

NETROPY® NETWORK EMULATOR

USER'S GUIDE

Firmware Version 3.0



APPOSITE
— TECHNOLOGIES

Netropy® Network Emulator User's Guide

Revision 3c, August 2017

Apposite® Technologies LLC

17835 Ventura Blvd., Suite 211

Los Angeles, CA 91316 USA

1.310.477.9955

info@apposite-tech.com

www.apposite-tech.com

Copyright ©2005-2018 Apposite Technologies LLC. All rights reserved. Apposite, Linktropy, and Netropy are registered trademarks of Apposite Technologies LLC. The Apposite logo and "WAN Emulation Made Easy" are trademarks of Apposite Technologies LLC. Android and the Google Play logo are trademarks of Google Inc.

Part No. DOC-UG-NNE-3c

CONTENTS

1 Overview.....	5
1.1 Netropy Configuration.....	5
1.2 Netropy Operation.....	5
1.3 Netropy Models and Firmware Versions	6
2 Installation and Set-Up.....	7
2.1 Preparation	7
2.2 Hardware Installation	7
2.3 IP Address Configuration via the Netropy GUI.....	8
2.4 IP Address Configuration via the Serial Console	8
2.5 Network Installation	9
3 Configuration	10
3.1 Emulation Engine	10
3.2 Paths	11
3.3 Classifiers	11
3.4 GUI and CLI.....	11
4 Emulation Engine	12
4.1 Overview of Emulation Engines.....	12
4.2 Engine Locking	14
4.3 Engine Self-Monitoring.....	14
5 Paths	15
5.1 Overview of Paths	15
5.2 Single-Link and Per-Client Emulation Paths	16
5.2.1 Single-Link Mode	16
5.2.2 Per-Client Emulation Mode	16
5.3 Path Types	17
5.4 Configuring Paths	18
5.5 Adding and Deleting Paths.....	18
6 Path Parameters: Configuring Emulation Conditions	19

6.1 Configuring Path Parameters	19
6.2 WAN Access Parameters.....	20
6.2.1 Link Type.....	21
6.2.2 Bandwidth.....	21
6.2.3 Background Utilization.....	22
6.2.3.1 Random Background Traffic	22
6.2.3.2 Packet Replay Using PCAP Files	23
6.2.4 Queue Limit	25
6.2.5 Queuing Strategy	26
6.2.6 MTU Limit	27
6.2.7 Frame Overhead	28
6.3 WAN Parameters	29
6.3.1 Delay	30
6.3.2 Loss	32
6.3.3 Network Outage	33
6.3.4 Corruption	34
6.3.5 Reordering	34
6.3.6 Duplication.....	35
7 Packet Classifiers.....	36
7.1 Overview of Packet Classification.....	36
7.2 Classification Methods	37
7.2.1 Classification Off.....	38
7.2.2 IP Address Classification	38
7.2.3 IPv6 Address Classification	39
7.2.4 VLAN Classification.....	40
7.2.5 MPLS Label Classification.....	41
7.2.6 MAC Address Classification	42
7.2.7 TCP/UDP Port Classification.....	43
7.2.8 Combination Classification	45
7.3 Rule Order	48
8 Recordings	49
8.1 Overview of Recordings.....	49
8.2 Managing Recordings.....	50
8.3 Recording Playback.....	51

8.4 Recording File Format.....	53
8.5 Example Recording File.....	54
9 Monitoring & Statistics Download.....	55
9.1 Graphs.....	55
9.2 Statistics.....	56
9.2.1 Statistics Download.....	58
10 Administration.....	60
10.1 Network Settings.....	60
10.2 Users.....	61
10.3 Bridge/Route.....	62
10.3.1 Bridging.....	62
10.3.2 Routing.....	63
10.4 Ethernet Settings.....	64
10.5 Recordings.....	65
10.6 Packet Captures.....	65
10.7 Save and Restore Engine Configurations.....	65
10.8 Management Network Status.....	66
10.9 Engine Log.....	66
10.10 Firmware.....	67
10.11 License Key.....	68
10.12 Date and Time.....	68
11 Command Line Interface.....	69
11.1 CLI Help.....	69
11.2 Navigation within the CLI.....	70
11.3 CLI Top Level Commands.....	71
11.3.1 LDAP Management Commands.....	75
11.3.2 Capture Command.....	76
11.4 CLI Engine Mode Commands.....	77
11.5 CLI Path Mode Commands.....	79
11.5.1 CLI Set Emulation Parameters.....	80
12 Security.....	85
12.1 Users and Passwords.....	85
12.2 Engine Locking.....	85
12.3 Recovering from a Lost Admin Password.....	85
12.4 SSL.....	86

12.5 SSH.....	86
12.6 LDAP	86
13 Apposite Support.....	88
13.1 Registration	88
13.2 Customer Support.....	88
Appendix A: Apposite End User License Agreement	90
Appendix B: Third Party Licenses	93

1 OVERVIEW

Thank you for purchasing the Apposite Technologies Netropy network emulator. This *User's Guide* describes the installation, configuration, and operation of the Netropy product. A companion *Hardware Guide* describes the Netropy hardware for each specific model. A separate *Quick Start Guide* provides a walk-through for first time configuration.

The Netropy network emulator attaches to an Ethernet network and simulates the bandwidth, delay, loss and other conditions of the wide-area network to test application performance.

1.1 Netropy Configuration

The Netropy network emulator is usually configured through the browser-based Netropy GUI (Graphical User Interface). The GUI is accessible through a dedicated management port from any PC or other device with a standard web browser using HTTP or HTTPS.

In addition to the GUI, the Netropy network emulator includes a command line interface (CLI) that can be accessed via a serial console port or over the network using Telnet or SSH. The CLI can be used to set the IP address of the management interface if the GUI is not accessible over the network, and to modify emulated link conditions for integration with scripting and test automation tools.

1.2 Netropy Operation

Configuration and operation of the Netropy network emulator via the browser-based GUI requires only a few simple steps:

1. **Open the GUI**
Connect to the Netropy unit through the dedicated management port using a standard web browser.
2. **Select the Emulation Engine**
Depending on hardware model, the Netropy unit will include between one and four separate Emulation Engines. Each Emulation Engine acts as an independent network emulation system connecting a pair of Ethernet ports.

3. **Add Paths**
Create separate WAN paths to carry packets between the two Ethernet ports in each Engine. Each Emulation Engine can simulate up to 15 separate paths.
4. **Configure WAN conditions for each path**
Configure each path with the bandwidth, delay, loss, and other WAN conditions.
5. **Classify Packets**
Assign packets to the paths by IP source and destination address range, VLAN, MAC address, TCP/UDP port numbers, MPLS label, or other packet identifier.
6. **Start the Emulation Engine**
Turn on emulation to begin testing.
7. **Monitor traffic**
View the graphs and link statistics to monitor application performance.
8. **Change configuration**
The configuration can be changed on the fly by adding or deleting paths, modifying path conditions, or updating the packet classification rules.

1.3 Netropy Models and Firmware Versions

The six current Netropy models, N61, N91, 10G1, 10G2, 10G4 and 40G, can all run the firmware described in this *User's Guide*. The various models offer identical functionality and differ only in capacity and number and type of network interfaces. The Netropy N90 model has been superseded by the N91 but can run the firmware described in this manual. Earlier Netropy models, the N60, N80, and 10G, are limited to running firmware versions up to v2.2.

Depending on when the Netropy unit was produced and if the firmware on the unit was updated, the unit may be running an earlier version of the firmware than described in this *User's Guide*. Please consult the appropriate *User's Guide* for the version of firmware running on your unit.

2 INSTALLATION AND SET-UP

To configure and operate the Netropy network emulator through its browser-based GUI, the dedicated Ethernet management port must first be configured with an appropriate IP address and subnet mask. For convenience, the MGMT interface comes pre-configured with an IP address of 10.0.0.10, and is accessible from a directly-connected host on the 10.0.0.0/255.0.0.0 subnet. The IP address and subnet mask of the MGMT interface can be changed through the Netropy GUI or through the command-line interface.

2.1 Preparation

Management of the Netropy device through the GUI requires a PC running a web browser with Flash version 10 or later installed.

Initial configuration of the management interface requires either:

- ▶ a PC running a supported web browser that can be configured and placed on the 10.0.0.0/255.0.0.0 network.
- ▶ a PC with an RS-232 serial port running terminal emulation software such as HyperTerminal or PuTTY.

2.2 Hardware Installation

Plug in a standard power cord (a U.S. power cord is supplied with the unit) and turn on the power. Depending on the particular model, the system will be available for use within 1 to 4 minutes.

For additional hardware installation details, please see the *Hardware Guide* for your model.

2.3 IP Address Configuration via the Netropy GUI

To configure the MGMT interface using the Netropy GUI:

- 1 Configure a PC running a supported web browser with the IP address 10.0.0.2 or other address on the 10.0.0.0/255.0.0.0 subnet.
- 2 Connect an Ethernet cable between the PC and the MGMT port on the Netropy unit.
- 3 Open the browser on the PC and enter `http://10.0.0.10`. Review the End User License Agreement.
- 4 Netropy units are shipped with a temporary license pre-installed for initial operation. Follow the instructions on the screen to download and install a regular license. This step may be skipped, and the license installed later.
- 5 Click on the Administration link at the top of the page and select the Network Settings tab. Set the IP address, subnet mask, and optional default gateway for the management interface, then click the *Apply Changes* button.
- 6 After the management interface has been configured, use the Ethernet cable to connect the MGMT port to the management network.

2.4 IP Address Configuration via the Serial Console

To configure the MGMT interface using the serial console:

- 1 Using the provided serial cable, connect the serial port of a PC running terminal emulation software to the CONSOLE port of the Netropy WAN emulator. Set the serial port parameters to 9600 baud, 8 bits, no parity, 1 stop bit, and disable flow control. See the *Hardware Guide* for your model for more details on connecting to the serial console.

- 2 Press `[ENTER]` to display a login prompt. At the prompt, log in as “admin”. There is initially no password.

```
netropy login: admin
```

- 3 Use the following commands to set the IP address, netmask, and default gateway of the MGMT port:

```
mgmt set addr <ip-address> netmask <mask>
```

```
mgmt set gw <default-gateway>
```

IP addresses and subnet masks are entered in dotted-decimal format. For example:

```
[netropy]> mgmt set addr 192.168.1.1 netmask 255.255.255.0
```

- ④ After the management interface has been configured, use an Ethernet cable to connect the MGMT port to the management network. Open a browser and enter the IP address of the MGMT port. Review the Netropy End User License Agreement.
- ⑤ Netropy units are shipped with a temporary license pre-installed for initial operation. Follow the instructions on the screen to download and install a regular license. This step may be skipped and the license installed later.

2.5 Network Installation

Each Netropy Emulation Engine is installed between two LAN segments and acts as a bridge or router between those two segments. Packets received on one port of the Emulation Engine are subjected to configured emulation conditions before being forwarded or routed to the opposite port.

If configured as a layer 2 bridge, install each Engine on an Ethernet network in a location where the traffic that is to be sent over the emulated WAN will be forced to flow through the device. If configured as a router, install the Engine between two separate subnets and configure static routes to pass traffic through the Engine. Each Engine is configured separately as a bridge or router.

3 CONFIGURATION

Configuration of the Netropy network emulator is aided by understanding a few basic concepts and terminology.

3.1 Emulation Engine

The Netropy Emulation Engine forwards packets and applies the configured emulation conditions between a pair of Ethernet ports.

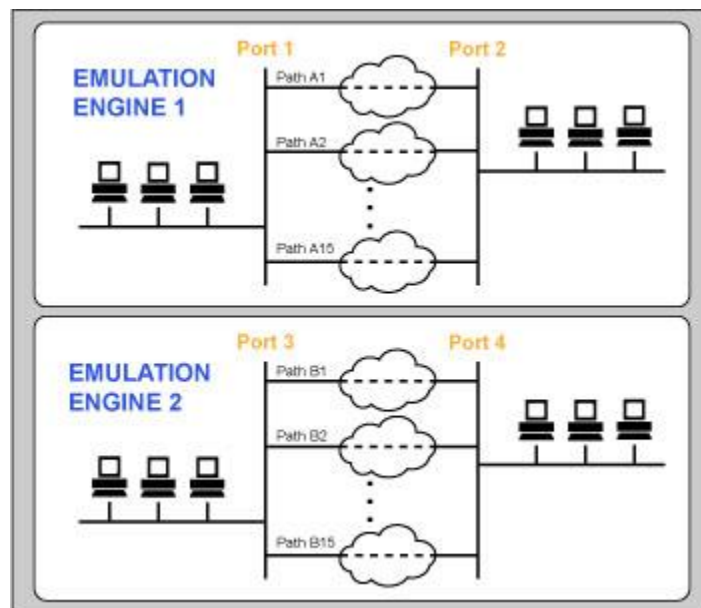


Figure 1: Two separate Emulation Engines, each with 15 paths between each pair of ports.

