it adequately. Therefore, in addition to our fairness rules, NetEqualizer offers **Connection Limits** as a way to reduce peer-to-peer (P2P) traffic. Connection Limits may also lessen the effects of Distributed Denial of Service (DDoS) attacks. We will discuss both briefly here.

P2P traffic attempts to create hundreds, or possibly thousands, of simultaneous connections to absorb a lot of your network bandwidth. Setting Connection Limits effectively blocks or reduces P2P traffic, by not allowing connections over the limit that you specify.

Also, by their very nature, *Connection Limits are effective against both encrypted and unencrypted P2P traffic.* We believe this mechanism to be superior to managing policy files of known P2P traffic types, which will not help with encrypted P2P in any case.

In a DDoS attack, storms of incoming connections are generated by hackers, with the intention of overwhelming a server or servers. An attacker will spoof requests, sending storms of erroneously addressed connection requests to your server. These request storms create overwhelming administrative overhead, crippling the server and requiring a reboot by IT staff. While there are techniques that attempt to validate the incoming requests by sending queries back to the sending IP address for verification, these approaches create *more traffic* on the network. Instead of this approach, we chose to address the issue by offering DDoS protection via Connection Limits. You may also use our [DDoS Monitor](#) to assess potential DDoS threats.



Connection Limits Defined

We define a Connection as an inbound or outbound data stream (IP pair). Connection Limits (CL) enable you to define how many connections each user (IP address) can open simultaneously. Connection Limits can be set on all or part of your network, for specific IP addresses or entire subnets (range of IPs). Once a CL is set for an IP address or subnet, the NetEqualizer will keep a count of the total active connections (of any type) for that IP or each IP in the subnet. Once the total active connections equals your CL setting, additional connections are blocked.

Connection Limits control the number of inbound and outbound data streams (IP pairs or "connections") that each user on your network can create. Connection Limits are bi-directional; any limit you set is divided in two and applied. For example, a Connection Limit of sixty (60) would be turned into two connection limits: thirty (30) inbound and thirty (30) outbound connections.

Connection Limits can be set per individual IP or for an entire subnet at one time. If you set a Connection Limit for a subnet, all IP addresses in the subnet range will receive the specified Connection Limit, as *Connection Limits are not a shared limit.* For example, if you set a Connection Limit of 40 for a subnet with four different IP addresses, each IP address will get 20 inbound/20 outbound connections. Connection Limits can be set up for a Class B subnet, Class C subnet, or any legal subnet value 1-32.



Viewing Connection Counts



We recommend monitoring your installation for several days before setting Connection Limits, to better understand how many connections you need to support valid network activities. To monitor your connections, from the Main Dashboard or Navigation Menu, *Click on -> [RTR] -> Active Connections -> View Connection Counts.*

You will see IPs and their associated inbound, outbound, and total connections. As you view the Connection Counts Reports, if you see a lot of large connection counts (IPs with Total Connections greater than or equal to one hundred (≥ 100)), you may want to consider setting Connection Limits. You can view the Connection Count Report in the [View Connections Count](#) section of this User Guide.

Setting Connection Limits



From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Limit P2P Traffic.* The Limit P2P Traffic screen opens, as shown below. In our example, we have created two Connection Limit (CL) rules.

	Host IP	CIDR	Connection Limits
1	192.168.1.107	32	80
2	192.168.1.112	32	80

Limit P2P is a Batch Entry Screen, where you can add or edit many rules at once and then save your changes, saving you time in the setup process. Click on the blue "+" buttons to add entries, red "x" to delete entries (circled in blue below), and the blue "up/down arrows" (circled in orange) to reorder Connection Limits. Once you have entered, modified, or deleted Connection Limits, *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard changes.

In the light blue box above the entry fields, you will find notes to help you in creating Connection Limits. Links are in orange, which you can click on to either get more information, or to move to another screen.

Once you save changes, you will be prompted to Restart Equalizing. Your Connection Limits will NOT take effect until you restart equalizing, even if you see them in your NetEqualizer Configuration.

We have also maintained your ability to add or delete a rule without restarting



Equalizing. In Software Update 8.4, this maintained in our Quick Edit screens. From the Limit P2P Traffic batch entry screen, you can move to the [Quick Edit](#) P2P Traffic screen, to add or delete P2P Traffic Rules without having to restart equalizing. Quick Edit cannot be used to modify rules.

Connection Limit Rules

Before creating Connection Limits, it is important for you to understand several rules that apply. We will discuss two key rules below: 1) Connection Limit Order, and 2) Subnetted Connection Limits Cannot Overlap.

Connection Limit Order is Important

Order is important in setting up Connection Limits. If you set up Connection Limits for an entire subnet, and want to have a *different* Connection Limit apply to an IP address within that subnet, you would need to do the following:

- Set up Connection Limit for an individual IP address (/32)
- Set up Connection Limit for subnet (/16, /24, etc.)

For example, this is NOT allowed (as shown in the screen below) as the individual connection limit is AFTER the subnet limit.

CONNECTION 10.1.1.0/24 1600
CONNECTION 10.1.1.45/32 60

		Host IP	/	CIDR	Connection Limit
1		10.1.1.0	/	24	1600
2		10.1.1.45	/	32	60

Save Changes Reset

You can use the arrow buttons to re-order rows. In our case, we clicked on the dark blue arrow to move row 1 down. This resulted in a valid configuration, as the as individual connection limit is now BEFORE the subnet limit, as shown below.

CONNECTION 10.1.1.45/32 80
CONNECTION 10.1.1.0/24 1600

		Host IP	/	CIDR	Connection Limit
1		10.1.1.45	/	32	60
2		10.1.1.0	/	24	1600

Save Changes Reset

You might have an individual connection limit exception like this if you had an e-mail or DNS server within the subnet range, which would require additional connections during network operation. We recommend setting Connection Limits = 3,000 for email and DNS servers if they are within a connection-limited subnet range. Otherwise, you do not need to set connection limits for them.



Subnetted Connection Limits Cannot Overlap

Make sure when setting up your Connection Limits that you do not create an overlap between subnetted Connection Limits (/24, /16, etc.). The NetEqualizer does not support an overlap between subnetted Connection Limits.

For example, this is NOT allowed, as the subnet ranges overlap.

CONNECTION 10.1.1.2/24 60

CONNECTION 10.1.0.3/16 40

			Host IP	/	CIDR	Connection Limit
 	1	 	10.1.1.2	/	24	60
 	2	 	10.1.0.3	/	16	40

Save Changes

Reset

Creating a Connection Limit

To create a Connection Limit, *click on your computer's TAB to put the focus into the Host IP address field. Type in: Host IP address in 11.22.33.44 format.*

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B use /16, for a Class C use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will get the Connection Limit.

TAB to Connection Limit field. Type in: *a Positive Even Integer*. TAB off the field to complete your entry. Your Connection Limit value will be divided in two for inbound/outbound limits. Remember, CL applies to each IP or each IP in the subnet.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. This will be fixed in a future release.

Click on any of the "+" icons (circled in blue) to continue adding Connection Limits. Once you have entered all of your limits, *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard changes.

Limit P2P Traffic

Home / Setup / Limit P2P Traffic

Use this page to set connection limits for P2P traffic.

			Host IP:	/	CIDR	Connection Limit
 	1	 	192.168.1.107	/	32	80
 	2	 		/	32	



In our example above, we have created an individual CL for 192.168.1.107 (/32) of 80, which is 40 inbound/40 outbound connections.

Any unsaved fields are shown in yellow, as a visual reminder that you have not yet saved your changes. If you were to click off this screen before saving changes, your data would be cleared.

We then clicked on one of the dark blue “+” icons to continue adding Connection Limits. A second blank row appears on the screen, as shown above. You can click on the Host IP field to start adding your 2nd Connection Limit. Once we have added all CLs, we would **Click on -> [Save Changes]** to save our changes. We are now prompted to Restart Equalizing, for the new Connection Limits to take effect. Once we **Click on -> [Restart Equalizing]**, our new Connection Limits will be available to the NetEqualizer process.

Note: If you plan to set up a large number of subnet-ranged Connection Limits (≥ 32 subnet ranges), you will need to set several tuning parameters. Please see [Appendix #4](#) for detailed instructions.

Note: If you have online gamers on your network, you may need to set your Connection Limit as high as sixty (60) to facilitate online game playing.

Note: When you first set up a Connection Limit for an IP address, NetEqualizer will not drop existing connections over the limit. We wait until the overages die off or finish, but in the meantime the IP address cannot open any more connections.

Modifying a Connection Limit

As of Software Update 8.4, we provide you with the ability to modify your Connection Limits, using the Limit P2P Traffic screen. Additionally, you can make multiple edits at once, and then **Click on -> [Save Changes]** to save changes or **Click on -> [Reset]** to discard. In our example below, we first changed the Connection Limit from 60 to 100. The Connection Limit field turns yellow, showing that we have updated the field but not saved it.

			Host IP	/	CIDR	Connection Limit
	1		192.168.1.107	/	32	80
	2		192.168.1.112	/	32	100

Save Changes Reset

We then reordered the Connection Limits, by clicking on the dark blue “up arrow” button, to move the 192.168.1.112 to the 1st row, as shown below.

			Host IP	/	CIDR	Connection Limit
	1		192.168.1.112	/	32	100
	2		192.168.1.107	/	32	80

Save Changes Reset



As we are done with our changes, we will now [Click on -> \[Save Changes\]](#) to save them. As shown below, our CL rows show no remaining yellow fields, as our changes are saved. We are now prompted to Restart Equalizing, for the updated Connection Limit rules to take effect. Once we [Click on -> \[Restart Equalizing\]](#), our new CL settings will be available to the NetEqualizer process.

✔ Success! Connection limits have been updated. Restart Equalizing to complete connection limit update.
✕

▶ Restart Equalizing

			Host IP	/	CIDR	Connection Limit
	1		192.168.1.112	/	32	100
	2		192.168.1.107	/	32	80

Save Changes
Reset

Deleting a Connection Limit

In order to remove a Connection Limit, click on the red "x" button on the row that you wish to delete. The CL row will disappear. To completely remove the CL from the configuration file, you need to [Click on -> \[Save Changes\]](#). We are now prompted to Restart Equalizing, for the Connection Limit rule to be completely removed from the NetEqualizer Configuration. Once we [Click on -> \[Restart Equalizing\]](#), our CL will be permanently deleted.

In our example, we have clicked on the red "x" next to the 2nd row, to delete the CL for 192.168.1.107 /32 of 80. You can see on the screen below that only the CL for 192.168.1.112 remains.

ℹ Use this page to set connection limits for P2P traffic.
^

			Host IP	/	CIDR	Connection Limit
	1		192.168.1.112	/	32	100

Save Changes
Reset

Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a Connection Limit has been removed, see the following [View Connection Limits](#) section of this User Guide.



Viewing your Connection Limits

There are two places in the NetEqualizer where you can view your P2P Traffic Limits. We will discuss both here.



View All Traffic Limits

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Manage Traffic Limits -> View All Traffic Limits*. Scroll down to the P2P Traffic Limits section of the report. You will see a list of your Connection Limits, as shown in the following excerpt from the report:

P2P TRAFFIC LIMITS					Edit
	Host IP	/	CIDR	Connection Limit	
1	192.168.1.112	/	32	100	
2	192.168.1.107	/	32	80	

From this report, you can also navigate back to the P2P Traffic Limits batch entry, by clicking on the dark blue "Edit" button.



Configuration File

You can also see your Connection Limits in the NetEqualizer Configuration File. From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [RTR] -> Configuration*. If your configuration file is more than 25 lines, you may need to click on the "Last" button to go to the end of your file, as we did in our example below.

36	CONNECTION	192.168.1.112/32 50 0
37	CONNECTION	192.168.1.107/32 40 0
Line #	Parameter	Value
Showing 26 to 37 of 37 entries (processed in 0.00040006637573242 s)		
First Previous 1 2 Next Last		

There will be one row for each Connection Limit, encompassing both an inbound and outbound limit, listing half the value that you selected (i.e. for value=100, you would see 50, as in the 1st row above). The rows will start with "CONNECTION" and also show the IP address or subnet that is being connection limited. In our example, we have two individual IP (/32) Connection Limits.

Most users typically peak out at 20 to 30 connections per second each for INBOUND and OUTBOUND traffic, so a Connection Limit of 60 would suffice in most cases. Setting a **Connection Limit = 60** is a good recommendation and excellent at controlling most P2P traffic.



Consider Setting Bandwidth Priority

NetEqualizer's default equalizing rules, which are set to "on" by default, are able to handle congestion-related traffic flow problems for most organizations. Most types of traffic that organizations want to prioritize are *prioritized by default* just by using the Equalizing Rules.

However, some organizations need to setup **Bandwidth Priority Rules** for specific traffic types. There are two types of priority that you can set up on the NetEqualizer: 1) traffic with priority over equalizing, and traffic hidden from equalizing.

We will discuss these in detail below:

1. [Priority Traffic](#)
 - Predefined IPs or subnets with priority over equalizing.
2. [Masking Off Traffic](#)
 - Traffic hidden from equalizing.

Priority Traffic gives known IP addresses and their associated streams preferential treatment. Bandwidth Priority Rules are most often used for streaming-video traffic. For example, if a business is streaming training videos into corporate offices, a "Priority Traffic" Rule would need to be set up to prioritize the IP address of the server or site hosting the training videos.

Masked Traffic is "invisible" to the NetEqualizer. Typically this is used to *exclude local traffic* (i.e. a computer talking to a server on your network) crossing the NetEqualizer link from being considered for any shaping decisions.



Defining Priority Traffic

[\(back\)](#)

How does NetEqualizer grant priority for IP addresses? NetEqualizer recognizes two classes of traffic:

1. Priority Traffic
2. Data Traffic


When **Priority Traffic** is detected, the bandwidth allocation for rest of the Data Traffic is reduced. When NetEqualizer identifies a priority IP address, it typically performs the following process:

1. A priority IP address becomes active
2. NetEqualizer dynamically reduces the data congestion ratio ([RATIO parameter](#)) by a few percent
3. This action (b) forces the PENALTY mechanism to kick in a bit sooner for non-priority streams, thus reserving space for your priority traffic
4. Priority traffic is given immunity to flow control. These streams will not be slowed by PENALTIES applied in [Equalizing](#). However, any Bandwidth Limiting Rules, such as [Hard Limits](#), will remain in effect.

Priority Traffic is assured bandwidth, up to the size of your network pipe. When you set up Priority Traffic, all your other traffic is dynamically pushed into a smaller bandwidth window. Note that if you set too much priority traffic, you will push all your remaining traffic into a very small window.

Factory delivered, NetEqualizer defaults are set to perform congestion control on your trunk when it becomes 85 percent full. In most cases, important business applications, such as VoIP, citrix, blackboards, web browsing, and e-mail will receive preferential treatment, and therefore there is no need to assign priority. In general, we find that only video servers require priority treatment.

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Manage Priority Traffic*. The Manage Priority Traffic screen opens, as shown below. In our example, we have created one Priority Traffic rule for individual IP 192.168.1.188 (/32).



	Host IP	/	CIDR	
	1	192.168.1.188	/	32

Priority Traffic allows you to select a specific IP address for priority treatment. Once set, this IP address, and any connection it is part of, will receive priority.

In [Software Update 5.8](#), we expanded our Priority Traffic feature to enable you to prioritize an entire IP subnet. This feature is useful if you need to prioritize a section of your network, for example, a subnet where your video-streaming servers are hosted.

If you use this feature to prioritize an IP or entire subnet and are concerned with these priorities using too much bandwidth, we also recommend using [Hard Limits](#) to add a hard limit for the IP or subnet, so that it does not take an unlimited amount of bandwidth.

Note: Use Priority Traffic sparingly. The most common mistake for new installations is to try to give priority to all important business applications. This is rarely actually needed, as most business applications will already be getting preferential treatment from [Equalizing](#).

Creating a Priority Limit

To create a Priority Limit, *click on your computer's TAB to put the focus into the Host IP address field*. Type in: *Host IP address* in 11.22.33.44 format.

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B use /16, for a Class C use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will get the Connection Limit.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. This will be fixed in a future release.



Click on any of the "+" icons (circled in blue) to continue adding Priority Limits. Once you have entered all of your limits, [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes. We are now prompted to Restart Equalizing, for the new Priority Limits to take effect. Once we [Click on -> \[Restart Equalizing\]](#), our new Priority Limits will be available to the NetEqualizer process.

Modifying a Priority Limit

As of Software Update 8.4, we provide you with the ability to modify your Priority Limits, using the Manage Priority Traffic screen. Additionally, you can make multiple edits at once, and then [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard. In our example below, we added a 2nd Priority Limit. The Host IP field turns yellow, showing that we have updated a field but not saved it.