Depending on hardware model, the Netropy unit contains between one and four separate Emulation Engines. Each Engine operates independently of the others, and can be thought of as a separate emulation device. Each engine has its own Ethernet ports, a network architecture that may include multiple paths and classifiers, and separate traffic statistics and graphs.

See Section 4 for more details on the Emulation Engine.

3.2 Paths

A path is a single configuration of bandwidth, latency, loss, and other network properties. Each Emulation Engine can have up to 15 separate paths.

In the default single-link mode, each path represents a single WAN link, and all traffic using the path is aggregated over the path. In per-client emulation mode, each path represents one set of emulation parameters, but the traffic to and from each separate IP address is allocated its own separate bandwidth.

Each path consists of three components: a WAN access link connecting the LAN to the WAN on each side and traversal over the WAN in the middle.

Each WAN access link connects a LAN to the WAN. Bandwidth constraints and conditions that affect bandwidth availability are configured in the WAN access link. The traversal over the WAN is characterized primarily by its latency, jitter, and loss conditions.

See Section 5 for more details on configuring paths.

3.3 Classifiers

Classifiers are sets of rules or filters that specify which packets are sent over which paths. Each port has its own classifier to direct the packets that arrive on that port. Most users will classify packets by IP source and destination address range, but packets can instead be classified by IPv6 address, VLAN, MPLS label, MAC address, TCP or UDP port number, or any other packet field or combination. Each classification rule includes an action that specifies whether matching packets are sent over one of the configured paths, dropped, or forwarded without emulation.

See Section 7 for more details on configuring classifiers.

3.4 GUI and CLI

Most users will find the browser-based graphical user interface to be the most convenient way to configure and operate the Netropy network emulator. However, a CLI is also available for integration with test automation tools.

The device can be managed and any of the path emulation parameters can be set and modified through the CLI. However, the configuration of the emulated WAN architecture, including creating paths and building classification rules, must first be completed through the GUI.

See Section 11 for more details on the CLI.

4 EMULATION ENGINE

4.1 Overview of Emulation Engines

The Netropy Emulation Engine forwards packets and applies the configured emulation conditions between a pair of Ethernet ports.

Depending on hardware model, the Netropy unit contains between one and four separate Emulation Engines. Each engine operates independently of the others, and can be thought of as a completely separate emulation device. Each engine has its own Ethernet ports, a network architecture that may include multiple paths and classifiers, and separate traffic statistics and graphs.

To configure a particular Emulation Engine from the Netropy GUI, click on the corresponding Engine button at the top of the main page.

Each Emulation Engine can be turned on or off independently. Emulation is initially turned off on all engines after reboot or power cycle. Emulation can be turned on or off from the main page of the GUI or through the CLI. When emulation is off, all packets are forwarded directly between the Emulation Engine's two ports, bypassing any emulation.

Throughput graphs and statistics can be viewed for emulated paths, as well as for the bypass traffic.

The entire configuration of any Emulation Engine can be downloaded to a local file from the Save tab of the Administration window. This configuration file can then be used to reconfigure any Engine.

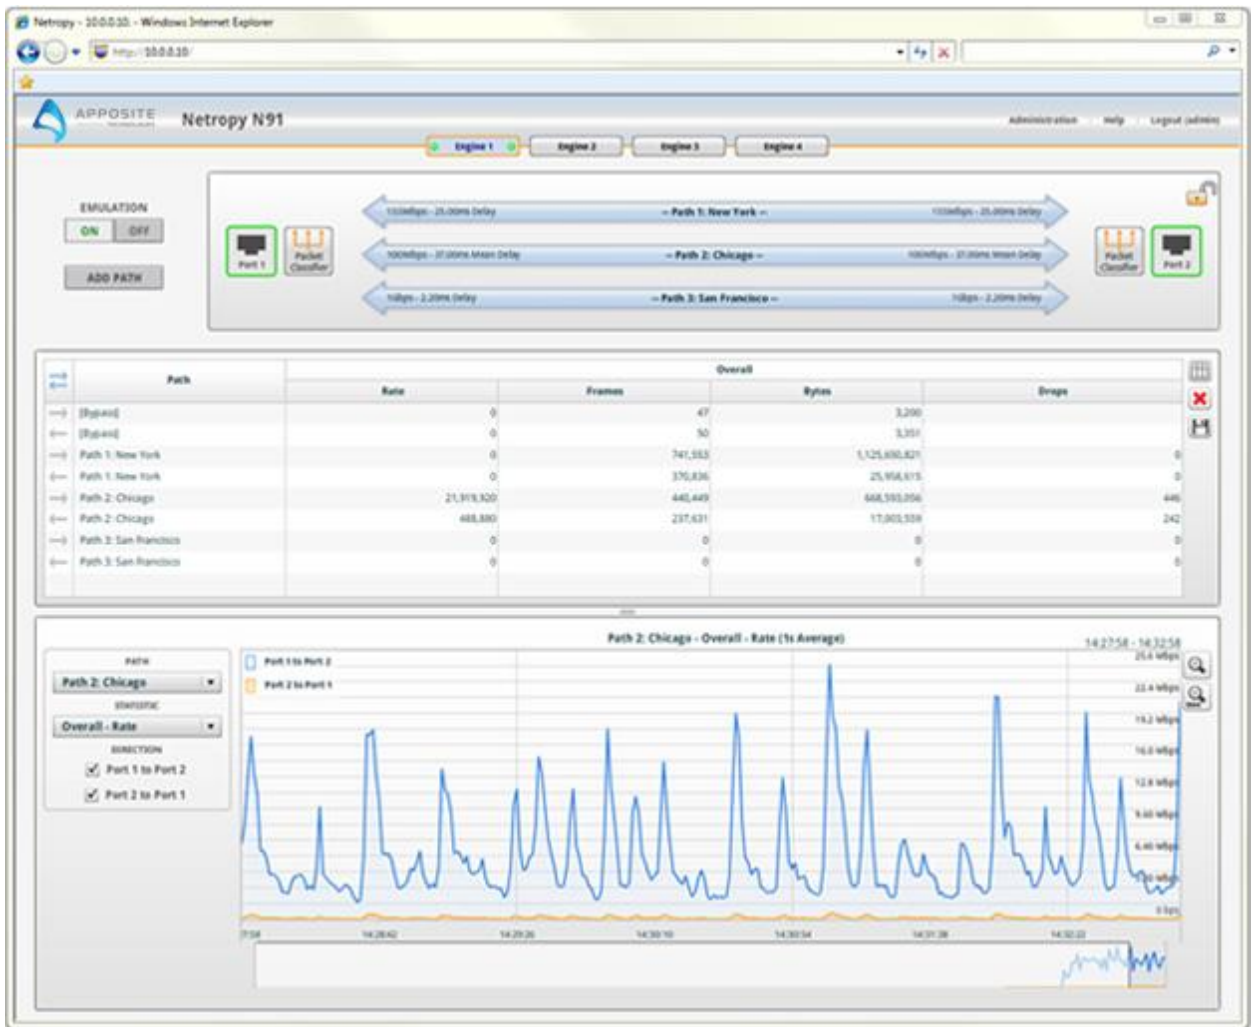



Figure 2: Main Page of the GUI

4.2 Engine Locking

Users can lock individual Emulation Engines to prevent any changes made to the configuration by other users. Only the user who locked the engine and the *admin* user can modify the configuration of a locked engine or unlock the engine. Locking is specific to each individual engine. Different users can lock different engines, or a single user can lock multiple engines.

To lock an Emulation Engine, click the  lock icon. To unlock, click again.

For details on creating and administering users, see Section 10.2.

4.3 Engine Self-Monitoring

The Netropy Emulation Engine performs continuous self-monitoring to ensure that test results have not been affected by limitations on the processing or buffering resources of the Netropy engine.

The current status of the Netropy engine is indicated by the LEDs on each engine selection button. The LED on the left shows engine processing resources and the LED on the right shows engine buffering resources. Resource availability is indicated by the LED color:

Green: Engine is functioning normally with sufficient resources for the current processing and buffering load.

Yellow: Resources are running low. Engine is functioning without error and tests are valid, but if processing or buffering load increases further, the engine is in danger of reaching overload conditions.

Red: Overload fault. Sufficient processing or buffering resources were not available and test results may have been affected. If a fault occurs, please review the error message in the log.

If an overload of processing or buffering occurs, the LEDs will remain red to indicate an error until the log message has been marked as read or the log has been cleared.

Hovering over the engine button displays a tool tip with details of the current status.

5 PATHS

5.1 Overview of Paths

A path is a single configuration of bandwidth, latency, loss, and other network properties. Each Emulation Engine can have up to 15 separate paths. In the default single-link mode, each path represents a single WAN link, and all traffic using the path is aggregated over the path. If per-client emulation mode is enabled, each path represents one set of emulation parameters, but the traffic to and from each separate IP address uses a separate WAN access link with separate bandwidth. Except where otherwise indicated, the descriptions in this User Guide refer to the default single-link mode.

Each path consists of three sections:

- WAN access link connecting a local network (in single-link mode) or individual devices (in per-client emulation mode) to the WAN line or cloud.
- WAN line or cloud.
- WAN access link connecting the WAN line or cloud to a local network.




Figure 3: Path Components

The WAN can be any type of wide-area network connection between two sites including terrestrial private lines, shared networks such as the Internet, and specialized satellite or wireless networks. The WAN is characterized primarily by its latency, jitter, and loss conditions.

Each WAN access connects a local network or individual client connected to the WAN. Bandwidth constraints and conditions that affect bandwidth availability are configured for the WAN access.

5.2 Single-Link and Per-Client Emulation Paths

Netropy offers the choice of treating each path as either:

Single Link Mode: A single WAN link with all traffic using the path sharing the bandwidth.

Per-Client Emulation Mode: A set of emulation parameters that is applied to all traffic using the path, with each IP address assigned a separate WAN access link with separate bandwidth.

Single-link mode is enabled by default. Use the Link Type parameter in the WAN Access panel to switch to per-client emulation model. This path mode setting can only be changed while emulation is turned off.

5.2.1 Single-Link Mode

In the default single-link mode, each path represents a single WAN link. All traffic using the path is aggregated over a single WAN access link on each side of the WAN and shares the bandwidth. In this mode, each Emulation Engine can simulate up to 15 individual WAN links.

Single-link mode is appropriate for emulating a network in which end clients share WAN bandwidth, such as a few local offices connecting to a datacenter or two datacenters connected over the WAN.

To use single-link mode, set the Link Type parameter for the outbound wan access on both ports to single-link.

5.2.2 Per-Client Emulation Mode

In per-client emulation mode, each path represents a set of emulation parameters, but the traffic to and from each IP address uses a separate WAN access link with its own separate bandwidth. In this mode, each path can simulate thousands of WAN links, and additional paths can be used to provide up to 15 different sets of emulated conditions. Per-client emulation mode is appropriate for testing large numbers of home users or mobile users independently connecting to one site or one server.

To use per-client emulation mode, set the Link Type parameter on the client side outbound wan

access to Per IP Source and the Link Type on the server side outbound WAN access to Per IP Destination.

Although each IP address has its own WAN access link with its own bandwidth allocation, the packets are aggregated over the WAN. Consequently, the traffic from each client is not treated separately, which may affect some WAN parameters such as variable delay or periodic loss. For example, if periodic loss is configured to drop every tenth packet, then every tenth packet from all clients together will be dropped rather than every tenth packet transmitted by each individual client.

As packets arrive for a path, Netropy automatically creates a new WAN access link for each new IP address. Netropy can create separate WAN access links for up to 65,536 IP addresses on each path. Once a WAN access link is created, it is not removed until emulation is disabled. If this limit is reached, packets for new IP addresses are dropped and an error is logged.

When using the per-client emulation mode, only the WAN access parameters of bandwidth, queue limit, and frame overhead can be emulated. The statistics and graphs display aggregate values across all clients, but since queue depth is only meaningful for individual clients, the statistics display a value of 0.

5.3 Path Types

Most network connections can be characterized as either point-to-point or cloud-based.

Private lines such as T1 or OC-3 lines directly connecting two sites are point-to-point connections. The bandwidth out-bound from one site is the same as the bandwidth in-bound at the other site, and typically the latency is constant. On these links, the bandwidth is throttled in the outbound direction from each site, and there is generally no need to configure the in-bound WAN access parameters.




Figure 4: Point-to-Point Line

Network connections that consist of an access link to a shared network such as the Internet, an MPLS network, or a Metro Ethernet ring, can be thought of as a cloud. Cloud networks typically have limited bandwidth access to a high-speed WAN, and frequently different speed access to the WAN at each site. In these situations, traffic can hit a bandwidth bottleneck both out-bound from a site to the WAN and in-bound from the WAN to the other site, making it necessary to configure both out-bound and in-bound WAN access parameters. Similarly, if there is variable delay in the WAN cloud, it may be necessary to configure the in-bound WAN access parameters to rate-limit the resulting flow.




Figure 5: Cloud Network

5.4 Configuring Paths

Each path is displayed on the main page of the Netropy GUI using an arrow labeled with its name. To configure a path, click on the arrow to open the Path Configuration window.

The Path Configuration window is used to rename the path, choose whether the path is bi-directional or uni-directional, and configure the WAN and WAN access conditions of the path.

To rename a path, click the existing name to edit it.

To set a path to be bi-directional or uni-directional, use the direction selectors on the top button bar. A path can be set to uni-directional only if the unused direction is not referenced by any classifier actions. The WAN and WAN access conditions of a path are configured independently in each direction of a bi-directional path.

When finished, click the *Apply Changes* button. If the Path Configuration window is closed without clicking the *Apply Changes* button, any unapplied changes will be lost.

Changes to the WAN and WAN Access conditions can be made from the CLI, but the path name and directionality can only be set from the GUI.

5.5 Adding and Deleting Paths

To add a new path, click the *Add Path* button on the main page.

To delete an unneeded path, click the path to enter the Path Configuration window and use the delete button on the top button bar to remove. Paths cannot be deleted while they are being used in any classifier rule.

6 PATH PARAMETERS:

CONFIGURING EMULATION CONDITIONS

6.1 Configuring Path Parameters

Each path consists of three sections:

- ▶ WAN access link connecting a LAN to the WAN
- ▶ Traversal of the WAN line or cloud
- ▶ WAN access connecting the opposite LAN to the WAN

The active configuration is not modified until the *Apply Changes* button is pressed. The *Apply Changes* button is grayed out when there are no changes to apply or if there are any invalid entries.

The screenshot shows the 'Path Configuration' window with the title 'HQ to Regional Office'. It features three main configuration panels:

- WAN Access - Port 1 (OUTBOUND to WAN):** Includes settings for Link Type (Single Link), Bandwidth (Fixed), Rate (1 Gbps), Background Utilization (Off), Queue Limit (Default), Queuing Strategy (Default (FIFO)), MTU Limit (Off), and Frame Overhead (18 - Ethernet HDR+FCS).
- WAN (PORT 1 to PORT 2):** Includes settings for Delay (Off), Loss (Off), Network Outage (Off), Corruption (Off), Reordering (Off), and Duplication (Off).
- WAN Access - Port 2 (INBOUND from WAN):** Includes an 'ENABLE' checkbox and settings for Link Type (Single Link), Bandwidth (Fixed), Rate (1 Gbps), Background Utilization (Off), Queue Limit (Default), Queuing Strategy (Default (FIFO)), MTU Limit (Off), and Frame Overhead (18 - Ethernet HDR+FCS).

At the bottom, there are 'Apply Changes' and 'Clear Changes' buttons.

Figure 6: Path Configuration Window

Final validation of the configuration is performed when the *Apply Changes* button is pressed. If there are no errors in the configuration, the new configuration immediately takes effect. If there are any errors in the configuration, a red box is displayed around the invalid panel. Hovering over the panel displays a description of the error.

6.2 WAN Access Parameters

The WAN Access panel configures the parameters that affect bandwidth availability for traffic outbound from the LAN to the WAN, and optionally for traffic inbound onto the LAN from the WAN. There are separate panels for the WAN Access link on both sides of the WAN.

The WAN access link configuration can either be a single amount of bandwidth shared by all traffic over the WAN access link (single-link mode) or a set of emulation parameters that are applied independently to the traffic from each separate IP address (per-client emulation mode.) See Section 5.2 for more information on link types.

In single-link mode, all WAN access emulation parameters are available. In per-client emulation mode, only the WAN access parameters of bandwidth, queue limit, and frame overhead can be emulated. Background utilization, queuing strategy, and MTU limits cannot be emulated and are not displayed in the GUI.

WAN Access - Port 1

OUTBOUND to WAN >>>

Bandwidth - Fixed

Rate 155 Mbps

Background Utilization - Random

Rate 10 %

Burst 1500 bytes

Queue Limit - Drop Tail

Queue Depth 100 KB

Queuing Strategy - Default (FIFO)

MTU Limit - On

MTU Limit 576

Send ICMP Error ☒ Enabled ☐ Disabled

IP Fragmentation Standard

Frame Overhead - 18 - Ethernet HDR+FCS

<<< INBOUND from WAN

☐ ENABLE

Figure 7: WAN Access Configuration Panel

6.2.1 Link Type

The Link Type panel is used to select between single-link mode and per-client emulation mode. In per-client emulation mode, set the client-side outbound WAN access link type to Per IP Source and the server side outbound WAN access to Per IP Destination.

Inbound WAN access, if enabled, can only be configured as a single link.




Figure 8: Link Type Panel

6.2.2 Bandwidth

The Bandwidth panel is used to configure the rate of the WAN access link.

The bandwidth is set in increments of 1 bit per second, with a minimum rate of 100 bps and a maximum rate determined by the maximum port speed (1 Gbps, 10 Gbps, or 40 Gbps depending on hardware model) unless limited to a lower rate by the license key. (See Section 10.11 for details on the license key.)

A series of values for bandwidth changing over time can be added in a Recording file and played using the Playback feature. See Section 8 for details.

If the entered bandwidth is higher than the maximum port speed or higher than the license key, a red error box will be drawn around the panel when the *Apply Changes* button is pressed and the changes will not be applied.



Figure 9: Bandwidth Panel

6.2.3 Background Utilization

The Background Utilization panel is used to create extra traffic that competes for bandwidth with the real application traffic passing through the WAN access link. Background traffic only affects the WAN access link on which it is configured, and is not transmitted through the other components of the path or outside the Netropy device. To have background traffic compete for bandwidth on the opposite WAN access link, create an identical background traffic configuration on the WAN access link inbound to the opposite port.

Background traffic can be useful for testing the performance of particular applications over links that are congested with other traffic, and for inducing jitter to test real-time applications. Background traffic can either be created with random packets based on an average link utilization rate or by replaying imported PCAP packet capture files.


6.2.3.1 Random Background Traffic

The Netropy Emulation Engine can generate random background traffic to compete with the real application traffic for bandwidth across the emulated WAN link. Random background traffic is specified as a link utilization rate and a traffic burst size.

The link utilization rate specifies the average percentage of bandwidth consumed by the background traffic. The link utilization rate can be set to 0 – 100% of the bandwidth in increments of 0.01%.

The burst size sets the size of the background traffic blocks and is specified in bytes from 64 – 2,000,000 bytes. The default value is 1500 bytes. Larger bursts of background traffic induce greater jitter in the actual traffic.

Random background traffic is modeled as a Poisson process in which bursts of data of a fixed size are transmitted at an average rate such that the bandwidth will be occupied at the specified link utilization rate. Because it is a random process, over short periods the actual background utilization rate may vary from the configured value.



Background Utilization - Random	
Rate	50 %
Burst	1500 bytes

Figure 10: Background Utilization - Random

6.2.3.2 Packet Replay Using PCAP Files

The Netropy Emulation Engine can replay PCAP packet capture files as background traffic to compete for bandwidth with the real application traffic.

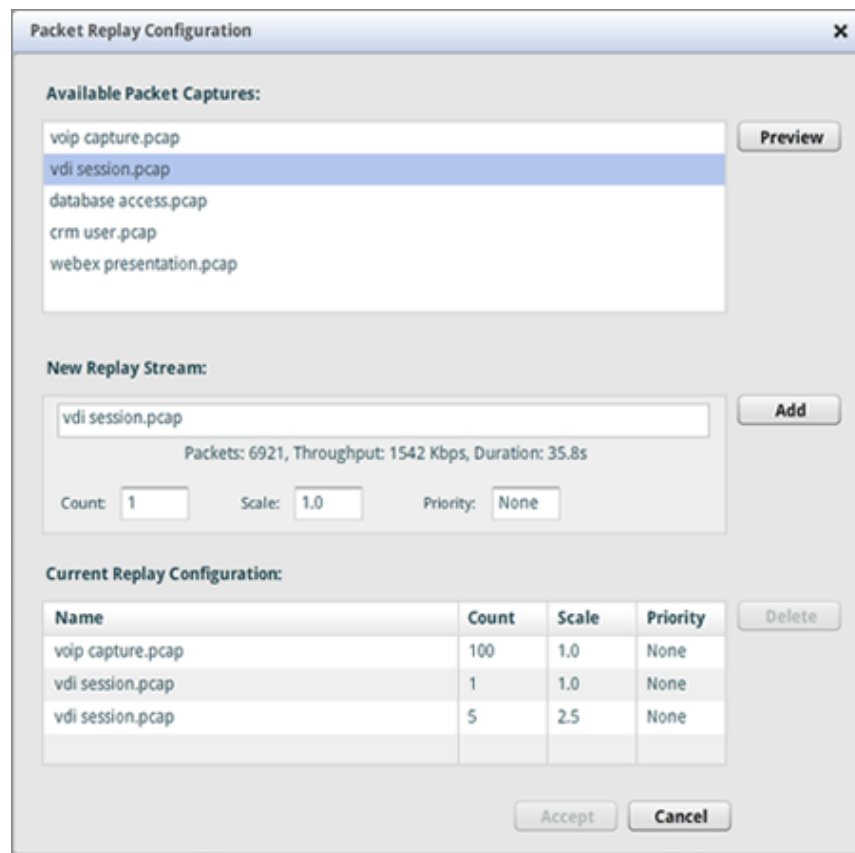


Figure 11: Packet Replay Configuration

Network traffic can be captured and saved to a PCAP file by Wireshark, tcpdump, or other protocol analyzer or network monitor tool. PCAP files are imported into Netropy through the Packet Captures tab of the Administration window (Section 10.6).

Any valid PCAP file can be used to generate background traffic, including individual streams, all traffic between two devices, all traffic from one device to any other device, or all traffic on the network over a period of time.

All packets in the PCAP file are replayed in a single direction. Bidirectional packet captures generally need to be split into two unidirectional capture files.

Multiple packet capture files can be replayed simultaneously, and each file can be replicated to simulate large numbers of streams or scaled to adjust the timing and bandwidth usage.

Each packet capture file can be up to 10 MB in size, with the total for all packet capture and recording files limited to 40 MB.

To configure PCAP replay background traffic, select packet replay on the Background Traffic panel and click *edit* to open the Packet Replay Configuration window.

The top of the Packet Replay Configuration window displays a list of available PCAP files previously imported into the Netropy device. Click a file to configure replay of a stream based on that packet capture. The *preview* button displays a graph of the data rate of the stream.




Figure 12: Background Utilization – Packet Replay

The New Replay Stream section of the window displays the number of packets in the selected file, its average data rate, and time duration of the capture. Configure replay with the following parameters:

- ▶ **Count:** Number of separate copies of the stream to run. By default, one copy is run but the same stream can be replicated up to 20,000 copies to simulate multiple users or clients. If multiple copies are run, each copy will start at a random location in the stream.
- ▶ **Scale:** Time scaling factor between 0.001 and 1000 that adjusts the speed of the packet reply. The default value of 1.0 replays the stream at the same speed as the original transmission. A value of 2.0 replays the stream at twice the original speed, thereby doubling the data rate. A value of 0.1 replays the stream at a tenth of original speed.
- ▶ **Priority:** A value between 0-7 used for IP Precedence or VLAN PCP priority level for all packets in the file when strict prioritization or round robin queuing strategy is enabled. If the default value of 'none' is selected, the original IP Precedence or VLAN PCP value of each packet is used.

Click the *add* button to add the file to the replay configuration. The list of files to be replayed is displayed in the Current Replay Configuration table. A file can be removed from the configuration by selecting the file and clicking the *delete* button. Click *accept* to complete configuration of packet capture replay and return to the Path Configuration window. Packet replay begins when the changes to the path configuration are accepted if emulation is already enabled, or when emulation is next turned on.

Up to 20 separate playback entries can be created. Multiple copies of the same file can be used with the same or different count, scale, and priority settings. A single PCAP file could therefore be used, for example, to simulate ten streams at half speed, two streams at the original speed, and five streams at twice the original speed.

Packet replay files and parameters can be changed on the fly while emulation is on. If the count, scale, or priority parameters for any stream are changed, playback for only that stream will start from the beginning when the path configuration changes are accepted, and any other running playback streams will not be affected.

Packet replay only generates background traffic to compete with actual application traffic from external sources. Packets from the PCAP files are not transmitted out of the Netropy device.

Packet replay consumes Netropy resources equivalent to traffic from external sources and must be included when considering packet forwarding and buffering limits of the device.

6.2.4 Queue Limit

The Queue Limit panel is used to select the queue management algorithm and configure the associated queuing parameters. The queue management algorithm controls the buffering and discarding of packets when they arrive faster than the rate of the WAN access link. The queue management algorithm and parameters can be set to match the configuration of an existing WAN access router.

There are three choices for the Queue management algorithm:

- Drop Tail

The Drop Tail algorithm (also called tail drop) is a simple FIFO queue of a configured maximum size. When the buffer is full, any additional packets that arrive are discarded. Using Drop Tail, specify the size of the buffer in KB or packets.