		Host IP	/	CIDR
	1	192.168.1.188	/	32
	2	192.168.1.200	/	32

[Save Changes](#) [Reset](#)

As we are done with our changes, we will now [Click on -> \[Save Changes\]](#) to save them. Once saved, our Priority Limit rows show no remaining yellow fields, as our changes are saved. We are now prompted to Restart Equalizing, for the updated Priority Limit rules to take effect. Once we [Click on -> \[Restart Equalizing\]](#), our new Priority Limit settings will be available to the NetEqualizer process.

Deleting a Priority Limit

In order to remove a Priority Limit, click on the red "x" button on the row that you wish to delete. The row will disappear. To completely remove the Priority Limit from the configuration file, you need to [Click on -> \[Save Changes\]](#). We are now prompted to Restart Equalizing, for the Priority Limit rule to be completely removed from the NetEqualizer Configuration. Once we [Click on -> \[Restart Equalizing\]](#), our Priority Limit will be permanently deleted.

In our example, we have clicked on the red "x" next to the 2nd row, to delete the Priority Limit for 192.168.1.200. You can see on the screen below that only the Priority Limit for 192.168.1.188 remains.

		Host IP	/	CIDR
	1	192.168.1.188	/	32

[Save Changes](#) [Reset](#)

Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a Priority Limit has been removed, see the following [View Priority Limits](#) section of this User Guide.



Masking Off Traffic



[\(back\)](#)

The masking features on NetEqualizer are intended to exclude Local Traffic crossing the NetEqualizer link from being considered for any shaping decisions. Masked traffic is "invisible" to the NetEqualizer. If you are utilizing the NetEqualizer to shape Internet Traffic going across your link, you should use the MASK feature to exclude Local Traffic (i.e. a computer talking to a server on your network).

Masking should not be used to prioritize traffic. [Priority Traffic](#) should be used to prioritize traffic, such as important video streams. Do not use the MASK feature.

There are two types of masking, "paired" and "absolute." A host or subnet assigned as a "paired" mask will only be ignored if it is talking to another host or subnet that is also registered as a paired mask. By design, a **Paired Mask** will cause NetEqualizer to ignore hosts talking to other paired mask hosts, while at the same time subject the same hosts to NetEqualizer's bandwidth shaping rules if they make a connection with a server on the Internet. **Absolute Masks** ignore all traffic to or from the masked host or subnet regardless of the connection.

From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[Setup\] -> Manage Traffic Limits -> Configure Masked Hosts](#). The Configure Masked Hosts screen opens, as shown below. In our example, we have created three Masked Hosts rules. Masks can be set for an individual IP address, an entire subnet, or any legal subnet value 1-32.

Note: In most cases, you will not need to use masking. NetEqualizer is typically setup on your Internet link, and does not see Local Traffic.

Note: If you plan to set up a large number of subnet-ranged Masks (≥ 32 subnet ranges), you will need to set several tuning parameters. Please see [Appendix #4](#) for detailed instructions.

Configure Masked Hosts

Home / Setup / Manage Traffic Limits / [Configure Masked Hosts](#)

Use this page to configure traffic that will be hidden from equalizing.

- Use [Quick Edit Masked Hosts](#) to add and delete masked hosts without having to restart equalizing.
- IP addresses are described using [CIDR notation](#) -- a CIDR of 32 means "this IP address only".
- Changes to masked hosts will require stopping and restarting the equalizing process to take effect.
- No changes are made to the masked hosts until they are explicitly saved.

		Host IP	/	CIDR	Type	
	1	12.155.72.147	/	32	☒ Absolute	☐ Paired
	2	33.52.185.213	/	24	☒ Absolute	☐ Paired
	3	192.168.2.0	/	24	☐ Absolute	☒ Paired



Creating a Masked Host

To create a Masked Host, *click on your computer's TAB to put the focus into the Host IP address field*. Type in: *Host IP address* in 11.22.33.44 format.

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B use /16, for a Class C use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will get the Connection Limit.

Click on -> "Absolute" or "Paired" radio button to define the Mask Type. Used Paired to only ignore traffic between paired mask hosts. Use Absolute to ignore all traffic to or from the masked host.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. This will be fixed in a future release.

Click on any of the "+" icons (circled in blue) to continue adding Masked Hosts. Once you have entered all of your masks, *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard changes.

VLAN Masking

In [Software Update 5.8](#), we expanded our masking feature to enable you to create VLAN Masks, using VLAN IDs. Our new VLAN Masking feature enables you to designate *entire local VLANs* that you want masked from Equalizing. To implement this feature, follow the instructions below.



From the Maintenance and Reference Menu, *Click on -> Maintenance -> [Run a Command]*.

To create a VLAN Mask (also known as VLAN Exclusion):

Type in: */sbin/brctl vlanexclusion my #*

This will exclude this VLAN ID (#) and store this VLAN # in the VLAN Exclusion Table. To control what the VLAN Exclusion Table does:

Type in: */sbin/brctl vlanflag my [0,1,2,3]*

[0,1,2,3] Values

0 = Turn off the feature (VLAN Masking is off).

1 = Mask all VLAN IDs in the table (*the specified VLANs are NOT equalized*).

2 = Mask all VLAN IDs EXCEPT what is in the table (*the specified VLANs ARE equalized*).

3 = Clear out the table and turn off the feature.

Note: Three (3) is different than 0 because zero (0) just turns off the feature without clearing the table. If you used zero (0) and then set it to 1 again, the members would stay the same.

Modifying a Masked Host

As of Software Update 8.4, we provide you with the ability to modify your Masked Hosts, using the Configure Masked Hosts screen. Additionally, you can make multiple edits at once, and then *Click on -> [Save Changes]* to save changes or *Click on -> [Reset]* to discard. In our example below, we updated the CIDR on the second Masked Host from /24 to a /22. The CIDR field turns yellow, showing that we have updated a field but not saved it.



		Host IP	/	CIDR	Type	
	1	12.155.72.147	/	32	☒ Absolute	☐ Paired
	2	33.52.185.213	/	22	☒ Absolute	☐ Paired
	3	192.168.2.0	/	24	☐ Absolute	☒ Paired

As we are done with our changes, we will now [Click on -> \[Save Changes\]](#) to save them. Once saved, our Masked Host rows show no remaining yellow fields, as our changes are saved. We are now prompted to Restart Equalizing, for the updated Masked Host rules to take effect. Once we [Click on -> \[Restart Equalizing\]](#), our new Masked Host settings will be available to the NetEqualizer process.

Deleting a Masked Host

In order to remove a Masked Host, click on the red "x" button on the row that you wish to delete. The row will disappear. To completely remove the Masked Host from the configuration file, you need to [Click on -> \[Save Changes\]](#). We are now prompted to Restart Equalizing, for the Masked Host rule to be completely removed from the NetEqualizer Configuration. Once we [Click on -> \[Restart Equalizing\]](#), our Masked Host will be permanently deleted.

In our example, we have clicked on the red "x" next to the 2nd row, to delete the Masked Host for 33.52.185.213 /22. You can see on the screen below that two Masked Host limits remain.

		Host IP	/	CIDR	Type	
	1	12.155.72.147	/	32	☒ Absolute	☐ Paired
	2	192.168.2.0	/	24	☐ Absolute	☒ Paired

Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a Masked has been removed, see the following [View Priority Limits](#) section of this User Guide.



Viewing your Priority Limits and Masked Traffic

There are several places in the NetEqualizer where you can view and validate your configuration for Priority Limits and Masked Traffic. These reports are useful to confirm that your adds/changes/deletes have been saved and are now part of your configuration.



View All Traffic Limits

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Manage Traffic Limits -> View All Traffic Limits*. Scroll down to the Priority Traffic Limits section of the report. You will see a list of your Priority Limits, as shown in the following excerpt from the report:

PRIORITY TRAFFIC				
	Host IP	/	CIDR	Edit
1	192.168.1.188	/	32	

Now scroll to the Masked Hosts section of the report, to see any Masked Hosts that you have configured.

MASKED HOSTS					
	Host IP	/	CIDR	Type	Edit
1	12.155.72.147	/	32	Absolute	
2	33.52.185.213	/	24	Absolute	
3	192.168.2.0	/	24	Paired	

From this report, you can also navigate back to the Configure Priority Traffic or Configured Masked Hosts batch entry, by clicking on the dark blue "Edit" button.



Configuration File

You can also see your Priority Traffic or Masked Hosts in the NetEqualizer Configuration File. From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [RTR] -> Configuration*. If your configuration file is more than 25 lines, you may need to click on the "Last" button to go to the end of your file.

This excerpt shows the three Masked Hosts that we displayed in the View All Traffic Report above. The rows will start with "MASK", followed by IP address or subnet, followed by the CIDR, here /32 and /24, and then "1001" for absolute limits and "1002" for paired limits.

40	MASK	12.155.72.147/32 1001
41	MASK	33.52.185.213/24 1001
42	MASK	192.168.2.0/24 1002

The following excerpt shows the Priority Limit that we displayed above in the View all Traffic Report. The rows will start with "PRIORITY", followed by the IP address and CIDR, here a /32 for an individual IP, then followed by "999999".

40	PRIORITY	192.168.1.188/32 999999
----	----------	-------------------------

Restricting Bandwidth Usage

Restricting Bandwidth Usage features encompass defining how much bandwidth to give a user over a specified time period (Setting User Quotas) and how to handle unauthorized access attempts (MAC Redirection).

1. [User Quotas](#) - GUI interface to define bandwidth usage limits over a time period. Also known as Professional Quota API.
2. [MAC Redirection](#) - Define authorized MACs on your network.



Establishing User Quotas

As of Software Update 8.4, Manage User Quotas replaces the Professional Quota API GUI Interface. Manage User Quotas will help you to quickly and easily utilize our NetEqualizer User-Quota API toolset commands. Quotas may be of interest to you if you run a business that charges customers based on bandwidth usage, or if you want to track bandwidth usage over a set time period. Manage User Quotas is a standard pre-written quota utility imbedded in each system. You can quickly plug in IP addresses from the GUI and have a monthly quota enforced right away. The GUI Interface enables you to:

- Track user data by IP
- Specify Quotas and Bandwidth Limits Rules by IP or an entire IP subnet (quotas are then applied to each IP individually within the subnet).
- Monitor real-time bandwidth utilization at any time
- Set up notification alarms when users exceed their bandwidth limits

Note: If you are looking to restrict bandwidth use to a set amount, and do not need to track bandwidth usage over time, we recommend setting a Bandwidth Limit. See the [Setting Bandwidth Limits](#) section of this User Guide to determine if Hard Limits, Pools, or VLAN Limits meet your needs.

From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[Setup\] -> Manage Traffic Limits -> Manage User Quotas](#). We then clicked on [Configure User Quotas Tab](#). The Configure User Quotas Tab opens. In our example, we have created two Quota rules.

START/STOP QUOTA SYSTEM

CONFIGURE USER QUOTAS

RESET QUOTA RULE

Use this page to configure user quotas.

- Use [Quick Edit User Quotas](#) to add and delete user quotas without having to restart equalizing.
- IP addresses are described using [CIDR notation](#) -- a CIDR of 32 means "this IP address only".
- Quota Amount -- number of total bytes for this quota to allow.
- Duration -- number of minutes between quota reset.
- Hard Limit Restriction -- number of Mbps to allow when exceeded.
- Contact -- who to contact when the alarm goes off. Use [Manage Alerts](#) to setup email for quota alerts.
- Changes to user quotas will require stopping and restarting the quota system to take effect.
- No changes are made to the user quotas until they are explicitly saved.

		Host IP	/	CIDR	Quota Amount (Bytes)	Duration	Hard Limit Restriction Mbps	Contact
☒	1	192.168.1.244	/	32	1000000000	1440	0.5	admin@example.cc
☒	2	192.168.1.245	/	32	500000000	1440	1	admin@example.cc

Save Changes

Reset



Manage User Quotas is a Batch Entry Screen, where you can add or edit many rules at once and then save your changes, saving you time in the setup process. Click on the dark blue "+" button to add entries, or the red "x" to delete entries. Once you have entered, modified, or deleted User Quota Rules, [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes.

In the light blue box above the entry fields, you will find notes to help you in creating User Quota Rules. Links are in orange, which you can click on to either get more information, or to move to another screen.

Once you save changes, you will be prompted to Restart the Quota System. Your User Quota Rules will NOT take effect until you Restart Quota, even if you see them in your NetEqualizer Configuration.

We have also maintained your ability to add or delete a rule without restarting the Quota System. In Software Update 8.4, this maintained in our Quick Edit screens. From the Manage User Quotas batch entry screen, you can move to the [Quick Edit](#) User Quota screen, to add or delete User Quota Rules without having to restart Quota. Quick Edit cannot be used to modify rules.

User Quota Rules

Before creating User Quota Rules, it is important for you to understand several rules.

Quota Hard Limit Restrictions are only enforced on Downloads

Quota Hard Limit Restrictions are only enforced on downloads. Uploads are not restricted. This is a simplification that we put in place as 99 percent of problems are download-related.

Hard Limits Interact with Quota

If you wish to have BOTH Hard Limits and Quota Rules in place, *the Hard Limits must be specified at the subnet level* (range of IPs). At this time you cannot set a Hard Limit for an individual IP (/32) with an individual or subnetted Quota Rule.

Quota and Hard Limits: Definition Rules

Allowed?	NO	NO	YES	YES
Hard Limit	IP	IP	subnet	subnet
Quota	subnet	IP	IP	subnet
For Example: <i>111 means not allowed</i>				
Hard Limit	10.99.100.02/32	10.99.100.02/32	10.99.100.01/24	10.99.100.01/24
Quota	10.99.100.01/24	10.99.100.02/32	10.99.100.02/32	10.99.100.01/24

How Hard Limits Interact with Quota

Hard Limits are used to set a fixed RATE of individual bandwidth to a single IP address or a subnet. If you set a Hard Limit for a subnet, the bandwidth assigned is not shared. For example, if you set up a 2Mbps up/1Mbps down for a subnet with four different IP addresses, each of the IP addresses will get 2Mbps/1Mbps to use.

Quotas are used to limit the AMOUNT of data that a single IP address or IPs within a subnet can use over a period of time.

Here is how they work together: If you have both a Hard Limit and a Quota Rule set for a IP address, if the Quota Amount is reached, the Quota System Hard Limit Restriction will kick in, replacing the Hard Limit set for the IP address. This can be used to ramp down your



Hard Limit when too much bandwidth is being consumed.

For example, if you had a Hard Limit for 10.99.100.01/24 of 5Mbps, each IP in the subnet would be limited to this rate of bandwidth consumption. If an IP in the subnet exceeded the 1 gigabyte Quota Amount during the week, the IP would now be limited to 1Mbps, instead of the 5Mbps. Other IPs in the subnet that had not exceeded quota would still get 5Mbps.

Set Date, Time, and Time Zone

Before starting the Quota System, make sure to set the NetEqualizer Time Zone, and also the Date and Time, so that the rules that you add to Quota System are aligned with your date, time, and time zone. For instructions, see the Set Date/Time and Time Zone section of the [Quick Start Guide](#).

We recommend that you limit the # of quota email addresses

To reduce the number of emails, use the same email address as your Contact for each Quota Rule. If a quota violation occurs, we will send one (1) email per hour to that address containing all violations. You will see one email per email address used as your Contact.

Creating User Quota Rules

	Host IP	/	CIDR	Quota Amount (Bytes)	Duration	Hard Limit Restriction Mbps	Contact
	1	10.99.100.1	/ 24	1000000000	10080	1	support@gmail.com
Save Changes Reset							

At the heart of the Professional Quota API are User Quota Rules. Above we have created one User Quota Rule. Our example would restrict each IP within the subnet to use 1 gigabyte of data over a one (1) week period, and then cap each IP at 1Mbps if they exceed their quota. You can read our settings as follows:

Parameter	Value
Host IP	10.99.100.1
CIDR	24 (subnet)
Quota Amount	1000000000 (1,000,000,000 bytes = 1 gigabyte)
Duration	10080 (1 week=7 days * 24 hours * 60 minutes)
Hard Limit Restriction	1 (Mbps)
Contact	support@gmail.com

To create a User Quota, *click on your computer's TAB to put the focus into the Host IP address field*. Type in: *Host IP address* in 11.22.33.44 format.

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B use /16, for a Class C use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will get the Connection Limit.

TAB to Quota Amount field. Enter an amount in bytes per second. This is the total amount of Download bytes allowed before the QUOTA restriction kicks in. See [Quota Rules Parameters Table](#) for more details.



TAB to Duration. Enter an amount in minutes. Quota Amount applies for this amount of time (duration). See [Quota Rules Parameters Table](#) for more details.

TAB to Hard Limit Restriction. Enter an amount, which will be displayed in the configuration units that you selected in [Preferences](#). See [Quota Rules Parameters Table](#) for more details.