Figure 13: Queue Limit – Drop Tail

► RED

Random Early Detection (RED) is an active queue management algorithm that monitors the average queue size and begins randomly dropping a small number of packets before the queue is full to create smoother flows and fairer drops. RED begins dropping packets at the configured minimum threshold, with the probability of drop increasing linearly until the configured maximum threshold, after which all packets are dropped. Configure the total buffer size, minimum threshold, and maximum threshold in KB or packets. For more details on RED, see <http://www.icir.org/floyd/red.html>. The value used for \max_p is 0.1 and for w_q is 1/512.



Queue Limit - RED	
Queue Units	<input type="radio"/> KB <input checked="" type="radio"/> Packets
Queue Depth	<input type="text" value="128"/> packets
Min Threshold	<input type="text" value="64"/> packets
Max Threshold	<input type="text" value="128"/> packets

Figure 14: Queue Limit – RED

► Default


The default option sets the queue management algorithm to Drop Tail and configures the queue depth to the equivalent of 250 ms at the currently configured bandwidth rate. For example, if the bandwidth is set to 100 Mbps, the default queue depth will be 3.125 MB. Changes to the bandwidth will automatically adjust the queue depth.

If priority queuing is selected under Queuing Strategy, the specified queue limits apply separately to the queue for each priority level.

All entries for queue depth and thresholds are limited to 100,000 packets or 100,000 KB.

6.2.5 Queuing Strategy

The queuing strategy panel determines the manner in which packets are queued and transmitted. The three options are a single FIFO queue, Priority queuing, and Round Robin.



Queuing Strategy - Default (FIFO)
Queuing Strategy - Default (FIFO)
Queuing Strategy - Priority
Queuing Strategy - Round Robin

Figure 15: Queuing Strategy

► Default (FIFO)

The default option uses a single FIFO queue. Packets are transmitted in the order they arrive, with no prioritization of packets.

► Priority

For class of service prioritization, incoming packets can be directed onto eight separate priority queues based on the priority setting in the packet. Packets are transmitted based on strict priority: if there are any packets on a higher priority queue, they will be transmitted before any packets on a lower priority queue. Queues are numbered from highest (7) to lowest (0).

The queue management algorithm and settings specified in the Queue Limit panel applies separately to each of the eight priority queues. For example, if Drop Tail is selected with a queue depth of 100 KB, each of the eight priority levels will consist of its own 100 KB queue.

There are two options for specifying the field to use for the priority level of the packets:

IP Precedence: the three bits of precedence in the ToS field of IPv4 packets, or the three bits of precedence in the traffic classifier of IPv6 packets.

VLAN PCP: the three bit Priority Code Point field in the VLAN header.

► Round Robin


Similar to Priority queuing, incoming packets are directed onto eight separate queues based on the IP Precedence or VLAN PCP priority value of each packet. Packets are pulled from each queue and transmitted in round robin order.

As in Priority queuing, the queue management algorithm and settings specified in the Queue Limit panel applies separately to each of the eight queues.

6.2.6 MTU Limit

The MTU limit panel allows the setting of a path MTU (Maximum Transmission Unit), specifies whether ICMP error messages are sent, and specifies whether IPv4 packets larger than the MTU limit are fragmented.

If MTU limits are enabled, any IPv4 packet that exceeds the MTU can either be dropped or fragmented, depending on the IP Fragmentation setting:



MTU Limit - On	
MTU Limit	1500
Send ICMP Error	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
IP Fragmentation	Standard

Figure 16: MTU Limit

► Standard

IPv4 packets without the Don't Fragment (DF) bit set are fragmented and all other packets are dropped.

- ▶ Never – Drop Only
Packets larger than the MTU limit are always dropped.
- ▶ Always – Ignore DF
Packets larger than the MTU limit are always fragmented regardless of the setting of the DF bit. This option should only be used when specifically required for testing.

All non-IPv4 packets larger than the MTU limit are dropped. IPv6 packets are never fragmented.

The MTU limit can be set to any value between 68 bytes and 9216 bytes.

The sending of IPv4 ICMP Destination Unreachable Fragmentation Needed or IPv6 PKTTOOBIG error messages to the originator of the packet can be enabled or disabled. When enabled, ICMP error messages are transmitted out the interface on which the original packet was received switching the source and destination Ethernet and IP addresses of the original packet.

ICMP error messages are limited to 15 packets per second (per path per direction), with short term bursts of up to 15 packets.

6.2.7 Frame Overhead

Frame overhead is the number of additional bytes required by a link-layer technology when transmitting a packet of data. Typically, the frame overhead consists of link-layer addressing and error checking information.

To emulate a link-layer technology with a particular frame overhead, select the value from the drop-down list, if available, or choose *Custom* and enter the specific value.

To emulate the traversal of a frame over the WAN, Netropy calculates the effective size of the frame as the payload of the Ethernet frame (without the Ethernet header or FCS) plus the specified frame overhead.

There are three choices for frame overhead:

- ▶ Ethernet (header + FCS)
This option emulates a WAN link layer with an Ethernet-like frame of 18 bytes of header and frame check sequence (FCS). This is the default option and is a reasonable choice if the properties of the link layer are unknown.
- ▶ Ethernet (header, FCS, preamble, pad)
This option emulates an actual Ethernet link, including the preamble and padding between Ethernet frames. Select this option to emulate an Ethernet-based WAN network.
- ▶ Custom
This option allows the specification of any link layer frame overhead in bytes per packet up to a maximum of 300 bytes. Select this option if the link layer frame overhead is known.

6.3 WAN Parameters

WAN delay, loss, link outage, corruption, reordering, and duplication parameters are configured on the WAN section of the Path Configuration window. The conditions are set separately for the two directions.

The image shows a 'WAN' configuration panel with two sections for different directions of traffic. The top section is for 'PORT 1 to PORT 2' and the bottom section is for 'PORT 2 to PORT 1'. Each section contains various parameters for delay, loss, corruption, reordering, and duplication.


Direction	Parameter	Value
PORT 1 to PORT 2	Delay - Constant	70 ms
	Loss - Random	0.1 %
	Corruption - BER	1 x 10 ⁻⁸
	Reordering - On	1.5 %
	Offset Range	5 packets
	Timeout	10 ms
	Duplication - Off	
PORT 2 to PORT 1	Delay - Normal	Mean: 70.0 ms, Std Deviation: 10.0 ms
	Allow Reordering	<input type="checkbox"/>
	Loss - Burst	Burst Probability: 1 %, Minimum Burst: 1 packets, Maximum Burst: 5 packets
	Corruption - Off	
	Reordering - Off	
	Duplication - On	Probability: 3 %

Figure 17: WAN Parameter Configuration Panel

6.3.1 Delay

The Delay panel sets the latency and jitter in each direction. For variable latency distributions, a short delay applied to a later packet may cause it to have a calculated transmission time prior to that of earlier packets with a longer delay. By default, packets are transmitted in the order received, which can skew the actual amount of delay applied. If “Allow Reordering” is selected, the order of the packets can be changed.


- Off: No latency added.
- Constant: A single, fixed value for latency.



The screenshot shows a panel titled "Delay - Constant". It contains a single input field labeled "Delay" with the value "50" and the unit "ms".

Figure 18:
Delay - Constant


- Uniform: A uniform distribution of latency ranging between the configured minimum and maximum values. The Minimum value must be less than or equal to the Maximum.



The screenshot shows a panel titled "Delay - Uniform". It contains two input fields: "Minimum" with the value "45" and "Maximum" with the value "55", both followed by "ms". There is also a checkbox labeled "Allow Reordering" which is currently unchecked.

Figure 19:
Delay – Uniform Distribution


- Exponential: An exponential distribution curve, with a specified minimum and mean.



The screenshot shows a panel titled "Delay - Exponential". It contains two input fields: "Minimum" with the value "35.00" and "Mean" with the value "50.00", both followed by "ms". There is also a checkbox labeled "Allow Reordering" which is currently unchecked.

Figure 20:
Delay – Normal Distribution


- Normal: A normal (Gaussian) distribution, with a specified mean and standard deviation (jitter). To avoid negative latencies, the mean must be at least 3 times the Std Deviation.



The screenshot shows a panel titled "Delay - Normal". It contains two input fields: "Mean" with the value "50" and "Std Deviation" with the value "5", both followed by "ms". There is also a checkbox labeled "Allow Reordering" which is currently unchecked.

Figure 21:
Delay – Normal Distribution

- **Accumulate & Burst:** Packets are held until either a packet count or time threshold is reached, then optionally delayed by an additional configured 'extra delay,' then transmitted as a burst. The timer for the time threshold is started when the first packet in the burst is received. The maximum packet count threshold is 1000 packets, and the maximum time threshold is 10000 ms (10 seconds).



Delay - Accumulate and Burst	
Burst After	10 packets
-or-	200 ms
Extra Delay	10 ms

Figure 22:
Delay - Accumulate and Burst

All delay values are specified in milliseconds in increments of 0.01 ms.

In addition to specifying the latency using this delay panel, a series of values for latency that change over time can be added in a Recording file and played using the Playback feature. See Section 8 for details.

By default, frames are not reordered even if subjected to differing delays using a uniform or normal distribution. To allow packets to be reordered, check the “Allow Reordering” option. For example, if the delay is set as a uniform distribution between 10 and 100 ms and the first frame is subjected to a 90 ms delay and the second frame is subjected to a 20 ms delay, by default, the second frame cannot be transmitted until after the first frame has been transmitted. If “Allow Reordering” is selected, (and assuming no congestion) the second frame will be transmitted 20 ms after arrival and the first frame will be transmitted 90 ms after arrival, causing the order of the packets to be switched.

To specify jitter, use either the normal or uniform distribution. Use the normal distribution to specify jitter as the standard deviation from the mean delay. Use the uniform distribution to specify peak-to-peak jitter between the minimum and maximum values of delay.

The various Netropy hardware models have different limits on the ability to process high packet rates combined with large latencies. For details, see the *Hardware Guide* for your model.



The end-to-end round trip time (RTT) is a combination of four separate delays in each direction: the propagation delay, transmission delay, queuing delay, and reordering delay.

6.3.2 Loss

The Loss panel configures packet loss each direction. The options are:

- ▶ Off: No packet loss.
- ▶ Random: Random packet loss. Specify a single value for the probability that each packet will be lost. Rates can be set from 0 – 100% in increments of 0.0001%.

Loss - Random ▼	
Loss Rate	<input type="text" value="25"/> %

Figure 23:
Loss - Random

- ▶ Burst: Burst loss. Specify the probability that each packet will begin a burst of lost packets, and a minimum and maximum number of packets that will be lost in sequence. For a fixed burst size, set the minimum and maximum to the same value. Probabilities can be set from 0 – 100% in increments of 0.0001%.

Loss - Burst ▼	
Burst Probability	<input type="text" value="5"/> %
Minimum Burst	<input type="text" value="5"/> packets
Maximum Burst	<input type="text" value="8"/> packets

Figure 24:
Loss - Burst

- ▶ Gilbert-Elliott: Gilbert-Elliott two-state loss. Specify the packet loss rates for the “good” and “bad” states, and specify the per-packet probability of transitioning from each state to the other. All rates are specified as percentages set from 0 – 100% in increments of 0.0001%. When Gilbert-Elliott loss is first configured and each time emulation is subsequently turned on, loss starts in the good state.

Loss - Gilbert-Elliott ▼	
Good State	<input type="text" value="0"/> % Loss
	<input type="text" value=".001"/> % Change
Bad State	<input type="text" value="10"/> % Loss
	<input type="text" value="1"/> % Change

Figure 25:
Loss – Gilbert-Elliott

- **Periodic:** Periodic packet loss. Specify the loss period and burst size in numbers of packets. For example, a period of 1000 packets with a burst size of 10 packets would result in a fixed pattern of 990 packets forwarded followed by 10 packets dropped.




Figure 26:
Loss - Periodic

- **BER:** Loss due to bit errors. Set the coefficient and exponent. Bit error rates can take values of 1×10^{-18} or greater and are entered in scientific notation. The coefficient of the rate must be entered as a value greater than or equal to 1 and less than 10. All packets that contain a bit error are discarded – to transmit corrupted packets, use the Corruption emulation.



Figure 27:
Loss - BER

In addition to the specifying the loss using this panel, a series of random loss rates that change over time can be added in a Recording file and played using the Playback feature. See Section 8 for details.

6.3.3 Network Outage

The Netropy Outage panel simulates a periodic link outage causing 100% packet loss for a specified duration and interval. The interval is the amount of time between the beginning of successive outages, not the amount of time between the end of one outage and the start of the next outage.

Both the duration and interval can be set as single values or ranges of values separated by a hyphen. Values can be set from 1 ms (0.001 seconds) to 1 week (604,800 seconds) in seconds in increments of 0.001 seconds. If the duration is set to a value that is longer than the interval, the outage will be continuous.





Figure 28: Network Outage

6.3.4 Corruption

The Corruption panel is used to insert bit errors into forwarded packets at the specified bit error rate. Set the BER coefficient and exponent. Bit error rates can take values of 1×10^{-18} or greater and are entered in scientific notation. The coefficient of the rate must be entered as a value greater than or equal to 1 and less than 10.




The screenshot shows a panel titled "Corruption - BER" with a dropdown arrow. Below the title is a label "Bit Error Rate" followed by a text input field containing "1" and a separate box containing "-7", with "x 10" between them.

Figure 29: Corruption

Corruption only affects the contents of received Ethernet frames. Neither the Ethernet header (including EtherType and optional VLAN tag) nor the Ethernet FCS will be corrupted.

6.3.5 Reordering

The Reordering panel specifies the probability for each packet that it is reordered, and how far back in the data stream the reordered packet is moved from its original position. If a packet is randomly selected for reordering, it is held until the offset number of packets arrive and reinserted into the data stream at that point. For example, if the offset is 5 packets, any packet that is reordered will be held and reinserted after the fifth subsequent packet.



The screenshot shows a panel titled "Reordering - On" with a dropdown arrow. Below the title are three rows of controls: "Probability" with a text input field containing "1" and a "%" label; "Offset Range" with a text input field containing "20" and a "packets" label; and "Timeout" with a text input field containing "10000" and a "ms" label.

Figure 30: Packet Reordering

To configure packet reordering, set:

- **Probability:** the likelihood that each frame will be reordered. Probability can be set from 0 – 100% in increments of 0.0001%.
- **Offset Range:** the number of packets that the reordered packet is moved back in the data stream. Either a single value or a range of values can be configured. To specify a range, input the minimum and maximum reordering offsets separated by a dash, i.e. 5-12.
- **Timeout:** the maximum amount of time to wait for the number of offset packets to arrive. For example, if the offset is set to 1000 packets and the timeout set to 5 ms, if 1000 packets do not arrive within 5 ms, the packet will be reinserted in the packet stream at that expiration of the 5 ms period. The default value for timeout is 10,000 ms. The timeout value is specified in ms in increments of 0.01 ms.

Only one packet can be held for reordering at any time. While a packet is waiting for reinsertion, the arriving packets are not subject to reordering. For example, if a packet is randomly selected for reordering with an offset of 5 packets, the next five packets that arrive cannot also be reordered.

6.3.6 Duplication

The Duplication panel specifies the probability for each packet that it is duplicated.




Figure 31: Packet Duplication

Duplicate packets are inserted into the data stream immediately after the original packet. Duplicate packets are then subjected to delay, loss, and reordering independently of the original packet.

The duplication probability can be set from 0 – 100% in increments of 0.0001%.

7 PACKET CLASSIFIERS

7.1 Overview of Packet Classification

Classifiers are ordered sets of rules or filters that specify which packets are sent over which paths. The first rule in order that matches the packet is used to specify the action for the packet.

Each classification rule consists of two components: match and action. Packets may be matched by source and destination IP address (IPv4 or IPv6), VLAN, MPLS label, MAC address, TCP or UDP port number, or any other packet field or combination. Packets that match the rule then follow the configured action. There are three options for the action of each rule:

- ▶ Use path: Packets matching the rule are sent over the specified path.
- ▶ Drop: Packets matching the rule are dropped.
- ▶ Bypass: Packets matching the rule are forwarded without emulation.

The default action specifies the action for packets that do not match any of the explicit rules.




Figure 32: Packet Classification

To configure the classifier, click the *Packet Classifier* button next to either port on the main window. Either button opens the Classifier Configuration window to set the classifier for both ports. Select the classification method for each port, then add rules. Up to 30 separate rules can be created for each port.

Only one of the classification methods can be used on each port. To make rules for more than one classification type, such as both IPv4 and IPv6 packets, to build complex rules with multiple criteria such as a specific IP address within a specific VLAN, or to create rules based on any arbitrary field or identifier in the packets, use the flexible Combination Classifier.

For symmetric configurations, configure rules on one port and use the *mirror* button to automatically create rules for the opposite port.

To forward ARP packets without impairment, check the “ARP Frame Bypass Emulation” option in the Bridge/Route tab of the Administration window (Section 10.3.1). When this option is enabled, all ARP packets are forwarded directly between the ports bypassing emulation regardless of any classification rules that would otherwise apply to the packets.

Classifiers can only be set in the GUI and cannot be created or modified in the CLI.

7.2 Classification Methods

Packets may be classified by source and destination IP address (IPv4 or IPv6) range, VLAN, MPLS label, MAC address, TCP or UDP port number, or any other packet field. Use the drop-down menu to select the classification method for each port. Only one classification method may be selected for all packets arriving on the port. The “Combination” classifier allows packets to be classified based on multiple classification types, a combination of fields, or to any arbitrary header field or packet identifier.

7.2.1 Classification Off

With classification turned off, a single specified action is applied to all packets.

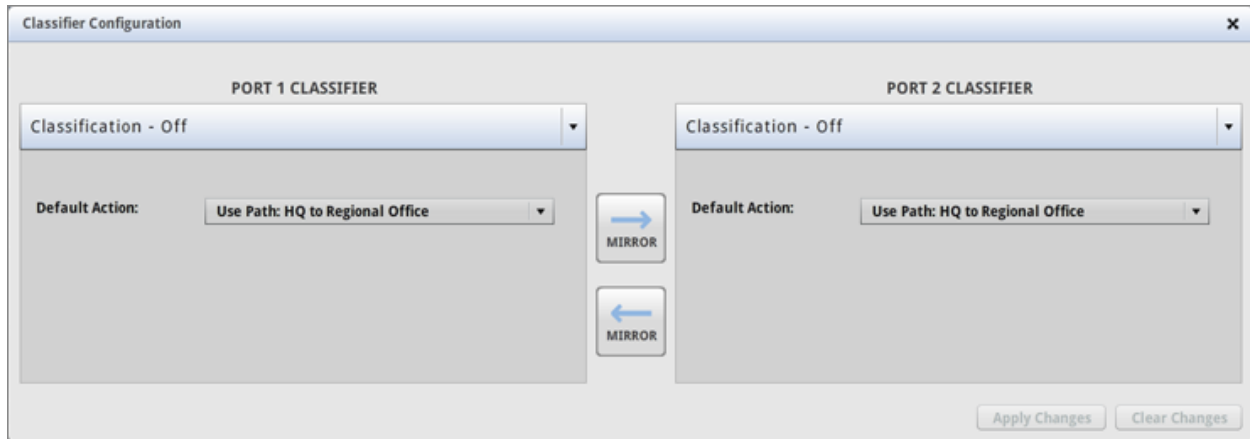


Figure 33: Classification Off

.2.2 IP Address Classification

IP address classification matches packets by source and destination IPv4 address or range of addresses. For a packet to match a rule, both the source and destination address of the packet must fall within the configured source and destination address ranges of the rule.

IP address ranges may be specified in the following formats:

- ▶ as a single address in dotted decimal notation
ex. 10.0.0.10
- ▶ as a range of IP addresses separated by a dash
ex. 10.0.0.0 – 10.255.255.255
- ▶ as a range of IP addresses represented by an address and prefix length
ex. 10.0.0.0/24
- ▶ using the wildcard 'any' to match any IP address

Non-IP packets and IP packets that do not match any of the numbered rules follow the separately configured Default Action.

For classification based on the Precedence or DSCP Type of Service (ToS) field in the IP header, use the Combination Classifier.

The screenshot displays the 'Classifier Configuration' window, which is divided into two main sections: 'PORT 1 CLASSIFIER' and 'PORT 2 CLASSIFIER'. Both classifiers are set to 'Classification - IP Address'. Each classifier contains a list of rules, numbered 1 through 4. The rules are configured as follows:

Classifier	Rule Number	Source IP Range	Destination IP Range	Action
PORT 1 CLASSIFIER	1	20.0.0.112	any	Use Path: Path 1 New York
	2	20.0.0.0/24	29.0.0.0/24	Use Path: Path 2 - Chicago
	3	20.0.0.0 - 20.0.0.255	any	Use Path: Path 3 - San Francisco
	4	21.0.0.112	any	Use Path: Path 4 - London
PORT 2 CLASSIFIER	1	any	20.0.0.112	Use Path: Path 1 New York
	2	29.0.0.0/24	20.0.0.0/24	Use Path: Path 2 - Chicago
	3	any	20.0.0.0 - 20.0.0.255	Use Path: Path 3 - San Francisco
	4	any	21.0.0.112	Use Path: Path 4 - London

At the bottom of each classifier section, there are buttons for 'Add Rule', 'Delete Rule', 'Move Up', and 'Move Down'. A 'Default Action' dropdown is also present, set to 'Bypass'. Between the two classifier sections, there are 'MIRROR' buttons with arrows pointing in opposite directions. At the bottom right of the window, there are 'Apply Changes' and 'Clear Changes' buttons.

Figure 34: Classification by IP Address

7.2.3 IPv6 Address Classification

IPv6 address classification matches packets by source and destination IPv6 address or range of addresses. For a packet to match a rule, both the source and destination address of the packet must fall within the configured source and destination address ranges of the rule.

IPv6 address ranges may be specified in the following formats:

- ▶ as a single IPv6 address
ex. 5001:0123:cccc::1
- ▶ as a range of IPv6 addresses represented by an address and prefix length
ex. 5001:0123:cccc::/48
- ▶ using the wildcard 'any' to match any IPv6 address

The standard abbreviation "::" can be used for one or more 16 bit all zero quantities.

Non-IPv6 packets and IPv6 packets that do not match any of the numbered rules follow the separately configured Default Action.

For classification based on the Traffic Class field in the IPv6 header, use the Combination Classifier instead.

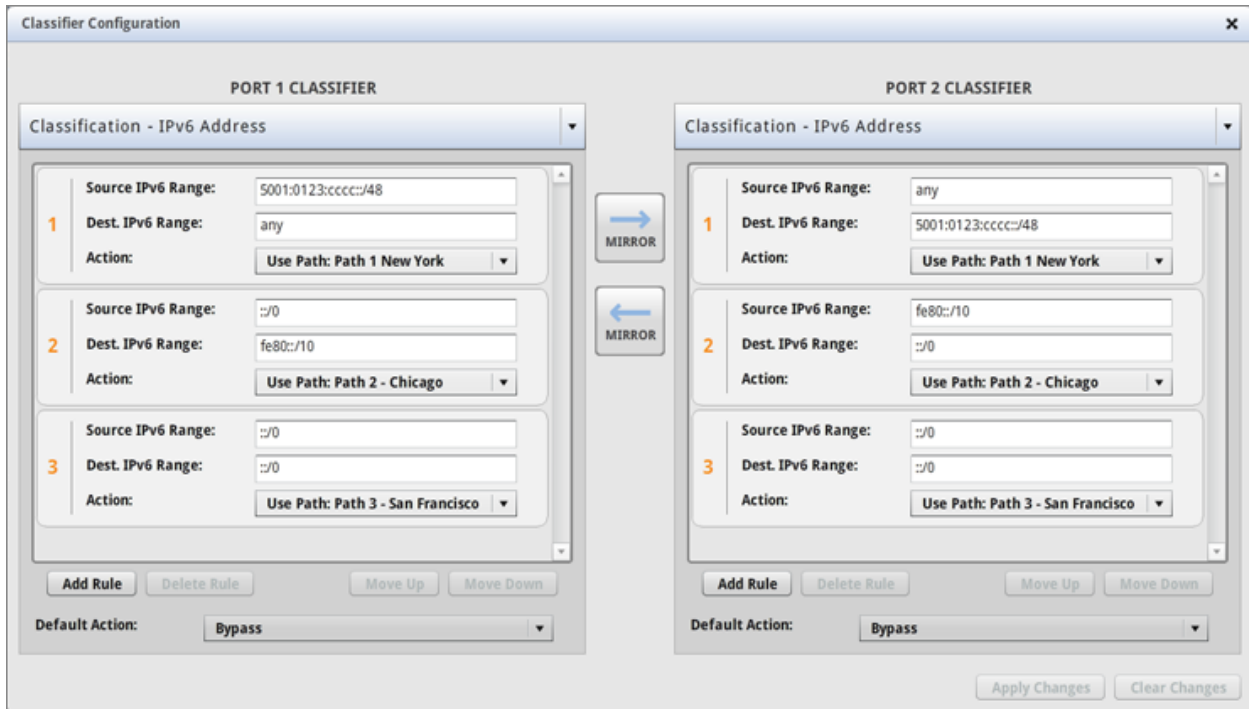


Figure 35: Classification by IPv6 Address

7.2.4 VLAN Classification

VLAN classification matches packets by the VLAN ID and PCP value in the IEEE 802.1Q VLAN tag.

The VLAN ID may be specified as a single value between 0 and 4095, or as a range of values separated by a hyphen. The wildcard 'any' can be used to match any value.

The PCP (priority code point) is a 3-bit field in the VLAN tag used for prioritization. The VLAN PCP is specified as a single value between 0 and 7, or the wildcard 'any'. To ignore the VLAN PCP, use 'any' so that any PCP value will match.

VLAN classification matches the 802.1Q VLAN header of which there can be only one in any packet. Any stacked service VLAN (SVLAN) 802.1ad headers are ignored by the VLAN classifier, but can be used for classification with the raw data filters of the Combination Classifier.

Packets that do not have a VLAN tag or do not match any of the numbered rules follow the configured Default Action.