TAB to Contact. Enter an email address, which must exist in the Configure Alert Emails. To reduce the number of emails, use the same email address as your Contact for each Quota Rule. If a quota violation occurs, we will send one (1) email per hour to that address containing all violations. If you had 100 Quota Rules set up, and 25 violations occurred, you would see one (1) email containing all 25 violations. For the above example, support@gmail.com would only get one (1) email each hour, even if multiple IPs in the range had quota violations.

If you have not already set up Alert Emails, please use [Manage NetEqualizer-> Manage Alerts -> Configure Alert Email](#) to set up your alert emails. See [Quota Rules Parameters Table](#) for more details.

Note: In 8.4, you must use "tab" to move between rows. If you hit "return", your data will be cleared. This will be fixed in a future release.

Click on any of the "+" icon to continue adding User Quotas. Once you have entered all of your User Quotas, [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard changes.

Once you Save Changes, you will be prompted to Restart Quota, as shown in the message here. [Click on -> \[Restart Quota\]](#) to have your new User Quota setting be available to the Quota process.

✔ **Success!** User quotas have been updated. Restart the quota system to complete the user quota update.



Restart Quota System

Resetting User Quota Rules

Resetting Quota on an IP has the effect of flushing out the data counted against the Quota Amount, setting it back to zero (0), and changing the Start Time for the Quota Rule to NOW(). In our example above, if we reset our Quota Rule for the **10.99.100.1/24** subnet, it would be back at 0 bytes against its one (1) gigabyte quota, and the quota would restart at the current date & time. From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[Setup\] -> Manage Traffic Limits -> Manage User Quotas](#). We then clicked on the [Reset Quota Rule Tab](#). The Reset Quota Rule Tab opens, as shown below.

START/STOP QUOTA SYSTEM	CONFIGURE USER QUOTAS	RESET QUOTA RULE
🔔 Use this page to reset the quota rule for an IP address. IP Address to Reset: Reset Quota Rule		



Tab to the IP Address to Reset field, and type in the IP address for the rule you would like flushed out, for our example that would be **10.99.100.1**, and *Click on ->[Reset Quota Rule]*.

If the Quota System is not started, you cannot reset rules, and will see the following warning message. *Click on ->[Start Quota System]* to start the Quota System. Once quota is started, you will be able to Reset Quota.

Warning! Quota rules cannot be reset while the quota system is stopped. Start the quota system if you want to reset quota rules.

▶ Start Quota System

Quota Rules Parameters Table

Parameter	Unit	Definition	What you can set to...	Tips
HOST IP	IP Address	Host IP is an IP address or subnet range for which the rule applies.	x.x.x.x (10.0.0.28) /32 = one IP address /24, /16 etcetera to apply to a subnet.	If you define for an entire subnet, the Quota Amount is a BYTE LIMIT FOR EACH IP (not shared across IPs).
Quota Amount	Bytes	The total amount of Download bytes allowed before the QUOTA restriction kicks in.	Any value in bytes. For 1 GB, you would enter 1000000000. (1,000,000,000 bytes)	Do not put commas in the rule.
Duration	Minutes	Quota Amount applies for this amount of time (duration). If the QUOTA amount goes over during that time the restriction will be enforced for the rest of that time period. Once the time expires, the quota restriction is reset, and the time starts over.	One (1) day would be 1440 (24*60). If Quota Amount was set to 2000000 (2meg), and Duration was 1440, the IP address would be restricted to 2meg over 1 Day.	Industry best practice for ISPs is to set Duration = 1 week. To span 1 Week= 10080 7 days * 24 hours * 60 minutes To span Days: # days * 24 hours * 60 minutes To span Hours: # hours * 60 minutes
Hard Limit Restriction	Bytes per second	The amount the IP or subnet will be set to if it exceeds its limit (if it goes over Quota Amount in Duration).	Any value in bytes per second.	Enforced on downloads only.
Contact	Google Email address (gmail)	Who to send an email to if an alarm is raised.	Any valid gmail address that is also set up as an Alert Email. Only 1 violation per hour is emailed out per IP address. IP addresses that share a Contact Email will go out in one notification.	gmail address must also be set-up in Manage NetEqualizer-> Manage Alerts -> Configure Alert Email window for this to work.



Modifying User Quota Rules

As of Software Update 8.4, we provide you with the ability to modify your User Quota Rules, using the Configure User Quotas Tab. Additionally, you can make multiple edits at once, and then [Click on -> \[Save Changes\]](#) to save changes or [Click on -> \[Reset\]](#) to discard. From the NetEqualizer Dashboard or Navigation Menu [Click on -> \[Setup\] -> Manage Traffic Limits -> Manage User Quotas -> Configure User Quotas Tab](#). The screen below opens.

In our example, we changed our 2nd User Quota rule IP address from 192.168.1.245 to 192.168.200.250, and also changed the CIDR from an individual IP (/32) to a subnet (/24). These fields turn yellow, showing that we have updated the Hard Limit but not saved it.

		Host IP	/	CIDR	Quota Amount (Bytes)	Duration	Hard Limit Restriction Mbps	Contact
	1	192.168.1.244	/	32	1000000000	1440	0.5	admin@example.cc
	2	192.168.200.250	/	24	500000000	1440	1	admin@example.cc

Save Changes Reset

As we are done with our changes, we will [Click on -> \[Save Changes\]](#) to save. We are prompted to Restart the Quota System, for our updated User Quota Rules to take effect. [Click on -> \[Restart Quota System\]](#) for the settings to be available to the Quota process.

Deleting User Quota Rules

In order to remove a User Quota Rule, click on the red "x" button on the row that you wish to delete. The User Quota row will disappear.

In our example, we have clicked on the red "x" next to the 2nd row, to delete the User Quota Rule for 192.168.200.250/24 that we edited above. You can see on the screen below that only the User Quota Rule for 192.168.1.244 remains.

		Host IP	/	CIDR	Quota Amount (Bytes)	Duration	Hard Limit Restriction Mbps	Contact
	1	192.168.1.244	/	32	1000000000	1440	0.5	admin@example.cc

Save Changes Reset

To completely remove the User Quota Rule from the configuration file, you need to [Click on -> \[Save Changes\]](#). We are now prompted to Restart the Quota System, for the User Quota rule to be completely removed from the NetEqualizer Configuration. Once we [Click on -> \[Restart Quota System\]](#), our User Quota Rule will be permanently deleted.


Note: In Software Update 8.4, we do not mark rows for deletion, so it is important to Save Changes after you make deletions. To check your configuration to make sure a User Quota Rule has been removed, see the following [View Traffic Limits](#) section of this User Guide.




Starting the Quota System

In order for the quota commands to work, you must first start the Quota System. If not started, you will see error or warning messages when you try to Reset Quota. You will also be prompted to Start (or Restart) Quota after saving any User Quota Rules. You can use the NetEqualizer Dashboard to see if the Quota System is running.

If you are not already on the Dashboard, from the Navigation Menu, [Click on -> \[Home\]](#). The Status Indicator for the Quota System will show either ON (green), as shown at right, or OFF (red).



If the Quota System is OFF, you will need to start it. To open the Manage User Quotas screen, you can either click on the Quota System Status Indicator, or [Click on -> \[Setup\] -> Manage Traffic Limits -> Manage User Quotas](#). The following screen opens.



[Click on -> \[Start Quota System\]](#) You see the following message once the Quota System is started.

✓ **Success!** Quota system has been started.

Setting up Quota Email Notifications

In order to get email notifications, you must set up a **valid gmail account** to be used to send emails from the Quota System.

In order to set up an email for notifications, [Click on -> Manage NetEqualizer -> Manage Alerts -> Configure Alert Email](#). See the [Email Notifications](#) section of this User Guide for details on how to set up an alert email.

Stopping the Quota System

When you Stop Quota, you are clearing out the byte counters, as they are stored in RAM at this time. Therefore, ALL your "Total Bytes Down" counters will be zero (0) once you Stop Quota.


The Quota System is designed to stay up for months at a time. In the future, we may incorporate a disk drive backup capability so in the event of a power failure you could resume quota counts. However, the current class of NetEqualizer systems do not have disk drives (for a variety of good reasons), and so we have not incorporated this into our Quota System.



Note: While the quota data does not persist, the Quota Rules are persistent. They will come up operational once defined, each time you start the Quota System.

If you still think you want to stop collecting quota data, use this menu item.

If you are not already on the Dashboard, from the Navigation Menu, **Click on -> [Home]**. The Status Indicator for the Quota System will show either ON (green), as shown at right, or OFF (red). If the Status Indicator is OFF (red), the Quota System is already off; stop here.



To open the Manage User Quotas screen, you can either click on the Quota System Status Indicator, or **Click on -> [Setup] -> Manage Traffic Limits -> Manage User Quotas**. The following screen opens.



Click on -> [Stop Quota System] You see the following message once the Quota System is stopped.

✓ **Success!** Quota system has been stopped.

Viewing User Quotas

There are several places in the NetEqualizer where you can view and validate your configuration for User Quotas. These reports are useful to confirm that your adds/changes/deletes have been saved and are now part of your configuration.



View All Traffic Limits

From the NetEqualizer Dashboard or Navigation Menu, **Click on -> [Setup] -> Manage Traffic Limits -> View All Traffic Limits**. Scroll down to the 5th section of the report, which shows User Quotas. You will see a list of your User Quotas, as shown in the following excerpt from the report:

USER QUOTAS								
	Host IP	/	CIDR	Quota Amount	Duration	Hard Limit Restriction Mbps	Contact	Edit
1	192.168.1.244	/	32	1000000000	1440	0.5	admin@example.com	
2	192.168.1.245	/	32	500000000	1440	1	admin@example.com	

From this report, you can also navigate back to the Configure User Quotas Tab batch entry, by clicking on the dark blue "Edit" button.



Configuration File

You can also see your User Quotas in the NetEqualizer Configuration File. From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[RTR\] -> Configuration](#). If your configuration file is more than 25 lines, you may need to click on the "Last" button to go to the end of your file.

This excerpt shows the two User Quota Rules that we displayed in the View All Traffic Report above. The rows will start with "QUOTA", followed by IP address or subnet, followed by the CIDR, then Quota Amount repeated twice, then Contact email, Duration and finally the Hard Limit Restriction. All amounts are in bytes per second in the Configuration File.

43	QUOTA	192.168.1.244/32 1000000000 1000000000 admin@example.com 1440 62500
44	QUOTA	192.168.1.245/32 500000000 500000000 admin@example.com 1440 125000



Viewing Quotas in Action

If you would like to see the quota amounts consumed for each IP, from the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[RTR\] -> Active Connections -> View Quota Report](#). The following screen report screen opens.

Status of All IPs

On the View Quota Report, you can see where every IP with a Quota Rule defined is against their current Quota Amount. This report will display *all* active IPs involved in Quota accounting. It shows the date & time when quota started being collected, the total bytes collected, if a restriction is currently in place, and when the restriction ends.

Interpreting the Report

For the 1st IP address (**192.168.1.244**) in the screen at right, here is what the report shows. You can first see the Quota Amount in bytes (1GB) and the Quota Duration (1 Day). Then the Total Bandwidth down in Megabytes (110MB) during the quota capture period, and the total Bandwidth up (14MB). You can also see when the quota counters will be reset (Reset time: Today at 11:59pm).

See [Quota Rules Parameters Table](#) for detailed definitions of Quota Amount and Quota Duration.

View Quota Report

Home / RTR / Active Connections / [View Quota Report](#)

Use this page to view the status of all IPs.

STATUS OF ALL IPS

IP: 192.168.1.244
Quota amount: 1000000000 B
Duration: 1 day
Total bandwidth down: 110 MB
Total bandwidth up: 14 MB
Reset time: Today at 11:59pm

IP: 192.168.1.245
Quota amount: 500000000 B
Duration: 1 day
Total bandwidth down: 330 MB
Total bandwidth up: 166 MB
Reset time: Today at 11:59pm



MAC Redirection

[\(back\)](#)


MAC Redirection is used to define MAC addresses that are authorized to be on your network. Any undefined MAC address is considered unauthorized and will be either: 1) "redirected" to a website of your choosing, or 2) dropped.

Case #1, redirection to a web site, is typically done to inform unauthorized users how to subscribe to your network or who to contact about your network and its use. Case #2 is used when you elect to just drop all unauthorized MACs, instead of redirecting them.




To access the MAC Redirect commands

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Maintenance] -> Manage Firewall*. The window below opens, defaulted to the Start/Stop Firewall tab.



To learn more about setting up MAC redirection

Click on -> Edit Rules File Tab. The screen shown below will open. *Then Click on -> Sample Firewall Rules*, (circled in blue below). This will take you to a webpage with Firewall Samples. Click on the links to view the samples.



You will find examples of setting up MAC Redirection. Should you need assistance please call our Support Team at 303.997.1300 x102 or email support@apconnections.net.

Note: MAC Redirection questions and support are not covered in the normal setup of the NetEqualizer product ([NSS](#)) and must be negotiated separately.

To set up all authorized MAC addresses

You need to add ALL authorized MAC addresses (the MAC addresses you wish to allow on your network). Make sure to include your DNS servers in the allowed list. Click on the Manage Macs.Allow Tab to open the window shown below.



Manage Firewall

Home / Maintenance / Manage Firewall

START/STOP FIREWALL EDIT RULES FILE EDIT CLEAR FILE **MANAGE MACS.ALLOW** VIEW SETTINGS

Use this page to manage the macs.allow file.

- MAC addresses are 6-byte hex addresses with leading 0s in each field -- 02:d2:04:29:f0:11 is valid, while 2:d2:4:29:f0:11 is not.
- No changes are made to macs.allow until they are explicitly saved.

		MAC to Allow	Name or Description
	1	00:1C:B3:09:85:15	Lab Server
	2	13:21:CC:69:F0:F1	Router

Click on the 1st field, MAC to Allow. Type in a valid MAC address to allow, with colons as the separator: **xx:xx:xx:xx:xx:xx**. MAC addresses are 6-byte hex addresses with leading 0s in each field -- 02:d2:04:29:f0:11 is valid, while 2:d2:4:29:f0:11 is not.

TAB to the Name or Description field. Type in: **name or description**, which will help you to find this in the future.

Click on any of the "+" icons (circled in blue) to continue adding MAC addresses to allow. Any unsaved fields are shown in yellow, as a visual reminder that you have not yet saved your changes. If you were to click off this screen before saving changes, your data would be cleared.

Once you have entered all of your limits, **Click on -> [Save Changes]** to save changes or **Click on -> [Reset]** to discard changes. We are now prompted to Restart the Firewall, for the new MACS to Allow to take effect. Once we **Click on -> [Restart Firewall]**, our new MACS to Allow will be available to the NetEqualizer Firewall.

Note: Each MAC and associated name or description must be unique.

To remove a MAC address from your authorized list:

In order to remove a MAC address from your authorized list, click on the red "x" button on the row that you wish to delete. The MAC to Allow row will disappear. To completely remove the MAC to Allow from the configuration file, you need to **Click on -> [Save Changes]**. We are now prompted to Restart the Firewall, for the MAC to Allow rule to be completely removed from the NetEqualizer Configuration. Once we **Click on -> [Restart Firewall]**, our MAC to Allow will be permanently deleted.

At this point only the authorized MAC addresses will pass through the system, the rest will be blocked. When enabled, MAC redirection looks at the macs.allow file when an outgoing connection is made from your network out to the Internet. If the user has a browser active, and the MAC address is unauthorized, it will drop the connection, unless you have redirected their browser to a website of your choosing.

To select the website to redirect to:

Follow the instructions in the Sample Instructions referred to in "To learn more about setting up MAC redirection" above.



Perform Quick Edits

The NetEqualizer has always supported the ability to add or delete rules without having to restart the NetEqualizer process. As of Software Update 8.4, we have added our batch entry capability, to better support the initial setup process.

We have also streamlined the add/delete process, by moving all rules to a "Quick Edit" area. These screens can be used when you need to add or delete one rule or a small number of rules. The main advantage to Quick Edits is that you do NOT need to restart equalizing after making your changes.

From the NetEqualizer Dashboard or Navigation Menu, *Click on -> [Setup] -> Perform Quick Edits*. The following screen opens, defaulted to the Hard Limits Tab.

You can select any of the seven (7) tabs to Add or Delete Rules: 1) Hard Limits, 2) By Pool, 3) By VLAN, 4) Masked Hosts, 5) User Quotas, 6) P2P Traffic, or 7) Priority Traffic.

Each screen in the Perform Quick Edits works in the same fashion. We will describe the Hard Limits screen here.

In order to perform any quick edits, Equalizing must be started. In most cases, Equalizing will be on. If Equalizing is not running in your environment, you will see the following warning message.

Warning! Hard limits cannot be modified while equalizing is stopped. Start Equalizing if you want to modify hard limits.

▶ Start Equalizing

Once you *Click on -> [Start Equalizing]*, equalizing will be started, and you can then edit your rules.



Quick Edit - Deleting a Rule

DELETE HARD LIMIT BY IP ADDRESS

Hard Limit:

In our example from the Quick Edit – Hard Limit Tab, we show the Delete Hard Limit By IP Address section of the screen. On the Hard Limit field, *Click on -> [dropdown arrow]* to select an existing Hard Limit to delete. The *Click on -> Delete*. The Hard Limit will be permanently deleted from the NetEqualizer configuration. You do NOT need to restart equalizing.

Quick Edit - Adding a Rule

ADD HARD LIMIT BY IP ADDRESS

Host IP	/	CIDR	Download Mbps	Upload Mbps	Burst Factor
	/	32			1

In our example from the Quick Edit – Hard Limit Tab, we show the Add Hard Limit By IP Address section of the screen.

To create a Hard Limit, *click on your computer's TAB to put the focus into the Host IP address field*. Type in: *Host IP address* in 11.22.33.44 format.

TAB to CIDR field. On the CIDR field, *Click on -> [dropdown arrow]* to select a CIDR value. For an individual IP, use /32, for a Class B use /16, for a Class C use /24, or any other subnet value from 1-32. If using subnets, each IP in the subnet will get the Hard Limit.

TAB to Download Mbps Hard Limit field. Type in: *a Positive Number*. TAB off the field to complete your entry.

TAB to Upload Mbps Hard Limit field. Type in: *a Positive Number*. TAB off the field to complete your entry.

TAB to Burst Factor. We recommend that you keep this set to "1". To read more about setting the Burst Factor to a value other than 1, see the [Add Bursting to Hard Limits](#) section of this User Guide.

Click on -> [Save Changes] to save changes or *Click on -> [Reset]* to discard changes. If you have saved changes, your new Hard Limit will be permanently added to the NetEqualizer configuration, and will take effect immediately.

In the light blue box above the entry fields, you will find notes to help you in creating Hard Limits. Links are in orange, which you can click on to either get more information, or to move to another screen.

Note: Perform Quick Edits screens cannot be used to modify rules. Please see the batch entry screens for each Rule type (Hard Limits, Pools, etc.) to modify rules.



Distributed Denial of Service Attack (DDoS) Tools

A Distributed Denial of Service Attack (DDoS) occurs when a hacker illicitly gains access to a system, takes it over, and then uses it to command many systems to flood a target network with traffic. The flood of traffic quickly overwhelms the target network, and causes the network to become inoperable for its normal purposes. Read more about [DDoS](#).

As the NetEqualizer is implemented near the network perimeter, and has visibility to all incoming and outgoing traffic, we are able to analyze traffic behavior and report on suspected DDoS attacks. Once identified as suspected DDoS, the NetEqualizer can then be used to block further traffic as needed.

As of [software update 8.2](#), we are offering a new [Distributed Denial of Service \(DDoS\) Monitor](#). The DDoS Monitor, which comes standard in 8.2, ***shows you some basic metrics on the outside intrusion hit rate into your network***. It can be used to spot anomalies that would indicate a likely DDoS attack in progress. See our [detailed blog article](#) on the subject for how this technology works.

If you decide you need something more proactive to mitigate a DDoS attack, we install our DDoS Firewall (DFW) feature, and provide consulting to help you configure it to block standard DDoS attacks. Our new DDoS Firewall tool (DFW) can be purchased as an add-on module. The goal of the NetEqualizer DDoS toolset is twofold: 1) to help you identify suspected DDoS attacks, and 2) to help you identify and block outside IP addresses until you have regained control over your network.



DDoS Monitor

The DDoS Monitor is used to analyze *unrequested incoming traffic* to look for inbound traffic that occurs both at high frequency and is repeated a large number of times, which is behavior typical of a DDoS attack. We then provide visibility into this traffic to help you monitor your network, through our new DDoS Monitor.

From the NetEqualizer Dashboard or Navigation Menu, [Click on -> \[DDoS\]](#). The following screen opens.

DDoS Monitor

DDoS MONITOR

Use this page to view all calculated requests running through the NetEqualizer.

Index	Src IP	Src Port	Port	Seconds	Count	App	Blocked
1	203.10.156	80	80	25	66	2	No
2	203.10.156	80	80	116	188	18	No
3	203.10.156	80	80	89	235	25	No
4	203.10.156	80	80	33	49	1	No
5	203.10.156	80	80	3	0	0	No
6	203.10.156	80	80	30	276	27	No
7	203.10.156	80	80	10	50	2	No

Showing 1 to 7 of 7 entries (generated at 20180105T10:07:00)



The DDoS Monitor displays all uninitiated requests coming into your network. You can see how persistent the request is (Seconds) and how often it is hitting your network in the last second (Rate), which then gives you an overall view (Count) of how active the attack is. For example, in our table above, Index 5 has been running for 99 seconds, hitting the network 27 times per second, for a total of 2,756 hits.

By analyzing the values of Count, Rate, and Seconds, you can identify which external IP addresses you want to block. In our example above, Index #1, #2, and #5 (circled in blue above) are candidates to consider blocking.

All fields of the DDoS Monitor are defined in detail below.

DDoS Monitor Fields

Field	Definition	Expected Values
Index	Table row #	0, 1, 2, 3....
SRC IP	The source IP for this connection	External IP address Suspect for a DDoS attack.
DST IP	The destination IP for this connection	Internal IP address
Port	Whether the traffic was initiated internally or externally.	1 = inside initiated 2 = outside initiated On the DDoS Monitor, this should always be 2.
Seconds	How long the requests have been running, in seconds.	Positive number greater than zero.
Count	The number of times the request has run over the Seconds indicated.	Positive number greater than zero.
Rate	Number of requests per second during the last second.	Positive number greater than zero.
Blocked	Indicates if the SRC IP address has been blocked using the DDoS FW.	yes = blocked no = unblocked

*Note: To keep this simple, Rate is just the **rate over the last second**. It is not an average over time. Due to this, the Count field will not equal Rate x Seconds, and may not even approximate it, depending on how variable Rate has been over time. However, the combination of Rate, Seconds, and Count are a great indicator of whether an external IP is involved in a DDoS attack on your network.*

DDoS Firewall

The DDoS Firewall (DFW) is an Add-on Module made up of a set of intelligent tools and consulting. The DFW is used to block external IP addresses that you suspect of being involved in a DDoS attack.

If you are reviewing results from the DDoS Monitor, and would like to take it further to block external IP addresses, please contact our Support Team at support@apconnections.net or 303.997.1300 x102.



Monitoring and Reporting

NetEqualizer provides both real-time and historical reporting through our **Dynamic Real-Time Reporting (RTR)** capability. Reports are offered in tabular and graphical formats. This enables you to see data in a format that is most meaningful to you, over a variety of timeframes.

Dynamic Real-Time Reporting (RTR) is imbedded within the NetEqualizer. It was first released in [software update 7.1](#). We have since updated it in [software update 7.4](#) by adding Traffic Reports, and further enhanced it in [software update 8.1](#) by adding Historical Data. In [software update 8.3](#), we expanded RTR to include Penalty Data, Top Usage Data, and the Export capabilities. With RTR, you are able to see real-time and historical data regarding bandwidth usage and traffic shaping (penalties) on your network, which you can sort, search, and graph, to help you better manage your network.

As of software update 8.3, all reporting in the NetEqualizer is accessed via the Dynamic RTR icon on the Common Tasks bar of the main Dashboard.



Dynamic RTR enables you to see all reports offered in our full reporting suite. There are two key categories of reports: Active Connections and Traffic History, which are described below. RTR also contains the NetEqualizer Log, Configuration, and Running Processes.

Active Connections Reports enable you to see what is going on in your network at this moment, in order to actively monitor and manage your network usage. Active Connections tabular reports are sortable and searchable real-time views of all connections or penalties active on the NetEqualizer. They also include Lookups by IP (country, DNS, rules), as well as links to the Traffic by IP Graph.

Traffic History Reports provides you a graphical view into the trends of bandwidth usage and traffic shaping on your network across time. This can help you in network design and planning activities, as well as to determine if your bandwidth level requirement is stable or increasing.

As of software update 8.1, you can store and view up from ten (10) minutes up to four (4) weeks of data on the NetEqualizer. In software update 8.3, you can also run [Export Data](#) to maintain additional history on a separate server. Contact support@apconnections.net if you need help with the Export Data function.

Notifications In [software update 5.5](#) and above, you can set-up Email Alerts, entering an email address that the NetEqualizer will send alerts & notifications to for selected events.


Before you start using reporting, it is **recommended that you set up your NetEqualizer to synchronize with either an internal NTP time server or an Internet Time Server**. That way, your date and time will always be accurate in your reports. Please follow the instructions in [Appendix #5](#) to sync your NetEqualizer Date/Time.



Dynamic Real-Time Reporting (RTR)

The RTR Menus are shown at right. This menu is displayed when you click on Dynamic RTR. The following reports and functions are available via RTR:

1. [RTR Dashboard](#) - View traffic in real-time flowing through NetEqualizer. Contains two dashboards, one for the entire network and the other by Pool.
2. [Active Connections](#) - Menu containing our four (4) key real-time reports: 1) IPv4 Active Connections, 2) IPv6 Active Connections, 3) Connection Counts, and 4) Active Penalties. Active Connections reports are sortable and searchable real-time views of all connections or penalties active on the NetEqualizer. Links to IP Lookups (country, DNS, rules), as well as Traffic by IP Graph.
3. [Traffic History](#) - Menu containing our five (5) traffic history reports: 1) General Traffic History, 2) Traffic History by IP/Pool/VLAN, 3) Top Talkers, 4) General Penalty Reports, and 5) Export Data. Traffic History reports are graphs or tables that show from 10 minutes up to 4 weeks worth of upload and download bandwidth usage (or penalties) for your entire network, Pool, VLAN, or tracked IP.
4. [NetEqualizer Log](#) - Displays key activity on the NetEqualizer, such as limits being applied, and penalties being added or removed.
5. [Configuration](#) - View how you have defined the key parameters, traffic limits, priorities, and P2P limits on your NetEqualizer. Use this to validate your settings.
6. [Running Processes](#) - Check out what processes are running on your NetEqualizer.
7. [RTR Preferences](#) - Select units in which to view your graphical data.
8. [Start/Stop RTR](#) - Turn RTR on/off and view RTR statistics.
9. [Autostart RTR](#) - Turn RTR on automatically upon a NetEqualizer reboot.



Notifications

In addition to RTR, we offer notifications via the Manage Alerts function. Set-up emails to notify, and select events that send alerts and notifications.

1. [Configure Email](#) - Set up email to send alerts to.
2. [Configure Alerts](#) - Select events to send email alerts on.



Dynamic Real-Time Reporting (RTR)

One of the things that has always differentiated the NetEqualizer from other monitoring and shaping tools is that *we have the actual data for every user accurately updated by the second*. Thus, we are able to make shaping decisions based on usage every second. This sets us apart from other network tools that report on traffic.

The reporting tools on most other devices tend to slog along and show you aggregate usage of five (5) minute averages. Even the charge back mechanisms that Internet providers use to figure out if you are over your allotted bandwidth do 95th percentile sampling, meaning they estimate your usage from sporadic sampling.

One thing we have not focused on, until now, is making this wealth of data available to our customers in a nice, organized, usable format. With the advent of our 64-bit release (7.0) and a more robust Apache web server, we are now able to display this data in real-time.

All data is current as of this second when displayed. You can click on Update Data where appropriate to refresh the data. From the Dashboard, *Click on ->[Dynamic RTR]*.



RTR Dashboard

(back)

The following RTR Dashboard opens, defaulted to the Real-Time General Traffic Graph, and below it, the Real-Time Pools Data Graph.

RTR reports are listed in a vertical menu on the left-hand side of the screen (circled above): 1) Active Connections, 2) Traffic History, 3) NetEqualizer Log, 4) Configuration, 5) Running Processes, 6) RTR Preferences, and 7) Start/Stop RTR.



You can compress this menu to just displaying icons, or expand it back to show full menu names by clicking on the compress icon, shown at left and circled above.




You can return to this screen from any other report by *Clicking on -> Dashboard -> RTR* (circled above).



Real-time General Traffic


In our graph below, we have compressed the menu down to icons, so that we can focus on the Real-Time Traffic (RTT) graph itself. On this graph you can see bandwidth consumed for all traffic flowing through your NetEqualizer. Click on a radio button to see this graph for Downloads, Uploads, or Both. Each time you come back to the RTR Dashboard, it starts the graph anew.

On the x-axis, you can see that the scale is set to every 10 seconds, and the Date/Time is shown below the graph. Date/Time matches whatever is set for your NetEqualizer in Manage NetEqualizer -> [Configure Date/Time].



Real-time Pool Data


We also have a bar chart of all the Pools defined on your NetEqualizer. Real-time Pool Data shows upload and download bandwidth usage for each Pool. If no Pools have been defined, you will only see the default pool, Pool 0. Pool 0 contains your entire network.



In our graph above, we have two defined pools, Pool 1, with shared limit of 20Mbps up/20Mbps down, and Pool 2, with 50Mbps up/50Mbps down. On the Pools chart you can see bandwidth consumed for each Pool, and whether a pool is being equalized. In this example, Pool 1 has reached peak for downloads (circled above), as the green download bar has touched the red equalizing line. The red "equalizing" line on each Pool will match your [RATIO](#) parameter. At the right is the key to the graph (circled).



This chart is a great way to see how your Pools are performing in real-time. You can read more about Pools in the User Guide [Bandwidth Limits/Setting Up Bandwidth Pools](#) section.



Dynamic RTR


RTR Active Connections Reports

[\(back\)](#)

Active Connections Reports enable you to see what is going on in your network at this moment, in order to actively monitor and manage your network usage. Active Connections tabular reports are sortable and searchable real-time views of all connections or penalties active on the NetEqualizer. They also include Lookups by IP (country, DNS, rules), as well as links to the Traffic by IP Graph.

To view the Active Connections Reports from RTR Menu, **Click on -> Active Connections** to open the menu.

The menu expands to show the options listed at right.



View Active Connections

[\(back\)](#)

The first report, Active Connections, shows the data streams (pairs of IP addresses) that are currently live on your network for all IPv4 and IPv6 traffic. This will show all connections your NetEqualizer is currently seeing. You can utilize this report to see what data streams are "hogging" your network by looking at the Wavg value. Wavg values over HOGMIN will be equalized when your network is congested.

To view Active Connections, **Click on -> View Active Connections**. You can also access this report directly from the Dashboard. On the Common Tasks bar, **Click on -> [View Current Activity]**. The following View Active Connections table opens.

View Active Connections

RTR - View Active Connections

ACTIVE CONNECTIONS

Use this page to view all active connections running through NetEqualizer.

25 records per page

Search:

Index	SRC Port	DST Port	Wavg	Avg	DST IP	SRC IP	Protocol	Port	Pool	TOS
0	49808	80	16	11	192.168.1.143	192.168.1.113	TCP	1	3	1
1	80	49763	247517	406660	192.168.1.113	108.171.213.134	TCP	2	3	1
2	49763	80	3465	8595	108.171.213.134	192.168.1.113	TCP	1	3	1
3	49761	50134	439761	397171	97.74.144.172	192.168.1.113	TCP	1	3	1
4	50174	49761	8833	7692	192.168.1.113	97.74.144.172	TCP	2	3	1
5	80	49808	15	0	192.168.1.113	192.168.1.143	TCP	2	3	1

Showing 1 to 6 of 6 entries (processed in 0.00060200691223145 s)

First Previous 1 Next Last

Update Data



The Active Connections Table is sortable and searchable, so that you can focus on the connections you are most interested in viewing. Use this to easily find bandwidth hogs by sorting on Wavg and displaying largest to smallest (descending). The report fields are defined in the [Fields of RTR View Active Connections](#) table below.

You can sort the report on any column by clicking on the arrows in the column header. Sort will default to ascending on 1st click, and then descending on 2nd click.

To refresh the data displayed, [Click on -> \[Update Data\]](#), (circled above).

You can also see a Search box above the table on the right. If you enter in an IP address or partial IP, the report will filter results to only display active connections for the selected IP.

Below each IP you can see links to open IP Lookups (circled above): "C" for IP to Country Lookup, "DNS" for DNS Lookup, "AR" for All Rules Lookup, and new in software update 8.1 "T" for Traffic History by IP Graph. IP Lookups are detailed below in the [IP Lookups](#) section.