The image shows a 'Classifier Configuration' window with two main sections: 'PORT 1 CLASSIFIER' and 'PORT 2 CLASSIFIER'. Both sections have a 'Classification - VLAN' dropdown menu. Below each menu is a list of four rules, numbered 1 to 4. Each rule has three fields: 'VLAN ID Range', 'VLAN PCP', and 'Action'. The 'VLAN ID Range' values are 5, 7-10, 1-99, and 100-149 respectively. The 'VLAN PCP' values are all 'any'. The 'Action' values are 'Use Path: Path 1 New York', 'Use Path: Path 2 - Chicago', 'Use Path: Path 3 - San Francisco', and 'Drop' respectively. Between the two classifier sections are two 'MIRROR' buttons with arrows pointing in opposite directions. At the bottom of each classifier section are buttons for 'Add Rule', 'Delete Rule', 'Move Up', and 'Move Down'. A 'Default Action' dropdown menu is located at the bottom of each section, set to 'Bypass'. At the very bottom of the window are 'Apply Changes' and 'Clear Changes' buttons.

Port	Rule #	VLAN ID Range	VLAN PCP	Action
PORT 1 CLASSIFIER	1	5	any	Use Path: Path 1 New York
	2	7-10	any	Use Path: Path 2 - Chicago
	3	1-99	any	Use Path: Path 3 - San Francisco
	4	100-149	any	Drop
PORT 2 CLASSIFIER	1	5	any	Use Path: Path 1 New York
	2	7-10	any	Use Path: Path 2 - Chicago
	3	1-99	any	Use Path: Path 3 - San Francisco
	4	100-149	any	Drop

Figure 36: Classification by VLAN

7.2.5 MPLS Label Classification

MPLS label classification matches packets by the MPLS label of each packet.

The MPLS label may be specified as a single value between 0 and 1048575, or as a range of values separated by a hyphen. The wildcard 'any' can be used to match any value.

Stacked MPLS labels are supported, but MPLS label classification is always based on the outermost label (the label acted on by the next router in line). Other MPLS labels can be used for classification with the raw data filters of the Combination Classifier.

Packets that do not have an MPLS label or do not match any of the numbered rules follow the configured Default Action.

The image shows a 'Classifier Configuration' window with two main sections: 'PORT 1 CLASSIFIER' and 'PORT 2 CLASSIFIER'. Both sections have a dropdown menu set to 'Classification - MPLS Label'. Each section contains a list of five rules, numbered 1 to 5. The rules are as follows:

Rule Number	MPLS Label Range	Action
1	10	Use Path: Path 1 New York
2	20-29	Use Path: Path 2 - Chicago
3	50	Use Path: Path 3 - San Francisco
4	1-99	Use Path: Path 4 - London
5	100-999	Drop

Below the rules, there are buttons for 'Add Rule', 'Delete Rule', 'Move Up', and 'Move Down'. A 'Default Action' dropdown is set to 'Bypass'. In the center, between the two classifier panels, are two 'MIRROR' buttons with arrows pointing in opposite directions. At the bottom right of the window are 'Apply Changes' and 'Clear Changes' buttons.

Figure 37: Classification by MPLS Label

7.2.6 MAC Address Classification

MAC address label classification matches packets by the source and destination MAC address of each packet.

The MAC address is specified as six groups of hexadecimal digits separated by hyphens. The wildcard 'any' can be used to match any MAC address.

Packets that do not match any of the numbered rules follow the separately configured Default Action.

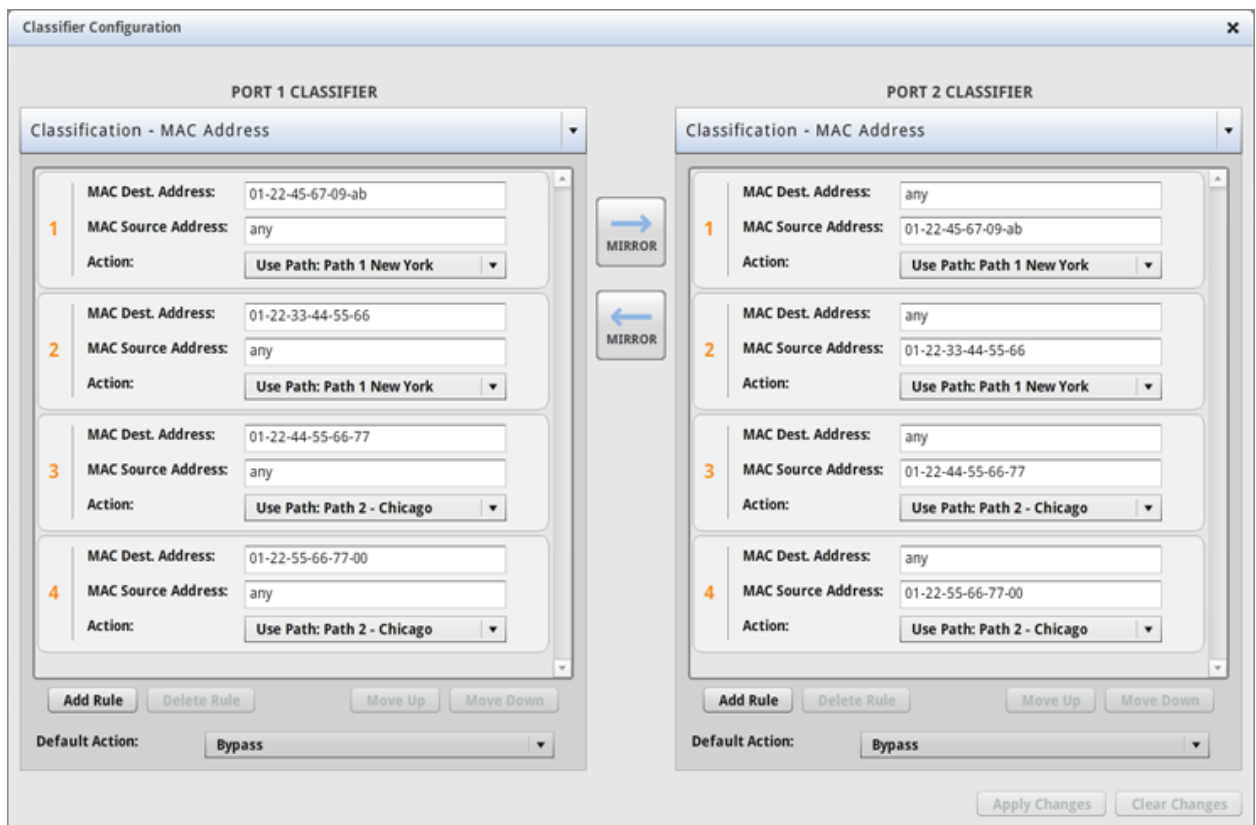


Figure 38: Classification by MAC Address

7.2.7 TCP/UDP Port Classification

TCP/UDP Port classification matches packets by the combination of source and destination IP address and TCP or UDP port number. The TCP/UDP Port classification is generally used to apply impairments selectively to traffic from one or more specific applications while allowing other application traffic to be forwarded without impairment.

For a packet to match a rule, the IP address, transport layer protocol (TCP or UDP), and application port number must all match the specified values for both source and destination.

Use the Protocol radio button to specify either a TCP or UDP packet. Each rule must be either TCP or UDP.

The source and destination addresses can be any valid IPv4 or IPv6 address or address range.

- ▶ as a single IPv4 or IPv6 address
ex. 10.0.0.10
ex. 5001:0123:cccc::1
- ▶ as a range of IPv4 addresses separated by a dash
ex. 10.0.0.0 – 10.255.255.255

- ▶ as a range of IPv4 or IPv6 addresses represented by an address and prefix length
ex. 10.0.0.0/24
ex. 5001:0123:cccc::/48
- ▶ using the wildcard 'any' to match any IPv4 or IPv6 address

The port number is a single value or range of values separate by a dash between 0 and 65535, or the wildcard 'any' to match any port number.

Non-IP packets and IP packets that do not match any of the numbered rules follow the separately configured Default Action.

The screenshot displays the 'Classifier Configuration' window, which is divided into two main sections: 'PORT 1 CLASSIFIER' and 'PORT 2 CLASSIFIER'. Each classifier has a 'Classification - TCP/UDP Port' dropdown menu. Between the two classifiers are 'MIRROR' buttons with arrows pointing in opposite directions.

PORT 1 CLASSIFIER:

- Rule 1:** Protocol: TCP (selected), Source Address: 10.0.0.155, Port: 80, Dest. Address: any, Port: any, Action: Use Path: Path 1 - New York.
- Rule 2:** Protocol: TCP (selected), Source Address: 192.168.0.155, Port: any, Dest. Address: any, Port: any, Action: Use Path: Path 2 - Chicago.
- Rule 3:** Protocol: TCP (selected), Source Address: any, Port: any, Dest. Address: 192.168.0.0/24, Port: 25, Action: Use Path: Path 3 - San Francisco.
- Rule 4:** Protocol: UDP (selected), Source Address: any, Port: 50243, Dest. Address: any, Port: any, Action: Use Path: Path 4 - London.

PORT 2 CLASSIFIER:

- Rule 1:** Protocol: TCP (selected), Source Address: any, Port: any, Dest. Address: 10.0.0.155, Port: 80, Action: Use Path: Path 1 - New York.
- Rule 2:** Protocol: TCP (selected), Source Address: any, Port: any, Dest. Address: 192.168.0.155, Port: any, Action: Use Path: Path 2 - Chicago.
- Rule 3:** Protocol: TCP (selected), Source Address: 192.168.0.0/24, Port: 25, Dest. Address: any, Port: any, Action: Use Path: Path 3 - San Francisco.
- Rule 4:** Protocol: UDP (selected), Source Address: any, Port: any, Dest. Address: any, Port: 50243, Action: Use Path: Path 4 - London.

At the bottom of each classifier section are buttons for 'Add Rule', 'Delete Rule', 'Move Up', and 'Move Down'. A 'Default Action' dropdown menu is set to 'Bypass' for both classifiers. At the very bottom of the window are 'Apply Changes' and 'Clear Changes' buttons.

Figure 39: Classification by TCP or UDP Port Number

7.2.8 Combination Classification

The Combination Classifier is a flexible classifier that allows packets to be classified using multiple types of rules, such as both IPv4 and IPv6 addresses, or to create complex rules combining multiple well-known fields such as a specific IP address range within one VLAN. In addition, the Combination Classifier includes a Raw Data classification feature which allows packets to be identified and classified based on any arbitrary matching criteria.

Each rule of the Combination Classifier consists of one or more filters and an action. Configured filters are shown in summary form in the Classifier Configuration screen. To build a rule, or to modify the filters of an existing rule, click on the *edit* link within the filter's summary box to open the Combination Rule Editor window.

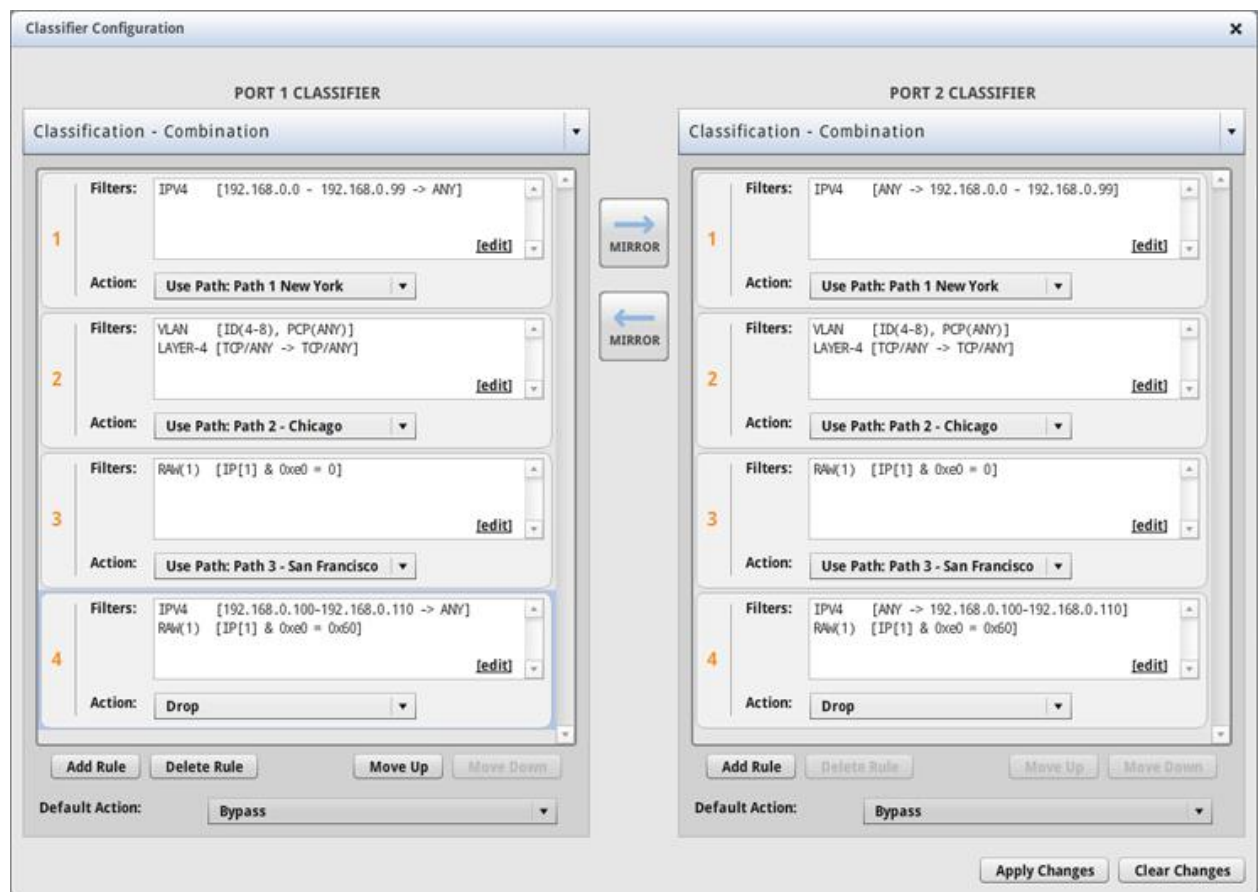


Figure 40: Combination Classification

The Combination Rule Editor window has separate tabs for each packet layer, along with a Raw Data filter tab. To add a filter, click on the tab for the desired packet layer, check the enable box, and enter the criteria to match. Valid entries for each layer are the same as described earlier in this chapter for the simple classifiers. Click the *accept* button when complete.