Fields of RTR View Active Connections

Field	Definition
Index	Table row #
SRC Port	The source port for this connection
DST Port	The destination port for this connection (the service being requested http, FTP, etc.)
Wavg	A weighted average of total bytes on this connection per second for the last eight seconds. Used to determine if the flow is a bandwidth hog (over Hog Minimum, which defaults to 12000 bytes).
Avg	The average in bytes per second since this IP pair came into the table
DST IP	Destination IP address involved in the connection.
SRC IP	Source IP address involved in the connection
Ptcl	The protocol (ICMP, TCP/IP, UDP). For IPv6 traffic mapped to an IPv4 address, this will show "- -"
Port	Outbound (value = 1) or Inbound (value = 2).
Pool	Pool #. Default is 0 (no bandwidth pools set-up). Otherwise, bandwidth pool #. If you have VLANs set-up, this will show the VLAN #.
TOS	0 if bit not set ("off"). Greater than 0 (>0) if bit is set ("on").

IP Lookups


([back](#))

On the View Active Connections report, you can click on a symbol below an individual IP address to view the associated IP Lookup.

For each individual IP, you can now select from one of four IP Reports: 1) Country Lookup "C", 2) DNS Lookup "DNS", 3) All Rules Lookup "AR", and 4) view Traffic History by IP Graph "T".

As of [software update 8.1](#), if an IP address is being "tracked" and therefore is eligible to be viewed in a traffic graph, we display a "T" below the IP address.

You can also go to several of the IP Lookups directly through the Active Connections Menu (circled at right).





Country Lookup for an IP

On Row 1 of the Active Connections table, if I **Click on** -> "**C**", (Country Lookup), for the SRC IP 108.171.213.134, the screen at right opens.

In this example, the country is the France.

Note: You need to use an external IP for this lookup to bring back a value.

NSLookup for an IP

On row 3 of the Active Connections Table, if I **Click on** -> "**DNS**", (DNS Lookup), for the DST IP 97.74.144.172, the screen at right opens. You can see in the resulting window a non-authoritative answer has been found.

Note: You need to use an external IP for this lookup to bring back a value.

Show All Rules for an IP

On the Active Connections Table, if I **Click on** -> "**AR**", (All Rules Lookup), for the DST IP on Row 1 192.168.1.113, the screen at below opens.


Show All Rules for an IP enables you to see the rules that have been set up for an individual IP. This is great way to determine if an IP has associated hard limits, or is a member of a Pool/VLAN limit. You can determine whether connection limits apply, and whether the IP has priority or has been masked. If you have set a quota for the IP, that is displayed as well. And most importantly, the validity of your configuration for that IP is checked, and you are warned if something is wrong.

In our example, 192.168.1.113 is a member of Pool 3 and has no validity issues.

Traffic History by IP Graph

For 192.168.1.113 on Row 1 of the Active Connections Table in our example, a "T" is displayed below the IP address. This means that 192.168.1.113 has been defined as a "tracked IP". See Traffic History -> Manage Tracked IPs section to learn how to set up a tracked IP or subnet. You also must have Started RTR in order to view Traffic History (Start/Stop RTR -> Start RTR).

If I **Click on** -> "T", (Traffic History by IP), for the DST IP on Row 1, 192.168.1.113, the screen below opens.



By default, a filter is applied so that the traffic displayed is only for the IP selected, in this case 192.168.1.113. You can see the filter that is set at the bottom left of the graph. In this case "by IP Address" is selected, and Value is 192.168.1.113. The traffic shown is in whatever units have been selected under RTR Preferences -> Traffic by IP/Pool/VLAN Graph Units. In this case, I have selected Megabits per second (Mbps).

The graph defaults to 1 hour for the last hour. Under "Show data for last" on the top left, you can change the graph to show 10 minutes, 30 minutes, 12 hours, 1 day, 1 week or 4 weeks of data. In this example, I have updated the graph to show the last 10 minutes of data.

Radio buttons on the top right can be used to select whether you display download traffic, upload traffic, or both on the graph.

And finally, below the graph on the x-axis, you can see the amount of data uploaded or download during the time period displayed for the selected filter, in this case 100.39MB downloaded and 89.18MB uploaded for 192.168.1.113.

This Traffic Report is a good tool to help you to investigate how much traffic any IP is consuming. You can change the time period displayed to show data from 10 minutes to 4 week old. In order to get data into this report, you must first have added the IP or IP



subnet to data being tracked, and also Started RTR by going to Start/Stop RTR and clicking on Start RTR. Both are described in detail in the Traffic Reports section of this User Guide.

Note: By design, the graph does not refresh. Click on ->[Update Data] below the graph to refresh your data.

IPv6 Traffic

As of [software update 8.2](#) and above, Active Connections includes IPv6 traffic, to help those running dual stacks (IPv4/IPv6) and to prepare for the migration over to full IPv6. When we examined real IPv6 traffic on a live network, as expected the upper bytes in the address rarely, if ever, changed. So by taking the lower 24 bits of the IPv6 address and mapping that into a locally unique IPv4 address, we can show and shape all the traffic in one table.

We will now look at IPv6 traffic that has been mapped to IPv4 address space.

Index	SRC Port	DST Port	Wavg	Avg	DST IP	SRC IP	Pttl	Port	Port	Port
39	39019	39019	219446	188437	10.0.18.2	10.0.18.108	10*	0	0	1
14	39019	39019	207943	188437	10.0.18.2	10.0.18.108	10*	0	0	1
21	39009	39019	189983	88988	10.0.18.2	10.0.18.108	10*	0	0	1
8	0	0	173834	75000	254.239.235.0	254.239.235.0	10*	0	0	1
9	39006	39009	173834	75000	254.239.235.0	254.239.235.0	10*	0	0	1
11	39012	39011	127360	152886	10.0.18.2	10.0.18.108	10*	0	0	1
5	39002	39017	49875	49800	254.188.28.79	254.188.28.79	10*	0	0	1
4	37864	35539	49800	49800	254.188.28.79	254.188.28.79	10*	0	0	1
7	39019	35539	39400	17888	254.239.235.0	254.239.235.0	10*	0	0	1
10	39018	39017	28486	15482	10.0.18.2	10.0.18.108	10*	0	0	1


This Active Connections table above is now supporting a dual stack. Rows where the Pttl is blank "--" (circled in the table above) are the IPv6 rows. You can also see that the DST IP and SRC IP for these rows are in IPv4 format.

As of [software update 8.2](#), equalizing is based on the total bandwidth across both sets of addresses, and not a separate decision for IPv6 and IPv4. Based on feedback from our customers, we are hearing that IPv6 traffic may now be a noticeable percentage (*reported at up to 10%*) of the traffic traversing your link to the Internet.

We began this transition by providing you with a way to see how much IPv6 traffic is passing through your network. We also offer a **View Active IPv6 Connections**, so that you can see the native IPv6 addresses as well.


View Active IPv6 Connections

([back](#))



The second report on the Active Connections Menu is used to view IPv6 Traffic protocols and to see unmapped (native) IPv6 addresses. To view IPv6 connections, *Click on -> Active Connections -> View Active IPv6 Connections.*

Active IPv6 Connections shows all IPv6 traffic streams (IP pairs) currently active on your network. This report gives you a sense of how much of your network traffic has converted over to IPv6.





In the example (circled below), you can see DST IP and SRC IP in IPv6 format, along with the traffic protocol. The report fields are defined in [Fields of View Active IPv6 Connections](#).

Index	SRC Port	DST Port	Wavg B/s	Avg B/s	DST IP	SRC IP	Ptcl
1	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
2	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
3	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
4	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
5	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
6	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
7	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP
8	0	0	077800	292546	fe80::0000:0000:0000:0000:0000:0000:0000:0000	fe80::0000:0000:0000:0000:0000:0000:0000:0000	TCP



Fields of View Active IPv6 Connections

Report Field	Definition
Index	Table row #
SRCP	The source port for this connection
DSTP	The destination port for this connection (the service being requested http, FTP, etc.)
Wavg	A weighted average of total bytes on this connection per second for the last eight seconds
Avg	The average in bytes per second since this IP pair came into the table
IP1	Destination IP address involved in the connection.
IP2	Source IP address involved in the connection
Ptcl	The protocol (ICMP, TCP/IP, UDP).

View Connection Counts

([back](#))

The third report on the Active Connections Menu, View Connections Count, is used to see if (and to what extent) P2P traffic is present on your network. Use this report to determine if you have IP addresses with unexpectedly high numbers of connections. You should also view this report to better understand how many inbound and outbound connections you need to support valid activities on your network, before you set any [Connection Limits](#).



We recommend monitoring your installation for several days before setting Connection Limits. For example, some online games require 60 or more total connections, without being P2P traffic. If you wish to allow this type of activity, you would need to set your connection limits to 60.

To monitor your connections, [Click on -> Active Connections -> View Connection Counts](#). The following report opens. The report fields are defined in [Fields of View Connection Counts](#).



View Connection Counts
[View Connection Counts](#)

CONNECTION COUNTS Update Data

Use this page to view connection counts for all IP's running through NetEqualizer.

10 records per page Search

IP	In	Out	Total
10.0.10.2	3	2	5
11.111.8.188	0	2	2
13.80.128.21	0	2	2
10.104.114.248	0	2	2
10.57.18.0	0	2	2
10.222.28.53	0	2	2
11.225.88.85	0	2	2
11.80.126.202	0	2	2
11.130.136.40	0	2	2
10.241.238.250	0	2	2
IP	In	Out	Total

Showing 1 to 10 of 90,900 entries (processed in 0.43837641857058 s)

Fields of View Connection Counts


Report Field	Definition
IP	IP address involved in the connection.
In	The number of inbound connections for this IP.
Out	The number of outbound connections for this IP.
Total	Total number of connections used by this IP.

View Active Penalties

([back](#))

The fourth report on the Active Connections Menu, View Active Penalties, is used to see equalizing in action. This report shows you equalizing in the current moment. You will see where equalizing is adding, increasing, or decreasing penalties on data streams on your network. There are three (3) columns: New Penalties, Increased Penalties, and Decreased Penalties.

To monitor your active penalties, [Click on -> Active Connections -> View Active Penalties](#). The following report opens. Each column shows the data flows (IP pairs) that are being equalized.



ACTIVE PENALTIES Update Data

Use this page to view all active penalized connections running through NetEqualizer, and their associated status (New Penalty, Increased Penalty, Decreased Penalty).

New Penalties	Increased Penalties	Decreased Penalties
10.0.10.2 to 10.0.10.104	10.0.10.2 to 10.0.10.104	10.0.10.2 to 10.0.10.104
10.0.10.2 to 10.0.10.103	10.0.10.2 to 10.0.10.103	10.0.10.2 to 10.0.10.103
		10.0.10.106 to 10.0.10.2

Last Update: June 18, 2015 10:32:03 AM MDT



RTR Traffic History Reports

[\(back\)](#)

Traffic History Reports enable you to quickly see how busy your network has been over a period of time. Reports are available in seven time increments, showing data from 10 minutes to up to 4 weeks. You can view your entire network, or hone in on an individual IP, Pool, or VLAN for analysis.


With Traffic History Reports, you can track NetEqualizer traffic (bandwidth usage), as well as look at patterns of upload and download usage. You can view the top bandwidth users, and see what penalties have been applied over time. You can also export data to review for longer periods of time.

With this tool, you can get a better handle on how busy your network is over time, which is useful in capacity planning, and also what IPs are consuming your bandwidth.

From the RTR Dashboard, *Click on -> Traffic History*. The screen at right opens. In the Traffic History menu, there are currently five (5) reports and two management capabilities.

In brief, General Traffic History enables you to see bandwidth usage for your entire network. Traffic History by IP/Pool/VLAN shows history for tracked IPs within selected entities. Top Talkers highlights the biggest bandwidth users (by IP) on your network. General Penalty Reports shows how equalizing has impacted your network over time. Export Data enables you to save off reporting data if you need to maintain > 4 weeks of data.

Manage Tracked IPs is used to set up which IPs or subnets will be available for reporting. Clear Reporting Data will erase all historical data.



We will review Traffic History reports and management capabilities in detail below.




Start RTR

([back](#))

By default, Traffic History reports are OFF. You must first turn these on in order to use Traffic History.

Click on -> *Start/Stop RTR*, and then Click on -> *[Start RTR]* (circled at right) to turn on Traffic History reporting.

Once you have turned on traffic reporting, the System Status 1st line will be green and say "Traffic History Reporting is currently ON." (circled at right).



Manage Tracked IPs

([back](#))

For many Traffic Reports, in order to populate data into the report, you need to let the NetEqualizer know which IPs you would like to track, or collect data against. We call this "tracked IPs".

To set up your tracked IPs, Click on -> *Manage Tracked IP's*. The screen on the right will come up. Type in an IP or subnet that you would like to track. Make sure to enter a NEWLINE after each entry. Once complete, Click on -> *[Save Changes]* to keep your edits or Click on -> *[Reset]* to discard your changes.



To see which IPs are being tracked, just come back to this screen at anytime. In our example, you can see that 192.168.1.1/24 is being tracked. Any IP address within this subnet will have data recorded for Traffic History Reports.

Note: In release 8.1, you need to enter ALL IPs or subnets associated with a Pool or VLAN, in order to accurately see all data within the tracked Pool or VLAN. We hope to change this in the future to enable you to just select the Pool # or VLAN # to track.




General Traffic History

([back](#))

The first report, General Traffic History (GTH), is a graph showing all traffic flowing through the NetEqualizer. As this graph uses data sampled every 1 minute, the graph is intended to show data use over time, not an exact bandwidth use per second.

General Traffic History defaults to showing the last 1 hour of traffic in MBps. You can change the time viewed from 10 minutes up to 4 weeks, and can view the graphs in megabytes (MBps) or megabits (Mbps). The graph will use whatever units you have selected in RTR Preferences -> General Traffic/Real-Time Graph Units.

You will see the "Both" radio button selected. You can click on Download, Upload, or Both, to change what bandwidth usage is displayed. Information and tips (above the graph) shows you useful information about the data displayed. If you want to zoom in on the graph, use your browser commands.




Note: By design, the graph does not refresh. Click on ->[Update Data], below the graph on the left-hand side, to refresh your data.

Traffic History by IP/Pool/VLAN

([back](#))

The second report we discussed in the [Traffic History by IP](#) section, focusing on the "by IP address" filter, which is used if you go to this graph from the Active Connections table. If you navigate to The Traffic History by IP/Pool/VLAN graph from the Traffic History menu, the following screen appears. There is no graph displayed until a value is entered in the Value field (circled below).





In our example below, we clicked on the “by Pool” radio button, entered “3” in the Value field, and then clicked on “Update Data” and also selected the Show data for last: “30 minutes” radio button. The Traffic by Pool graph is displayed for Pool 3. This graph shows all the traffic for tracked IPs in Pool 3.

Note: If you did not track an IP that is in Pool 3, the data will not be included here.

Traffic by Pool Graph



By default, the Traffic Reports by IP/Pool/VLAN show data in kilobytes per second (KBps) for the last hour. You can use the RTR Preferences menu to change the units displayed to kilobits per second (Kbps), kilobytes per second (KBps), megabits per second (Mbps), or megabytes per second (MBps).



Top Talkers


[\(back\)](#)

The third report, Top Talkers, highlights the biggest bandwidth users (by IP) on your network. This report is useful to see the "top 10" IPs consuming upload, download, or both on your network.

In the sample report below, showing data for a 10 minute period, you can see that 192.168.1.106 has the largest download usage, while 192.168.1.100 has the largest upload usage.

Just like the General Traffic History Report, Top Talkers defaults to showing the last 1 hour of traffic in MBps. You can change the time viewed from 10 minutes up to 4 weeks, and can view the graphs in megabytes (MBps) or megabits (Mbps). The graph will use whatever units you have selected in RTR Preferences -> General Traffic/Real-Time Graph Units.

You will see the "Both" radio button selected. You can click on Download, Upload, or Both, to change what bandwidth usage is displayed. Information and tips (above the graph) shows you useful information about the data displayed. If you want to zoom in on the graph, use your browser commands.





General Penalty Reports

[\(back\)](#)

The fourth report, General Penalty Reports, shows how equalizing has impacted your network over time. This report is useful to see how often penalties have been applied to your network. General Penalty shows the number of penalties applied, and can be viewed for intervals from 10 minutes to 4 weeks.

The report below shows an average of 1,500 penalties being applied to a network over a 30-minute interval.





Clear Reporting Data

([back](#))



If you would like to delete reporting data used for the Traffic History Reports, you can do so.

Click on ->Traffic History -> Clear Reporting Data. The Clear Reporting Data menu (below) opens.

When you click on "Clear General Traffic Data", the data used for the General Traffic History and General Penalty graphs will be deleted from the graphs and storage. General Traffic data will restart collection.

When you click on "Clear IP Traffic Databases", the IP tracking data used for Traffic by IP/Pool/VLAN graphs, Top Talkers, and Export Data will be deleted from the graphs and storage. IP tracking data will restart collection for all tracked IPs and IP subnets.

The screen below to the right will appear to warn you that you are about to delete data. *Click on ->[OK]* to continue or *Click on ->[Cancel]* to keep your data.






Export Data to a Reporting Data Warehouse

([back](#))

Once you have started up RTR reporting, we store up to 4 weeks (1 month) of reporting data on your NetEqualizer. However, as the reporting data is stored in NetEqualizer memory (RAM), we have limited our data storage to 1 month.


If you would like to store data to report against for longer periods of time, you can export data to a separate server to create your own Reporting Data Warehouse.



Click on ->Traffic History -> Export Data. The menu below opens.

Data is downloaded in comma separated value (.csv) format. You can choose to download in one of three ways:

- 1) 24 hours of traffic, which gives you minute-by-minute detail in bytes/second.
- 2) 4 weeks of traffic, which gives your hour-by-hour data from in bytes/second.
- 3) General Penalty data, which gives you minute-by-minute count of penalties over 4 weeks.



View NetEqualizer Log

([back](#))

The NetEqualizer Log File contains a record of the actions of the NetEqualizer. It displays key activity on the NetEqualizer, such as limits being applied, and penalties being added or removed. It is viewable from two menus in the NetEqualizer.

To view the NetEqualizer Log from RTR Menus


From the RTR Menus, Click on -> NetEqualizer Log.



To view the NetEqualizer Log from the NetEqualizer Menus

From the Management and Reporting Menu, Click on -> View Current Activity -> [View NetEqualizer Log].

The NetEqualizer Log File screen defaults to displaying 25 rows. You can change this to 10, 25, 50, or 100 rows (circled below at top left). If you are looking for item, such as an IP address, or the word "PENALTY", you can use the Search field to limit your view to just those items. And just like other reporting screens, if you want to refresh the data, click on the "Update Data" button at the top right.



In the NetEqualizer Log File, you will see three main types of entries, discussed below:

1. **Traffic Up and Down** - Traffic flowing on your network in bytes/second.
2. **PENALTY Entries** - Actual penalties being applied. Contains the word PENALTY followed INCREASE, DECREASE OR REMOVE. Also shows the connection (IP address pair).
3. **PENALTY THRESHOLD** - For informational purposes only. Not actual penalties.

Sample NetEqualizer Log

Value
09/08/14 10:04:39 Traffic up: 942 Traffic down: 38581 POOL 3
PENALTY THRESHOLD pool 3 up 7000000000 down 7000000000
09/08/14 10:04:42 PENALTY REMOVE: 192.168.1.113 108.171.213.134 POOL: 3 TOS: 1
09/08/14 10:04:43 PENALTY IP : 192.168.1.113 108.171.213.134 POOL: 3 Traffic Up: Traffic Down: 554905
Hogmin: 12000 VLAN: TCP WAVG: 278098 BUFF: 1499 DELAY: 5 TOS: 1
Value
First Previous 3 4 5 6 7 Next Last

Traffic Up and Down

Approximately every twenty seconds, the NetEqualizer Log will contain a date and time stamped entry for Traffic Up (outbound) and Traffic Down (inbound). This is instantaneous bytes per second of traffic in each direction flowing on your network.

The default pool 0 (entire network trunk) will always be displayed. If VLANs or POOLS have been defined, you will see a line for each of them as well, with their defined amount of bandwidth.

In the example above, the first line of the sample log shows this data for Pool 3.

PENALTY Entries

A PENALTY entry means that NetEqualizer has decided that a communication link between



two IP addresses (a connection) is using too much bandwidth, and so NetEqualizer has levied a PENALTY against this connection.

The penalty causes all data on this connection to slow down by [PENALTY_UNIT](#). If this connection continues to use too much bandwidth, NetEqualizer will increase the amount of this delay, up to your [MAX_PENALTY](#).

In the 3rd line of the log, a penalty is being removed. In NetEqualizer Log files you can see entries for penalties being applied (PENALTY), increased (PENALTY INCREASE), decreased (PENALTY DECREASE), and being removed (PENALTY REMOVE).

If you are under RATIO on your network, you will not see penalties being applied.

PENALTY THRESHOLD per Bandwidth Pool (internal use only)

You can ignore the PENALTY THRESHOLD lines in your log, as these are to be used by APconnections. The up and down will always be displayed as 7000000000. If VLANs or POOLS have been defined, you will see a line for each of them as well, with 7000000000 displayed for up and down as well.

In our sample NetEqualizer Log, the 2nd line shows a penalty threshold.

Note: PENALTY_THRESHOLD lines are NOT actual penalties being applied to your network.



View the Entire NetEqualizer Log or the Previous Log

While we do not suggest using the log to decide if someone or something is being a bandwidth hog, there may be times where you want to view it in its entirety.

To see the entire NetEqualizer Log:

From the Maintenance and Reference Menu, *Click on->Maintenance->[Run A Command]*.

Type in: cat /tmp/arbblog

To see the entire Previous Log:

Type in: *cat /tmp/arbblog.bak*

Remember that the penalties in the log file show IP connections that may have been just slightly over HOGMIN, and may not have been the reason penalties needed to be imposed at all. It could have simply been that you had an overabundance of good connections/traffic and the total amounted to more than equalizing would allow without something being penalized, so it penalized all connections over HOGMIN.

Configuration

[\(back\)](#)

View how you have defined the key parameters on your NetEqualizer. It is viewable from several places in the NetEqualizer.

To view from RTR Menus

From the RTR Menus, *Click on -> Configuration*.

To view from the NetEqualizer Dashboard

From the Dashboard, *Click on ->[Show Configuration]*.



To view from the NetEqualizer Menus

From the Setup and Configuration Menu, *Click on -> Manage NetEqualizer -> Manage Configuration -> [Show Configuration]*.

Use this report to validate your settings. For example, you could verify that HOGMIN was set to your expected value (default is 12000 bytes), when reviewing connections that are getting penalized in the NetEqualizer Log. You can also see all your traffic limits, connection limits (P2P), priorities, and masks. The lines that start with TRACK show the subnets that are being tracked for the IP-level traffic reports. In the example below, you can see on Line #29 that I am tracking the 192.168.1.1/24 subnet.

CONFIGURATION			Update Data
Use this page to view the NetEqualizer Configuration. Search the table or sort the columns to analyze the NetEqualizer Configuration.			
25 records per page	Search:		
Line #	Parameter	Value	
25	HARD	198.145.20.140 212000 212000 16 1	
27	HARD	3.3.3.3/32 640000 640000 100000 1	
28	HARD	192.168.1.108 32 0 200000 1	
29	TRACK	192.168.1.1/24	
30	HARD	192.168.1.113 32 0 200000 1	
Line #	Parameter	Value	

Running Processes

[\(back\)](#)

View this report to check out what processes are running on your NetEqualizer. Use this report to see how much CPU and memory is being consumed by each process, as well as how long a process has been running. This is useful in troubleshooting efforts. From the RTR Menus, *Click on -> Running Processes*. In the example, I have searched on "neteq", and so the process table only shows the neteq process.

Running Processes													
RTR / Running Processes													
RUNNING PROCESSES													Update Data
Use this page to view running processes on the NetEqualizer. Search the table or sort the columns to analyze the processes running. The table column names match those of the Linux command, 'top'.													
25 records per page	Search: neteq												
Line #	PID	USER	PR	NI	VIRT	RES	SHR	S	CPU	MEM	TIME	COMMAND	
36	2401	root	20	0	30620	6692	2884	S	0.0	0.2	4:13.12	neteq	
Line #	PID	USER	PR	NI	VIRT	RES	SHR	S	CPU	MEM	TIME	COMMAND	
Showing 1 to 1 of 1 entries (filtered from 128 total entries) (processed in 0.0015549659729094 s)													
First Previous 1 Next Last													



Start/Stop RTR



([back](#))

The Real-Time Traffic graph on the Dashboard is always running, as it populates in real-time whether RTR is on or off. Similarly, the Active Connections Table always has data. Neither of these are impacted when you "Clear Reporting Data".

To populate reporting data for the Traffic History Graphs, you need to start RTR. This is also documented in the [Start RTR](#) section of Traffic History.

From the RTR Menus, Click on -> Start/Stop RTR, and then Click on -> [Start RTR] to turn on Traffic History reporting.

Once you have turned on traffic reporting, the System Status first line will be green and say "Traffic History Reporting is currently ON." (circled above).



Autostart RTR

([back](#))

If you would like RTR to automatically start upon reboot, so that it will be collecting data, you will need to add a command to your AutoStart file.

To start RTR upon reboot:

From the Maintenance and Reference Menu, **Click on -> Maintenance -> Edit Autostart File -> [Edit].**

Type in the following command at the bottom of the file:

`php /var/www/newgui/RTR/start-rtr.php`




Email Notifications

You can set-up an email account to receive alerts and notifications from the NetEqualizer. Notifications (email alerts) can be sent either immediately, or if you prefer less email, they can be batched up and sent out once per day.

This feature enables you to select from a list of standard events, and then sends you an email notification when the event is triggered. Standard events initially will include the following:

- Bandwidth Up or Bandwidth Down reaching 100% capacity
- IPv6 traffic exceeding 1%
- License violation, running over your license limit



Configure Email

([back](#))

In order to receive alerts, you need to define what email will be used for the alerts.

To set up an email address to send alerts to:

From the Management and Reporting menu, [Click on -> Manage Alerts -> \[Configure Email\]](#).

Fill out the eight (8) fields, which are used to populate each email notification, as follows: type in a [valid SMTP Server](#), [Port](#), [Authentication](#), [Secure Transfer](#), [Username](#), [Password](#), [From email address](#), and [From Name](#) that you would like used for the alert emails. You will need to scroll to see the second half of this window.

To cancel your changes, [Click on -> \[Reset\]](#). All values will be reset to the Default Values (if you have not yet set values), or to your previous values as stored in your configuration file (if you have already set values previously).

To save your configuration, [Click on -> \[Submit Changes\]](#). The screen at the right opens.

To verify that your settings will work, send a test email by typing in a [To email address](#), [Subject](#), and [Message](#), and then [Click on -> \[Send Test Email\]](#). Validate that the email is received by the SENT TO address, using the SENT FROM address with the SUBJECT LINE that you set.

Setup Email Notices

☒ Set up email alerts using your local SMTP server or Gmail.
See notes on both processes beside each form field.

☒ To set up an SMTP server that does not require authentication, leave Secure Transfer, Username, and Password blank, and put false in Authentication.

SMTP Server: SMTP:
GMail:

Port: SMTP: 25 or 587 commonly
GMail: 465

Authentication: SMTP: true or false
GMail: true

Secure Transfer: SMTP: ssl or tls
GMail: ssl

Username: Username/Password only necessary if
Authentication is "true".

Password:

Your request is complete.

The current settings are now:

SMTP set to [192.168.1.143]
Port set to [25]
Authentication set to [true]
Secure Transfer set to []
Username set to [yourname]
Password is set.
From set to [yours@email.am]
From name set to [yourpas]

To:

Subject:

Message:

[Send Test Email](#)



Setup Email Notices Parameters

Value	Description	Default Values
Remote SMTP server	Valid email server address.	smtp.yourdomain.com
Port	Port the email server uses to send/receive mail.	SMTP: 25 or 587 commonly GMail: 465
Authentication	True if auth is used. False if not.	SMTP: true or false GMail: true
Secure Transfer	Security (if used).	SMTP: ssl or tls GMail: ssl
Username	Username/Password only necessary if Authentication is "true".	Blank
Password	Same as Username.	Blank
From	Valid email address that you would like to send the alert FROM.	youremail@yourdomain.com
From Name	Name FROM for all emails.	Blank



Configure Alerts

[\(back\)](#)

Once you have set up your email server, you can select events to be notified on and the notification period, which is the amount of time between alert emails.

To select events and set your notification period:

Click on -> **Manage Alerts** -> **[Configure Alerts]**. The following screen opens.

Simply select any number of events (from 1 to all) that you want the NetEqualizer to report on and also type in a **valid TO email address**, and **Subject** to be used for all alerts.

You set a notification period by typing in the time in *seconds* (for example, "3600" is 1 hour) for how often you want an alert email sent for any events that have been triggered.

Select Events to Email On

☐ **Disk Full Warnings**
Recommended when running ntpd or Caching server.

☒ **IPv6 Alert**
Will let you know if your IPv6 traffic exceeds 1 megabit.

☐ **Link Capacity >95% Full**
Capacity is determined by your settings for TRUNK_UP and TRUNK_DOWN.

Amount in seconds for email notices if needed:

Email alerts to:

Subject:

▲ You must have already configured email and checked at least one event above or this email alert will not run.

Note: By default the routine that does the actual emailing is disabled. If you Select Email Alerts prior to Configuring Email, they will not run until you have configured your email server.

Redundancy and Failover

If you are concerned about passing traffic if your NetEqualizer goes down, either for scheduled maintenance or due to an unplanned failure (i.e. power outage, equipment failure, etc.), you might want to consider building in a plan for either: 1) full redundancy, or 2) failover.


We recommend full redundancy, as this ensures both that your network is passing traffic and that traffic shaping (equalizing) is being performed. Failover only ensures traffic is passing on your network; it does not offer traffic shaping.

Setting up Full Redundancy

NetEqualizer's bridge architecture fully supports network redundancy. If you would like to ensure that equalizing is in place at all times, you can put two NetEqualizers in your network in active/passive mode.

NetEqualizer is designed to fail "closed", meaning that network traffic will not pass through the unit if it goes down. You can set up redundancy across your two (2) NetEqualizers in one of two ways, as described below. Your network setup should look something like this:

NetEqualizer – Active/Passive Redundancy Configuration



You may have a firewall or smart switches and the actual configuration will of course be different than the diagram above but the same concept will apply.

Note: For either #1 or #2, we recommend that you run through a failure scenario manually and make sure that the failover takes over correctly on a dry run.

#1) Use Redundancy Capabilities on your Firewall or Switches

If your Firewall or Switches offer redundancy capabilities, you can configure those to treat one NetEqualizer as the active path, and the secondary NetEqualizer as your passive (failover) path.


Note: Under this scenario, redundancy is configured from the Firewall or Switch. Redundancy via STP is NOT enabled on the NetEqualizers.



#2) Use Spanning Tree Protocol (STP) on your NetEqualizers

NetEqualizer takes advantage of a mature feature already built into the Linux operating system called [Spanning Tree Protocol](#) (STP). Two NetEqualizers placed in parallel can be configured for a master/slave relationship where one server will back the other.

Using your NetEqualizers for redundancy involves running [Spanning Tree Protocol](#) (STP) on each. First, setup each NetEqualizer so they have their own (and different) management IPs. For example, the NetEqualizer comes with 192.168.1.143 as the default, so setup the second NetEqualizer with 192.168.1.144. Same netmask and gateway for this example.



Second, on each of the NetEqualizers, from the Maintenance and Reference Menu, *Click on -> Maintenance -> Edit Autostart File -> [Edit]*.

Type in the following two lines at the bottom of the file:

```
/sbin/brcctl stp br0 on  
/sbin/brcctl stp my on
```



Finally, for STP to take effect, reboot each of the NetEqualizer's. From the Maintenance and Reference Menu, *Click on -> Maintenance -> [Reboot NetEqualizer]*.


Note: Under this scenario, STP is enabled on the NetEqualizers and NOT on your Switches or Firewalls.

Failover

If you do not need full redundancy, but would like a failover solution to ensure that your network continues to pass traffic if your NetEqualizer goes down, you can configure a STP-capable switch to bypass the NetEqualizer. You can use your own switch or try our [third-party programmable bypass switch](#), which can be used where other devices are using STP without conflicting with them.

Note: In the case of the NetEqualizer going down, this solution does not maintain traffic shaping on your network.

NetEqualizer – Failover Configuration with Bypass Switch





Maintenance Tasks



Powering Off the NetEqualizer

If you ever need to shutdown the NetEqualizer (not just a reboot), here is a graceful way to do that:

Go to the Maintenance and Reference Menu, *Click on -> Maintenance -> [Run A Command]*. Type in: */sbin/shutdown -h now*
Then wait about 20 seconds and push the power button to shut it off completely.




Backing Up Your Configuration Settings

We recommend that you Save and Download Your Configuration. While we include a backup CF card with each NetEqualizer shipped, this does not contain your custom configuration settings.


Click on -> Setup and Configuration -> Manage NetEqualizer -> [Manage Configuration]. The following screen opens.

To back up your configuration, *Click on -> [Save Configuration]*. Save the **NetEq.cfg** file to a backup location.



Getting Software Updates for the NetEqualizer

We release Software Updates typically two (2) times per year. All customers that have current [NSS](#) contracts are eligible to receive Software Updates. If you are not sure if you are current on NSS, contact us at admin@apconnections.net or 303.997.1300 option 5.



To get the latest Software Update, you first need to generate a System Diagnostic file to email to Support. If you are not already on the Dashboard, from the NetEqualizer Navigation Menu, *Click on-> [Dashboard]*, then *Click on -> [Run Diagnostics] -> [Show Results]*.