The Combination Classifier allows filtering on multiple criteria, for example, packets on VLAN 12 with an IP source address between 192.168.0.100 to 192.168.0.110. For this example, click the "IPv4/IPv6 Header" tab, check the *enable* box to turn on IP layer filters, select address type of IPv4, set the source address of 192.168.0.100 – 192.168.0.110 and set the destination address to "any". Then click on the VLAN Header tab, check the *enable* box, and enter the VLAN ID range of 12 with a PCP value of "any".

The screenshot shows the "Combination Rule Editor" window. On the left is a vertical stack of tabs: "MAC Address", "VLAN Header", "MPLS Label", "IPv4/IPv6 Header" (which is selected and highlighted with a blue border), "L4 Protocol/Port", and "Raw Data". To the right of these tabs is a configuration area. At the top of this area is a checked checkbox labeled "Enable". Below it are four fields: "Address Type" with radio buttons for "IPv4" (selected) and "IPv6"; "Source Range" with a text box containing "192.168.0.100-192.168.0.110"; "Destination Range" with a text box containing "any"; and "Type of Service" with a dropdown menu set to "Precedence" and a text box containing "1". Below the configuration area is a section titled "Current Configuration" containing a table with two columns: "Layer" and "Filter".

Layer	Filter
VLAN	ID(12), PCP(ANY)
IPV4	192.168.0.100-192.168.0.110 -> ANY, Precedence 1

At the bottom right of the window are two buttons: "Accept" and "Cancel".

Figure 41: Combination Rule Editor

The Raw Data classifier allows any portion of the packet contents to be used for filtering. The layer, offset, and length values together specify the packet data on which to filter. This data is combined in a bitwise AND with the supplied mask, and the result is compared with the supplied range.

Combination Rule Editor

MAC Address
VLAN Header
MPLS Label
IPv4/IPv6 Header
L4 Protocol/Port
Raw Data

☒ Enable

Layer: IP Header

1
Offset: 1
Length: 1
Mask: 0xe0
Range: 0x20-0x40

Layer: IP Header

2
Offset: 1
Length: 1
Mask: 0xe0
Range: 0x80

Add Filter
Remove Filter

Current Configuration

Layer	Filter
IPv4	192.168.0.100-192.168.0.110 -> ANY
RAW(1)	IP[1] & 0xe0 = 0x20-0x40
RAW(2)	IP[1] & 0xe0 = 0x80

Accept
Cancel

Figure 42: Raw Data Rules

If the specified layer does not exist, the filter does not match. For example, if TCP Header layer is selected, but the packet is a UDP datagram, the match will fail.

The offset can be any arbitrary length and can extend past the end of the header. Selecting the Ethernet layer allows selection of any set of bytes from the beginning of the packet. If the length extends beyond the end of the packet, the match will fail.

Each filter that has been created for the rule is listed in the Current Configuration. There can be filters for each of the five layers, plus up to eight Raw Data filters, though it is unlikely that filters for more than two different layers would be needed.

For a packet to match a rule, all filters must match. If any filter does not match the specified criteria, the rule is ignored and the next rule in order is tested for a match. The action specified in the first rule that matches all of the filter criteria is used.



Matching on multiple filters using the combination classifier is processor intensive and can reduce the maximum packet processing rate of the system.

7.3 Rule Order

Rules may overlap each other. The first matching rule in numerical order is used. The *Move Up* and *Move Down* buttons are used to adjust rule order. The default action specifies what happens to packets that do not match any numbered rule.

Example: consider the following set of IP Address Classification rules for Port 1:

Rule 1	Source IP Range:	10.0.0.10
	Destination IP range:	any
	Action:	Use Path: Link A
Rule 2	Source IP Range:	10.0.0.0/24
	Destination IP range:	any
	Action:	Use Path: Link B
Rule 3	Source IP Range:	10.0.0.0 – 10.255.255.255
	Destination IP range:	any
	Action:	Use Path: Link C
	Default Action:	bypass

In this example, any packet that arrives on Port 1 with source address of 10.0.0.10 will be sent over Link A. Any other packet with source address on the 10.0.0.x subnet will be sent over Link B. And any packet with a source address on the 10.x.x.x subnet that is not on the 10.0.0.x subnet will be sent over Link C.

All IP packets not on the 10.x.x.x subnet and all non-IP packets will be forwarded directly to the opposite port, bypassing the emulated WAN network.

8 RECORDINGS

8.1 Overview of Recordings

Path delay, loss, and bandwidth conditions that fluctuate over time can be simulated using a time-series of values from a user-supplied recording file usually generated by the Netropy Recorder for Windows and Linux or the Apposite Recorder for Android™ software. This provides a convenient method of capturing the conditions of a live network link and reproducing those conditions in the lab.

The Netropy Recorder is an application for Windows and Linux available for free download from the Apposite website. The Apposite Recorder for Android provides similar functionality on Android-based phones and tablets and is available for free download on Google Play™. The Netropy or Apposite Recorder sends ICMP Echo Requests from the device on which it is running to a specified destination on the other side of the network, then records the reported delay and loss values to a recording file. For more information on the Recorder software, please consult the *Netropy Recorder User's Guide* or the Help text within the Apposite Recorder for Android.

Users can also create their own recording files to generate any desired time sequence of emulated conditions. Although the Recorder software does not record bandwidth, the recording file can include values for bandwidth as well as latency and loss. To create a recording file, or to edit a recording file created by the Recorder, refer to the Recording File Format description in Section 8.4.

The use of recordings is a 2-step process. First, recording files are loaded into the Netropy network emulator for use by any Path of any engine. Then, available recordings can be selected for playback within the configuration of any Path.

Recordings represent the conditions for a single direction, and consequently, all data represent one-way, not round trip values.

8.2 Managing Recordings

Use the Recordings tab of the Administration window to add and delete recording files.

To add a new recording file, click the *add* button and browse to the recording file stored on the management PC or a locally accessible file server. To remove a recording file, select the file and click the *delete* button.

Selecting a previously loaded recording displays a summary of the delay and loss characteristics of the data, as well as the length of the recording. Clicking the *preview* link displays a graph of the data. Separate graphs will be displayed for each condition included in the file.

Each recording file can be up to 10 MB in size, with the total for all packet capture and recording files limited to 40 MB.

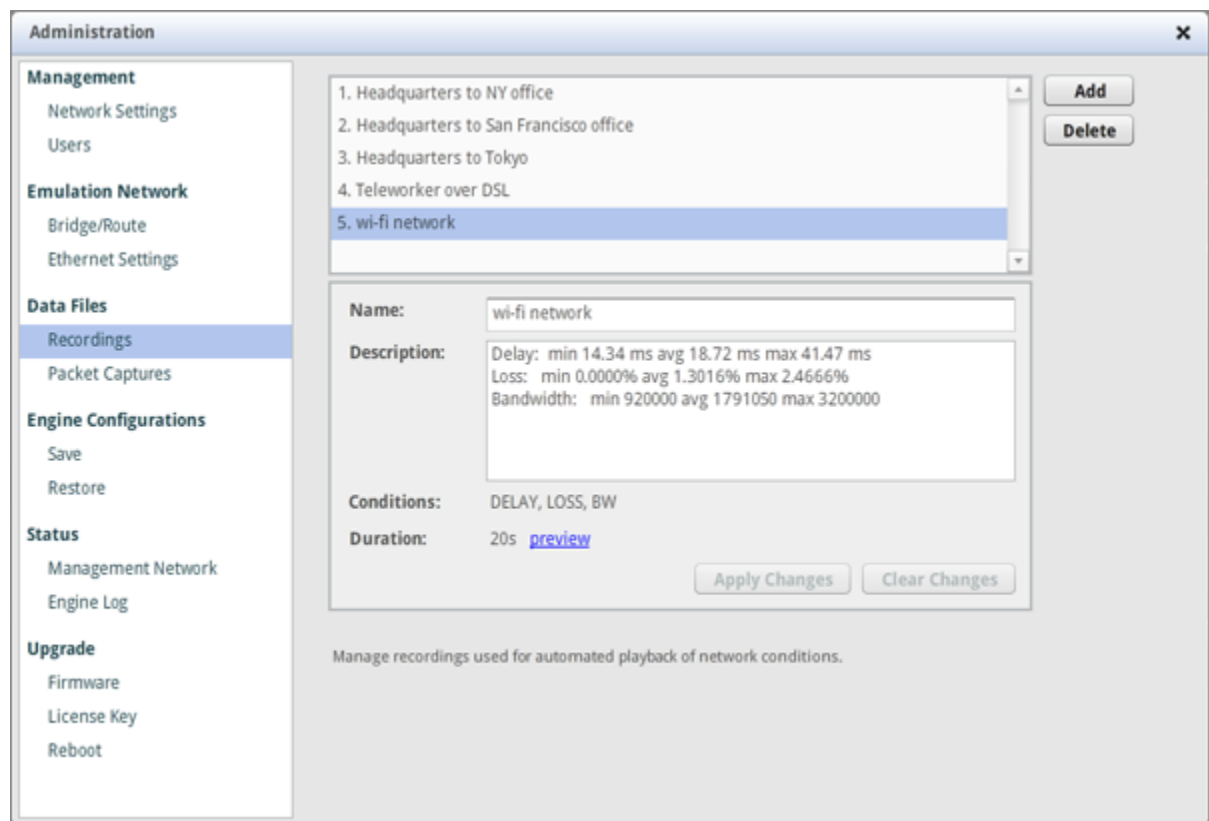


Figure 43: Recording Administration

8.3 Recording Playback

To playback a recording, first turn off emulation for the engine. Playback can only be configured while emulation is off. Then click on the Path to open the Path Configuration window. Click the Playback button on the top button bar to configure recording playback.

Select the recording to use for each direction. Playback is configured separately in each direction. The same recording can be used in each direction, different recordings can be used in each direction, or a recording can be used in one direction with non-recorded conditions in the other direction.