Wait approximately forty (40) seconds for the diagnostic file to be fully generated. You can check if you have the whole file by looking at the end of the file, where it will say `#####done#####` if complete. Once you have the full file, you can save in text file format by *right clicking to Save Page As... filename.txt*, and then enter a valid text file filename. Attach the file to an email and send to support@apconnections.net.



Once your file is received, Support will review your diagnostic file, and let you know of any special instructions you will need to follow to upgrade. Otherwise, they will send you the standard instructions, along with a Quick Update Routine.

Note: Software Updates typically require a brief reboot of your NetEqualizer box. You may want to schedule this work during a planned downtime window.

Note: You will need Internet access to the NetEqualizer to perform Software Updates.

Software Updates are applied to the NetEqualizer in either one of two ways:


1) By a simple Quick Update Script (RECOMMENDED)

This runs against your current Compact Flash (CF) card. The first step of the Quick Update process creates backups of your key files, in case you need to restore back to your previous state.

2) By re-flashing or replacing your Compact Flash (CF) Card

This process should only be used if recommend by Support. Typically you would only use this process if it is recommended that you move to a new CF card.

Two CFs are provided with each NetEqualizer (one in the unit, and a backup CF in the Accessories Box). You can round-robin your CFs. Please be aware that when you replace the CF you replace every file and setting, as this is no different than putting a new hard drive into a system and removing the old one. In order to not lose your custom changes, you will need to retain the following files:



Save off your Custom Configuration (NetEq.cfg file)

You need to first save off your existing configuration file (NetEq.cfg) and then replace it after you swap out the CF card. Before replacing the CF, from the NetEqualizer Navigation Menu, *Click on -> [Dashboard] -> [Show Configuration]* and copy out the NetEq.cfg file's contents to a notepad file on your computer.



Save off your Shell Scripts

You should also create backups of any shell scripts that you have modified: such as settime.sh and crontab (used to sync date/time), and ntopdump.sh, ntop2mysql.sh, and ntopshort.sql (used to create a reporting data warehouse). For each file that you need to save, perform the following:

From the Maintenance and Reference menu, *Click on -> [Maintenance] -> [Edit Any Text File]* and copy out the file's contents to a notepad file on your computer.



Save off your Autostart File

From the Maintenance and Reference menu, *Click on -> [Maintenance] -> [Edit Autostart File]* and copy out the Autostart file's contents to a notepad file on your computer.



Record your Final Key

If you don't have your final key available then from the Maintenance and Reference menu, *Click on [Maintenance] -> [Run A Command]*

Type in: *cat /root/floppy/finalkey*
and record the key.



Once you install your re-flashed CF card, you will need to re-enter your Final Key, restore your NetEq.cfg file, copy back in your Autostart File, and copy back in any shell script files. The instructions that you receive from Support will walk you through this process in detail.

Note: Support will send detailed instructions to walk you through re-flashing or replacing the Compact Flash card.



Troubleshooting

This section of the User Guide contains some ideas to troubleshoot your NetEqualizer. For our full [Support Archive](#), please go to our NetEqualizer News blog site. You can also review our [Advanced Tuning](#) library, recommended for NetEqualizer power users. Finally, you can contact Support at support@apconnections.net or 303.997.1300 x102.

- 1) [My pools show no traffic](#)
- 2) [I cannot get traffic through the NetEqualizer](#)
- 3) [I would like to check my NetEqualizer Log](#)
- 4) [I would like to send a Diagnostic File to Support](#)
- 5) [I would like to run commands to troubleshoot my system](#)

My pools show no traffic

([back](#))

If you have set up pools, and your cables are reversed, you will see no traffic flowing through your pools. The remedy will be to swap your LAN and WAN cables. Review the diagram below to identify the WAN and LAN ports.

Once you swap cables, you will need to one of the following: 1) reboot your box, or 2) restart the equalizing process.




- 1) **To reboot your box:** From the Maintenance and Reference Menu, *Click on -> [Maintenance] -> [Reboot NetEqualizer]*.



- 2) **Stop/Start Equalizing:** Alternatively, you can just stop and restart the equalizing process. From the Dashboard, *Click on -> [Start/Stop Equalizing] -> [Start/Stop Equalizing]*, and stop equalizing by *Clicking on -> [Stop Equalizing]*, and restart by *Clicking on -> [Start Equalizing]*.

The easiest way to figure out the ports on the NetEqualizer 3000 and 4000 series is that if you are facing the back of the unit, the LAN port is on your right and the WAN port is on your left. Ports are not labeled on the Series 3000/4000.

Where to Install the NetEqualizer





I cannot get traffic through the NetEqualizer

([back](#))

Can you put a keyboard and monitor on the NetEqualizer and see if there are any errors showing up on screen?



From the Troubleshooting and Support menu, *Click on -> Troubleshooting*, then *Click on -> [Run a Command]*. Then run any of the following commands:

To see if your Compact Flash Card is corrupt:

cat /etc/arbdefault.conf

Make sure the top few lines exist in the NetEqualizer Configuration File. This ensures that your CF card is not corrupted (with a corrupt CF there most likely will be no config output). With older machines, this could happen because we used to run ntop on the CF. ntop on newer machines is now run in RAM.

To check for errors:

dmesg

Look at the end of the output. If you see any error messages, save them off and email to support@apconnections.net (errors like: can't read disk sector or out of memory or Duplicate IP).

Ideally, you should not reboot the NetEqualizer, as that will clear out the NetEqualizer Log File (although we do save the previous version to a .bak file).

I would like to check my NetEqualizer Log

([back](#))

The NetEqualizer Log File contains a record of the actions of the NetEqualizer. Here we highlight how to view the file. For a detailed description of how to read the NetEqualizer Log File, read our [Show the NetEqualizer Log File](#) section of the User Guide.



To view the NetEqualizer Log:

From the Management and Reporting Menu, *Click on -> View Current Activity -> [View NetEqualizer Log]*.




To see the ENTIRE NetEqualizer Log:

Go to the Maintenance and Reference Menu, *Click on -> Maintenance -> [Run a Command]*.

Type in: *cat /tmp/arbblog*

To see the entire PREVIOUS Log:

Type in: *cat /tmp/arbblog.bak*



I would like to send a Diagnostic File to Support

([back](#))

You can generate a System Diagnostic File to email to Support. From the Troubleshooting and Support menu, *Click on -> Troubleshooting*, then *Click on -> [Run Diagnostics] -> [Run Diagnostics] -> [Show Results]*.



Wait approximately forty (40) seconds for the diagnostic file to be fully generated. You can check if you have the whole file by looking at the end of the file, where it will say #####done...##### if complete. Once you have the full file, you can save in text file format by **right clicking to Save Page As... filename.txt**, and then enter a valid text file filename. Attach the file to an email and send to support@apconnections.net, along with a description of the problem you are encountering. Once your file is received, Support will review your diagnostic file, and contact you to help troubleshoot your system.



I would like to run commands to troubleshoot my system

[\(back\)](#)

We offer a command line interface to help you to troubleshoot your system. From the Troubleshooting and Support menu, **Click on -> Troubleshooting**, then **Click on -> [Run a Command]**. The following window opens.

You can either enter a Command in the text box, or scroll down to see Common Commands (see List of Common Commands Table below) and run any of them by clicking on the associated button.

Note: You cannot run any command that requires user input.

List of Common Commands

Command	Description
Date	To get the current date and time.
shutdown -h now	To shutdown the NetEqualizer unit gracefully.
Uptime	To get the current uptime and load average.
/art/count	To get a dump of the active connection report.
df -h	To get the current amount of disk space left on the system.
cat /root/rebooted	To print the log of reboots.
ps ax	To get a list of the current running processes.
/art/toponce	To print output of the program top one time.
/sbin/ifconfig	To get information about the network interfaces.
ethtool eth0;ethtool eth1	To get the current connection speed of the network interfaces.
ip route show	To show the current network route information.
cat /proc/cpuinfo	To get CPU information about the unit
cat /proc/meminfo	To get the unit's installed memory and memory usage information.
Dmesg	To get the low-level OS warning messages Note: some warnings are normal.



Frequently Asked Questions (FAQs)

This section of the User Guide contains a sampling of tips, based on FAQs (frequently asked questions). For our full [Support Archive](#), please go to our NetEqualizer News blog site. You can also review our [Advanced Tuning](#) library, recommended for NetEqualizer power users.

- 1) [How to Name your NetEqualizers](#)
- 2) [How to Enable Speed Tests](#)
- 3) [Check your Maximum Bandwidth with our Speed Log Tool](#)
- 4) [How to Test Bandwidth Limiting Rules](#)
- 5) [How to Tune Bandwidth Limit Precision](#)
- 6) [How to Monitor Bandwidth Hogs](#)

How to Name Your NetEqualizers

[\(back\)](#)

If you have more than one NetEqualizer, you may want to give each an individualized name, in order to keep them straight while you are working with them. The NAME parameter is displayed prominently on the Dashboard, so that you can easily tell which NetEqualizer you are administering.



From the Maintenance and Reference menu, *Click on ->Maintenance-> [Edit Any Text File]*.

To open the config file, type in: */etc/arbdefault.conf*

Go to the "NAME=" parameter line in the file and change whatever is listed after the equal sign to what you want to call this NetEqualizer. For example, if you decided to name by location, you might have a *NAME=BoulderNetEQ*.

Click on -> [Post Changes] to save or *Click on -> [Reset]* to cancel your changes.

How to Enable Speed Tests

[\(back\)](#)

In order to ensure that speed test sites are not equalized, you can give them priority treatment on your network. This is done through setting up each speed test site as [Priority Traffic](#).



To set up Priority Traffic:

From the Setup and Configuration Menu, *Click on -> Manage Priority -> [Manage Traffic with Priority over Equalizing]*

Check your Maximum Bandwidth with our Speed Log Tool

[\(back\)](#)

If you want to verify that you are getting the bandwidth expected from your provider, you can run these commands. This is useful during troubleshooting, if you feel that you are hitting RATIO too soon or not hitting RATIO often enough. It could be that you are not



getting your expected bandwidth - you might be getting more or less.



From the Maintenance and Reference menu, *Click on -> Maintenance -> [Run a Command]*. To start up the speedtest logging routine:

Type in: *nohup /art/speedtest 1 24 1 3600 1>/dev/null 2>&1 &*

Blank screen will come up when command is processed.

Click on Back Arrow to return to Run a Command and then Close the window.

Once you do this and wait at least an hour (we recommend that you run this during peak), you will be able to run the following:



From the Maintenance and Reference menu, *Click on -> Maintenance -> [Run a Command]*

Type in: *cat /tmp/speedlog*

You will see something like:

Bandwidth Speed Test Warning Thu May 5 08:01:56 MDT 2011, expecting a min speed of 102400, peak reached only 46098

This is dependent on what you have set for Bandwidth Up (TRUNK_UP) and Bandwidth Down (TRUNK_DOWN). Sometimes it will have a warning and sometimes the message will be different. In this example, maximum bandwidth should have been 102400 and only 46098 was achieved (45% of expected bandwidth). If this test was run during peak hours, either your pipe is not saturated or you would contact your bandwidth provider to find out why you are not able to access all of your promised bandwidth.

Click on Back Arrow to return to Run a Command and then Close the window.



To stop the process you run the following:

From the Maintenance and Reference menu, *Click on -> Maintenance -> [Run a Command]*

Type in: *ps ax*

Look for a line like: *20167 ? SN 0:00 /usr/bin/perl -I -w /art/speedtest 1 24 1 3600*

Then you need to stop the process. Type in: *kill 20167* where 20167 is the process number of the line you found.

If you want to verify that the process is NOT still running, run *ps ax* again and look for it.

How to Test Bandwidth Limiting Rules

[\(back\)](#)

Because NetEqualizer *adjusts to traffic over several seconds*, attempts to set limits on short traffic bursts will have limited affect. NetEqualizer is designed to allow short bursts of traffic through. For most users, allowing these bursts is the desired effect. Short bursts have relatively little effect on overall traffic and should be given priority.

When you do your initial testing on Bandwidth Limits ([bandwidth limiting rules](#)), use file transfers that persist for more than 15 seconds to allow NetEqualizer to come to a steady rate of data transfer.



How to Tune Bandwidth Limit Precision

[\(back\)](#)

Note: This assumes that you are NOT already using [bursting](#) on your bandwidth limits.

NetEqualizer is designed to do a good job over time (five minute averages) of keeping bandwidth within specification. However, the NetEqualizer will allow some bursts through. As NetEqualizer takes a few seconds to adjust to changing traffic situation, if you are testing with one or two large downloads, the bursts will be more pronounced than traffic on a busy network.

Some tuning may be required to override the background shaping rules (which may be more restrictive than your desired limits). On higher speed networks, the default tuning in NetEqualizer may reduce traffic rates more than an acceptable margin of error (acceptable error margin to us is 10 percent; we do not claim to have billable quality rate limiting).

We recommend reducing the size of your [PENALTY_UNIT](#) to compensate. Click on the link to go to the PENALTY_UNIT section of this document, where we offer detailed recommendations on tuning PENALTY_UNIT.

How to Monitor Bandwidth Hogs

[\(back\)](#)

Below is a step-by-step process to create a new script that will just show you network connections over Hog Minimum (HOGMIN). This does not mean that they are currently being penalized because the script doesn't know if you are in a penalty situation or not. If you are using RATIO of your Bandwidth Up (TRUNK_UP) or Bandwidth Down (TRUNK_DOWN), then they are being penalized.



Click on -> Maintenance -> [Run a Command]

Type in *touch /art/showhogs*

A blank screen will come up when command is processed.

Click on Back Arrow to return to Run a Command.

Now run the following command.

Type in: *chmod a+x /art/showhogs*

A blank screen will come up when command is processed.

Click on Back Arrow to return to Run a Command and then Close the window.



Click on -> Maintenance -> [Edit Any Text File]

Type in */art/showhogs* to open this file.

Copy and paste the following contents into that edit window and then save it.

Note: You may need to hit <return> to force an EOL character after each line.

```
#!/usr/bin/perl -w
@y=`/art/BROWSE_CONFIG PARAM HOGMIN|head -n 1`;
foreach $line1 ( @y) {
  chomp($line1);
  @s1=split(" ", $line1);
  $hogmin=$s1[1];
}
print "HOGMIN is currently set to $hogmin. ";
```




```
print "The connections below are over HOGMIN...\n";
print "SRCP DSTP Wavg Avg IP1 IP2 Ptcl Port Pool\n";
@x=`/sbin/brctl getbrain my 0|grep -v Wavg`;
foreach $line ( @x) {
  chomp($line);
  @specials=split(" ", $line);
  if ( $specials[5] > $hogmin)
  {
    print "$specials[1] $specials[2] $specials[3] $specials[4]
    $specials[5] $specials[6] $specials[7] $specials[8] $specials[9]\n";
  }
}
```

Click on -> *[Post Changes]* to save changes. On the next screen, you will see "Your request is complete".



Click on -> *Maintenance* -> *[Run a Command]*

Type in */art/showhogs*

All connections over Hog Minimum will be displayed when command is processed.





Appendix 1 - Equalizing Parameters, Units, & Defaults

Key Equalizing Parameters				
Parameter	Unit	Default	What you can set to...	Tips
Ratio (RATIO)	Percentage	85	<i>1-100</i>	Default of 85% works for most networks. To be more aggressive, use 70 or 75%. To be less aggressive, raise it to 90%.
Bandwidth Up (TRUNK_UP)	Mbps <i>Preferences-> Configuration Units = Mbps</i>	15360 Mbps	<i>Size of your outbound network pipe. Traffic from the LAN to the WAN (Internet).</i>	Always set to your License Level or lower. License Level is shown as ## in the message "The system is authorized to pass ## Mbps."
Bandwidth Down (TRUNK_DOWN)	Mbps	15360 Mbps	<i>Size of your inbound network pipe. Traffic from the WAN (Internet) to the LAN.</i>	Always set to your License Level or lower. License Level is shown as ## in the message "The system is authorized to pass ## Mbps."
Hog Minimum (HOGMIN)	Mbps	1 (1000 Kbps)	<i>For networks of size: < 50Mbps .5 ≥ 50Mbps & < 200Mbps .75 ≥ 200Mbps & < 1Gbps 1.0 ≥ 1Gbps 2.0</i>	Default of 0.5Mbps (500 kbps) is set so that VoIP, email, web-based applications, web surfing, and chat is below Hog Minimum.
Equalizing Rules (DEFAULT_RULES)	On/Off toggle	On	<i>Leave at Default of "on".</i>	Must be "On" for Equalizing to kick in. Turn off during installation if you want to run throughput tests.
Additional Equalizing Parameters				
Maximum Penalty (MAX_PENALTY)	Hundredths of seconds	140	<i>Rarely changed from Default value.</i>	Should be greater than PENALTY UNIT and less than 200.
Penalty Unit (PENALTY_UNIT)	N/A	5	<i>For networks of size: ≥ 10Mbps to < 50Mbps 2 - 3 ≥ 50 Mbps 1</i>	PENALTY_UNIT is the minimum penalty that will be inflicted on a packet when a penalty is set up on an IP address. Values for this variable are integers in the range 1 - 100, with 1 being the least restrictive.
Connection Tracking Table Size (BRAIN_SIZE)	Number of Connections (IP pairs) to track in one (1) second.	10000	<i>For networks of size: < 100Mbps 20000 ≥ 100Mbps & < 1Gbps 30000 ≥ 1Gbps 40000</i>	How many connections (IP pairs) to keep track of at one time during any given second.
Inactive Tics (INACTIVE_TICS)	Hundredths of seconds	200	<i>100-800 Rarely changed from Default value.</i>	1 (100) to 8 (800) seconds. Time something tracked in Connection Tracking Table without activity.



Appendix 2 - Setting/Forcing LAN Speeds and Duplex

Occasionally you need to manually set LAN Port Speed and Duplex in order for the NetEqualizer to operate at the expected port speeds and in the correct duplex mode.

The NetEqualizer LAN ports auto-negotiate 95% of the time. However, the NetEqualizer may need to be manually set to work with some Routers or Switches. Symptoms that you need to set your LAN Port Speed and Duplex are that you are having collisions and/or dropping packets. Both these symptoms will make your network throughput less than expected.

Note: Although dropped packets are not a good thing, if you are seeing less than 1/10 of a percent (< 0.1%) of the total packets transmitted it will have no adverse effect on your network. If it starts to approach 1 percent (1%), you should follow these instructions.

To Check Your Current Port Speeds

From the Maintenance and Reference Menu, *Click on -> Maintenance -> [Run a Command]* to run the following commands:

To see if your ports are dropping packets or having collisions, run: */sbin/ifconfig*
To see what your ports' details are run the following commands: */usr/sbin/ethtool eth0* and */usr/sbin/ethtool eth1*

You can also run these commands by clicking on "To get the current connection speed of the network interfaces:" button.

To Set Your Port Speed and Duplex Mode

From the Maintenance and Reference Menu, *Click on -> Maintenance -> [Run a Command]* to run the following command:

```
/usr/sbin/ethtool -s DEVNAME \ [ speed 10|100|1000 ] \ [ duplex half|full ] \ [ autoneg on|off ]
```

Here are some examples to force a WAN interface to a certain speed and full duplex:

```
/usr/sbin/ethtool -s eth0 speed 1000 duplex full autoneg off  
/usr/sbin/ethtool -s eth1 speed 1000 duplex full autoneg off
```

eth0 should be WAN and pointed towards the Internet. eth1 should be LAN and going into your internal network.

To Put Your Port Speed and Duplex Mode in Auto Startup File (recommended)

We recommend that you add these commands to your Autostart file, after you have verified that the change is set the way that you would expect and works on your network.

From the web GUI:

From the Maintenance and Reference Menu, Click on -> Maintenance -> [Edit Autostart File] -> [Edit].



Insert your changes at the bottom of the Autostart File, and then *Click on the [Post Changes]* to save or *Click on the [Reset]* to discard your changes.

From the console or SSH:

If you would like to put these commands in the Autostart File, you can put them into */art/autostart* by editing the file from the console or SSH. Login as "*root*", using the default password (unless you changed it previously).

From vi or nano:

You can also use nano or vi to edit the */art/autostart* file. Start your editor by typing in the following:

```
nano -w /art/autostart
```

The command is formatted as follows:

```
ethtool -s DEVNAME \  
    [ speed 10|100|1000 ] \  
    [ duplex half|full ]  \  
    [ autoneg on|off ]
```

At the very bottom of */art/autostart*, put in your new command lines, such as:

```
/usr/sbin/ethtool -s eth0 speed 1000 duplex full autoneg off  
/usr/sbin/ethtool -s eth1 speed 1000 duplex full autoneg off
```

Use the backspace and delete and arrow keys just like in Notepad. Save with Ctrl-o and Enter and exit with Ctrl-x. There is a menu at the bottom of nano that shows these commands.



Appendix 3 - Packet Capturing for taps such as CALEA

NetEqualizer is a CALEA Probe

The NetEqualizer acts as a CALEA Probe via packet capturing & forwarding. We provide a network probe with the following capabilities:

- It will allow an ISP or other operator to comply with a basic warrant for information about a user by capturing and sending IP communications in real time to a third party.
- Communication may be captured by headers or headers and content.
- We provide basic descriptive tags identifying headers, data, and time stamps, along with HEX or ASCII representation of content data.

Note: The NetEqualizer does not do any analysis of the data. We are providing the probe function ONLY.

CALEA Compliance

As best we can tell at this time, there is no one government agency that can fully declare our technology CALEA compliant. However, we do pledge to work with our customers should they be faced with a warrant for information to adjust and even customize our solution; however additional consulting fees may apply.

Although the law (see [CALEA](#) sections 103 and 107(a)(2)) is fairly specific on **what** needs to be done, the **how** is not addressed to any level of detail to which we can engineer our solution.

We believe that the law and specifications on "how" to deliver to a law enforcement agency are somewhat ambiguous. The FBI has created some detailed specifications, but the reality is that there are some 40,000 law enforcement agencies, and they are each given autonomy on how they receive data. We do provide samples (see below) on how to receive NetEqualizer-captured data on a third party server, but are unable to guarantee definite compliance with any specific agency.

Many people are following the **ATIS specification** that was put forth by the FBI, and we *have read and attempted to comply with the probe portion of that specification*. But, the reality is that there is no one agency given the authority to test a solution and bless it as compliant.

So, if faced with a warrant for information, the law enforcement agency in charge may indeed want something in a slightly different format. If this is the case, contact our Support Team at support@apconnections.net or 303.997.1300 x102 for help in complying. Please note that as the CALEA module is not covered under [NSS](#), consulting charges may apply.

For additional information on CALEA, go to: <http://transition.fcc.gov/pshs/services/calea/>.

The NetEqualizer is set up as a CALEA Probe in two steps, which must be executed in the order below:

1. Setting up the Receiver for the tap.
2. Setting up the NetEqualizer to capture packets.



Step #1: Setting up the Receiver for the Tap

Install netcat (nc) onto a computer. Netcat can be installed on Ubuntu or Debian with:

```
apt-get update  
apt-get install netcat
```

Netcat can also be installed on Windows by finding the Windows version on the Internet and installing it.

Set up the port to listen on:

On the receiving computer, run the command line of: `nc -l -p XXXXX`

where XXXXX is the port you want to listen on, and that you setup on the NetEqualizer to send on.

Pipe results to a File (optional):

Netcat can be piped to a file using the > and | like any other command.

Step #2: Setting up the NetEqualizer to Capture Packets

To set-up packet capturing on the NetEqualizer:

From the Management and Reporting Menu, *Click on -> Manage Packet Capture -> [Manage Packet Capture]*. Fill in the fields and scroll down to *Click on -> [Start Packet Capture]*.

As packet capturing takes up both memory and CPU on the NetEqualizer, we recommend that you turn it off when you no longer need to capture data.

To stop packet capturing on the NetEqualizer:

From the Management and Reporting Menu, *Click on -> Manage Packet Capture -> [Manage Packet Capture]*. Scroll down to *Click on -> [Stop Packet Capture]*.

Note: You must have already started the service or capture routine on the receiving server (Step #1 above).

Note: Starting multiple packet captures will require multiple stops to turn packet capture off.



Appendix 4 - Tuning Parameters for a Large Number of subnet-ranged Limits, Pools, & Masks

Tuning for a Large Number of subnet-ranged (≥ 32 subnet ranges) Connection Limits, Hard Limits, Masks, Pools, and VLANs

The NetEqualizer currently simplifies your configuration set-up by enabling you to enter Hard Limits, Connection Limits, Masks, and Pools as "subnet ranges" (i.e. HARD x.x.x.x/24 or /16), instead of as individual rules.

In [Software Update 5.4](#) we added a tuning parameter that offers you further flexibility in your use of subnet-ranged Connection Limits, Hard Limits, Masks, and Pools. This parameter is used to increase the number of subnet-ranged definitions possible, from the default of 32 up to a maximum of 128.

This parameter does not need to be changed if you use a combined ≤ 32 subnet ranges in total across Pools, Hard Limits, Masks, and Connection Limits.

Note: You also need to consider the number of VLANs that you have set up in your total. If you have > 32 VLANs, or your number of VLANs + subnet-ranged items > 32 , you should set this parameter.

Set_Max_Table_Size Parameter

The tuning parameter, **set_max_table_size**, is used to increase the number of subnet-range definitions possible, from the default of 32 up to a maximum of 128.

To set this parameter, go to the Maintenance and Reference Menu, *Click on -> Maintenance -> [Run a Command]*.

Type in: */sbin/brctl set_max_table_size my XX*

Replace "XX" with the number of subnet-ranged entities that you want to support, such as *'/sbin/brctl set_max_table_size my 45'* to support 45.

When you save your changes, in order for them to take effect upon your configuration, you must stop and then restart equalizing. *Click on -> Maintenance -> [Start/Stop Equalizing] -> [Stop Equalizing]*. Then restart equalizing by *Clicking on -> [Start/Stop Equalizing] -> [Start Equalizing]*.

To persist Set_Max_Table_Size upon reboot

From the Maintenance and Reference menu, *Click on -> Maintenance -> [Edit Autostart File]* and add the following to the end of the Autostart file:

Type in: */sbin/brctl set_max_table_size my XX*

You also need the following commands ONCE at the end of the Autostart file:

```
/etc/init.d/arbitrate stop  
sleep 10  
/etc/init.d/arbitrate start
```

Note: NetEqualizer can support a maximum of 128 subnet-ranged entities. All subnet-ranged Connection Limits + Hard Limits + Masks + Pools + number of VLANs must be ≤ 128 . As always, the NetEqualizer continues to support an UNLIMITED number of individual Connection Limits, Hard Limits, Masks, and Pools (i.e. x.x.x.x/32).



Appendix 5 - Syncing NetEqualizer Date/Time

Keep the NetEqualizer Date/Time synchronized with either your own NTP Time Server (3A) or Internet Time Servers (3B)

Over time the NetEqualizer time will drift, like any server. You can configure the NetEqualizer to use your own NTP (Network Time Protocol) Time Server or an Internet Time Server.

Follow the instructions below, using either #3A (Sync to your NTP Time Server) or #3B (Sync to Internet Time Servers), but NOT both. If you prefer to use a command line interface, you can edit the /root/settime.sh and /root/crontab files from the command line or SSH with a text editor.

Note: You must have NetEqualizer setup so that it can access the Internet for this to function.

*Note: You must make sure that ntop is not running, or **stop ntop**, before changing the time on your NetEqualizer. Otherwise ntop will not function to create graphs. You can check this on the NetEqualizer Dashboard. The ntop process is running if the ntop button is ON (GREEN); ntop is off if the button is OFF (RED). If the ntop process is ON (GREEN), Click on -> View Historical Reports -> [Start/Stop ntop], and stop ntop.*

Set Time Zone

You should also set your time zone for the NetEqualizer.

Note: This does NOT need to be added to the autostart file or the cron job.

- 1) Login to your NetEqualizer console with the default root account credentials. Log into the unit with SSH or via a keyboard and monitor directly connected to the NetEqualizer.
- 2) At the command prompt, type in:
`dpkg-reconfigure tzdata`
and then follow the on-screen instructions.
- 3) Once you are done, you can type Ctrl-D to logout of the shell.

Keep the NetEqualizer Date/Time synchronized with either your own NTP Time Server (3A) or Internet Time Servers (3B)

1. Click on -> Maintenance -> [Run a Command].
Type in `touch /root/settime.sh; chmod a+x /root/settime.sh`
Blank screen will come up when command is processed.
Click on Back Arrow to return to Run a Command and then Close the window.
2. Click on -> Maintenance -> [Edit Any Text File].
Type in `/root/settime.sh` to open this file. The file is initially BLANK.
3. Type the following lines into the settime.sh file.

3A. Sync to your NTP Time Server...

`/usr/sbin/ntpdate 10.0.0.1 <return>` (replace 10.0.0.1 with your ntp time server IP)
`/sbin/hwclock --localtime --systohc <return>` (type in **2 DASHES** before each parameter)

3B. Sync to Internet Time Servers...



```
/usr/sbin/ntpdate-debian <return>
```

```
/sbin/hwclock --localtime --systohc <return> (type in 2 DASHES before each parameter)
```

You MUST make sure to hit **<return>** to force an EOL character after the line. *Click on -> [Post Changes]* to save changes. On the next screen, you will see "Your request is complete".

4. *Click on -> Maintenance -> [Run a Command].*

Type in *touch /root/crontab*

Blank screen will come up when command is processed.

Click on Back Arrow to return to Run a Command and then close the window.

5. *Click on -> Maintenance -> [Edit Any Text File].*

Type in */root/crontab* to open this file. Replace file contents with the following:

Type in ** /5 * * * * /root/settime.sh <return>*

You MUST make sure to hit **<return>** to force an EOL character after the line. *Click on -> [Post Changes]* to save changes. On the next screen, you will see "Your request is complete".

6. *Click on -> Maintenance -> [Run a Command].*

Type in *crontab /root/crontab*

Blank screen will come up when command is processed.

Click on Back Arrow to return to Run a Command and then close the window.

7. **To have the sync persist upon reboot, you must add this to the AutoStart file as well.**

Click on -> Maintenance -> [Edit Autostart File].

Type in *crontab /root/crontab* on a new line right ABOVE the line that says *thedata='date'*.

Click on -> [Post Changes] to save changes. On the next screen, you will see "Your request is complete".

To validate Date/Time is set correctly

The easiest way is to set your Date/Time to something in the past (Past Date/Time). From the Setup and Configuration menu, *Click on -> Manage NetEqualizer -> [Configure Date/Time]*. Enter a Past Date/Time and *Click on -> [Submit]* to save. Go to the NetEqualizer Dashboard, scroll down if needed, and refresh your screen. Make sure the Date/Time is set to your Past Date/Time. Now wait 5 or more minutes (the cron job runs every 5 minutes), go back to the Dashboard, and refresh your screen again. Make sure that the Date/Time is now Current. As the cron job will run every 5 minutes, your Date/Time will no longer drift.



Appendix 6 - Firewalling the NetEqualizer

This appendix is for customers that need to install the NetEqualizer outside of their firewall, on the public side of their Internet pipe.

Firewall rules are provided to prohibit unauthorized users from accessing the NetEqualizer IP and thus SSH access and the NetEqualizer Web GUI screen.

WARNING: The firewall rules can lock you out from the unit if you give it the wrong rule or do not know what you are doing. Be sure that you have access to the unit's console before testing rules if you are unsure. At the console you could login and clear firewall rules typically with: *iptables -F* or if using ebtables, *ebtables -F* at the command prompt and clear most rules a user would be testing.

Review Firewall Samples

The Firewall Samples (listed below) and available via the GUI, should be viewed before setting up a Firewall on your NetEqualizer.

From the Management and Reporting Menu, *Click on -> Manage Firewall Settings -> [Sample Firewall Rules] -> [Show Rules]*.

Note: The NetEqualizer has a bridging firewall installed so the FORWARD table is used for rules affecting things going "through" the unit. INPUT and OUTPUT tables are used to protect the unit itself.

Sample file	Description
<i>firewallprotectneteq.txt</i>	How to protect the NetEqualizer unit from unauthorized access.
<i>howtodroppacketsfromthisorthat.txt</i>	How to drop packets by IP or PORT or MAC address going through the NetEqualizer.
<i>redirectfw.txt</i>	A sample of how to create a capture portal using IPtables and the macs.allow file.

To set up the NetEqualizer Firewall:

From the Management and Reporting Menu, *Click on -> Manage Firewall Settings -> [Configure Firewall] -> [Edit Firewall Rules File] -> [Edit]*.

Put in the rules that you need, by copying and pasting from the Firewall Samples and then modifying for your environment. *Click on -> [Post Changes]* to save or *[Reset]* to cancel.

To view NetEqualizer Firewall Settings:

From the Management and Reporting Menu, *Click on -> Manage Firewall Settings -> [View Firewall Settings]*.

To start or stop the NetEqualizer Firewall:

After you have configured your Firewall, or made changes to it, you will need to start/restart it. From the Management and Reporting Menu, *Click on -> Manage Firewall Settings -> [Start/Stop Firewall]*. Then *Click on -> [Start/Restart Firewall]* to start.

If you decide for any reason to stop your firewall, *Click on -> [Stop Firewall]* to stop.