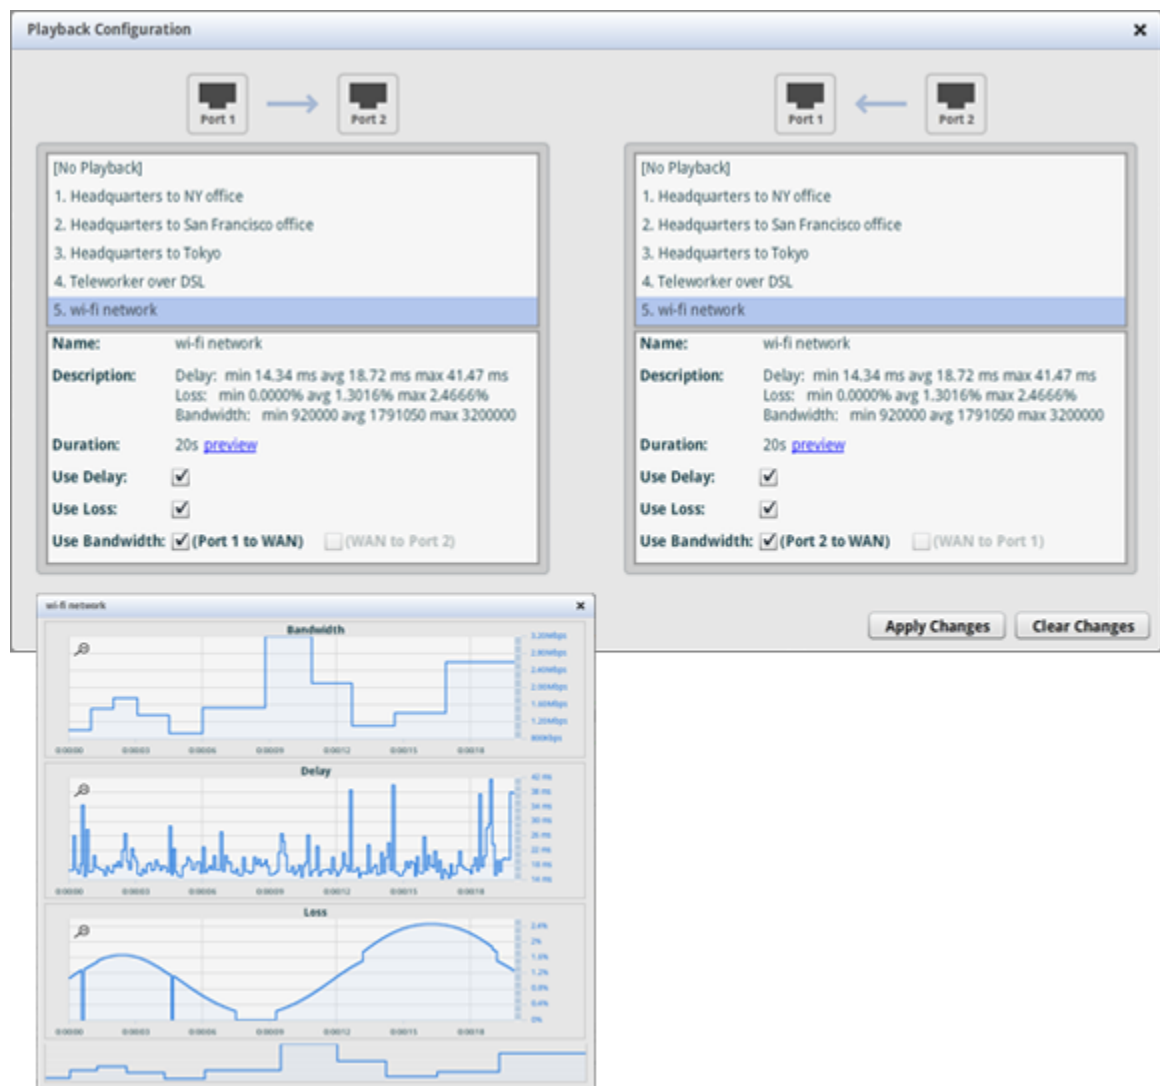


Figure 44: Recording Playback Configuration

Click the check boxes to select whether to use recorded delay, loss, and/or bandwidth. Any parameters not included in the recording are grayed out. By default, all available conditions are selected.

For each path direction, recorded bandwidth can be used to set bandwidth from a port outbound to the WAN, from the WAN inbound to a port, or both. The option to use recorded bandwidth from the WAN inbound to a port will be enabled only if the inbound side of the corresponding WAN Access has been enabled.

Apply changes to return to the Path Configuration window and configure any other emulation parameters if needed.

Playback of the recording begins when emulation is turned on. When playback reaches the end of the recording, it restarts at the beginning and continues looping until emulation is disabled.

To view the progress of the recording playback, return to the Playback Configuration window. While playback is running, a graph of the recording is displayed showing the progress of the playback.

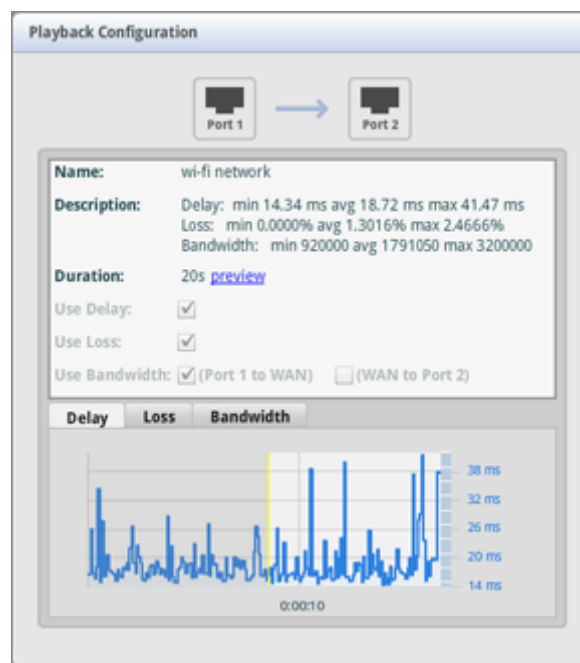


Figure 45: Recording Playback Progress

8.4 Recording File Format

Recordings are text files with the format described in the table below. The recording file may contain up to 100,000 lines of data.

Line Type	Syntax and Description
metadata	<pre># <attribute> : <value></pre> <p>A list of attributes of the recording and their values. Must be at the beginning of the file, prior to any recording data, and preceded by a “#” symbol. Defined attributes are:</p> <pre># name (or # title) : <recording name/title></pre> <p>The title of the recording displayed above the recording graphs.</p> <pre># description : <recording description></pre> <p>A detailed description of the recording. If no description is included in the file, the computed min/avg/max delay and loss values will be displayed as the description of the recording.</p> <pre># contents : <impairment 1, impairment 2, ...,impairment n></pre> <p>An ordered, comma-separated list of impairment names that defines the fields of the data section. Valid impairment names are limited to BW, DELAY, and LOSS. For example:</p> <pre># contents : DELAY,LOSS # contents : DELAY # contents : BW,DELAY,LOSS</pre> <p>If the contents are not specified, "DELAY,LOSS" is assumed.</p>
Data	<pre><start time> <impairment 1 value> [<impairment 2 value>] [<impairment 3 value>]</pre> <p>Start time is the time in seconds, relative to the beginning of the recording, when the delay and loss values of the line take effect. Start time is a floating point number and must be at least 0.001 seconds greater than the start time of the previous line. There cannot be more than 10 lines in any one second period. On the final line in the recording file, the start time is used only to determine the duration of the previous line with delay and loss values ignored.</p> <p>Delay values are specified as latency in milliseconds with a resolution of 0.01 ms.</p> <p>Loss values are specified as the packet loss rate in percent (without the percent sign) between 0.0000 and 100.0000.</p> <p>Bandwidth values are specified in bits per second.</p>

8.5 Example Recording File

The following is a short example recording file with bandwidth, delay, and loss parameters. Within the first second, the loss values change five times, then remain constant for 19 seconds, then the bandwidth and delay values change over the next 30 seconds.

```
# Name : Example Recording
# Contents : BW,DELAY,LOSS
0, 2000000, 52.5, 0
0.2, 2000000, 52.5, 25.0
0.4, 2000000, 52.5, 50.0
0.6, 2000000, 52.5, 75.0
0.8, 2000000, 52.5, 50.0
1.0, 2000000, 52.5, 25.0
20.0, 2000000, 52.5, 0
25.0, 4000000, 52.5, 0
30.0, 8000000, 60.0, 0
40.0, 8000000, 75.0, 0
50.0, 8000000, 100.0, 0
```


9 MONITORING & STATISTICS DOWNLOAD

The GUI displays real-time statistics and throughput graphs for the traffic over each of the emulated links.

9.1 Graphs

The data visualization section of the main window displays a graph of any of the traffic statistics, such as throughput, over any path. Use the drop-down menus to select the statistic and path to view. The graph can display the statistics in either direction of a path or overlay both directions.

Statistics for background traffic is shown separately from real traffic entering the Engine from external sources.



Figure 46: Statistics Graph

To zoom on any portion of the graph, place the cursor over the edge of the area of interest and click and drag to zoom onto that segment. Use the slider under the graph to pan to earlier or later time periods, and use the *zoom out* buttons to reduce the zoom level. When fully zoomed out, the graph displays the previous two hours.



A 24 hour graph of any statistic with zoom and pan controls is also available by clicking the Download Statistics button in the statistic panel, then clicking the Time Picker button.

9.2 Statistics

The statistics section of the main window displays statistics for each path and for bypass traffic. By default, only overall throughput rate, frames and bytes transmitted, and packets dropped are displayed. To view other statistics, click the column configuration icon to the right of the table.

Rates are displayed as averages over the past one second interval. Counters are displayed as cumulative values since the last reset. Rebooting or power cycling the device resets all values. The *reset totals* button resets values displayed in the statistics panel.

	Path	Overall			
		Rate	Frames	Bytes	Drops
→	[Bypass]	0	50	3,200	
←	[Bypass]	0	50	3,200	
→	HQ to Local Office	42,127,536	941,447	1,429,104,256	417
←	HQ to Local Office	992,320	493,957	35,018,784	0
→	HQ to Home via DSL	1,882,320	141,321	214,525,278	663
←	HQ to Home via DSL	43,680	79,244	5,719,076	188
→	HQ to Remote Backup	0	0	0	0
←	HQ to Remote Backup	0	0	0	0

Figure 47: Path Statistics



The *reset totals* button resets statistics for its browser window only. Opening a new browser window or reloading the current page will restore the statistics. Use the *erase statistics* button on the Statistics Selection and Download window to reset all values permanently. A reboot or power cycle of the Netropy unit will also reset all values.



Statistics cannot be recovered after a reboot or power cycle of the device.

Data available for display for each path in each direction are described in the table below.

Segment	Statistic	Description
Overall	Rate	current transmission rate for traffic delivered across the WAN over the previous 1 sec. interval.
	Bytes/Frames	cumulative number of bytes and packets delivered across the path.
	Drops	cumulative sum of packets dropped as a result of queuing limits on both the outbound and inbound WAN Access links and the frames dropped due to configured WAN loss parameters.
WAN Access	Tx Rate	current transmission rate for traffic delivered across the WAN Access link over the previous 1 sec. interval. Does not include background traffic.
	Tx Bytes/Frames	cumulative number of bytes and packets transmitted over the WAN access link. Does not include background traffic.
	Queue Length – Bytes/Frames	number of bytes and packets currently in the transmit queue, including estimated queue occupancy of background utilization traffic, when configured.
	Queue Drops	cumulative number of packets dropped due to configured queuing limits. Does not include drops of background traffic.
	Background Bytes/Frames	cumulative number of bytes and packets injected as background traffic on the WAN Access link.
	Background Queue Drops	cumulative number of packets of background traffic dropped due to configured queuing limits.
WAN	Loss Drops	cumulative number of packets dropped due to configured loss parameters.
	Frames Reordered	cumulative number of packets reordered.
	Duplicated	cumulative number of packets duplicated.
	Corrupted	cumulative number of packets that contain one or more errors.

9.2.1 Statistics Download

Depending on the number of configured paths, up to 24 hours of statistics, in 1 second intervals, are available for download to a file. To download, click the *Download Statistics* button to open the Statistics Download window.

In the Statistics Download window, select the paths and statistics. The time period for the statistics can either be entered manually or selected graphically using the Time Picker. The Time Picker displays a throughput graph of up to 24 hours with zoom and pan controls to assist in finding and selecting periods of interest.

After selecting the time period, click on the download button for per-second interval counters or cumulative counters, and choose a location to save the file.

The data can be downloaded as either:

- ▶ Per-second interval counters: separate values of the number of packets, bytes, or other counters within each one second interval
- ▶ Per-second cumulative counters: cumulative values for the number of packets, bytes, or other counters since the beginning of the selected time period reported each second.

Statistics Selection and Download

Path Statistics

Paths to Include	Overall	Port 1 WAN Access	WAN	Port 2 WAN Access
<input type="checkbox"/> All	<input type="checkbox"/> Rate	<input type="checkbox"/> Tx Rate	<input type="checkbox"/> Loss Drops	<input type="checkbox"/> Tx Rate
<input type="checkbox"/> [Bypass]	<input type="checkbox"/> Frames	<input type="checkbox"/> Tx Bytes	<input type="checkbox"/> Frames Reordered	<input type="checkbox"/> Tx Bytes
<input type="checkbox"/> Path 1 New York	<input type="checkbox"/> Bytes	<input type="checkbox"/> Tx Frames	<input type="checkbox"/> Frames Duplicated	<input type="checkbox"/> Tx Frames
<input type="checkbox"/> Path 2 - Chicago	<input type="checkbox"/> Drops	<input type="checkbox"/> Queue Length - Bytes	<input type="checkbox"/> Frames Corrupted	<input type="checkbox"/> Queue Length - Bytes
<input type="checkbox"/> Path 3 - San Francisco		<input type="checkbox"/> Queue Length - Frames		<input type="checkbox"/> Queue Length - Frames
		<input type="checkbox"/> Queue Drops		<input type="checkbox"/> Queue Drops
		<input type="checkbox"/> Background Frames		<input type="checkbox"/> Background Frames
		<input type="checkbox"/> Background Bytes		<input type="checkbox"/> Background Bytes
		<input type="checkbox"/> Background Queue Drops		<input type="checkbox"/> Background Queue Drops

Time Range

2012/11/21 15:52:06 - 2012/11/21 16:52:06 **Time Picker**

Per-second interval statistics: **Download**

Per-second cumulative statistics: **Download**

Erase Statistics

Figure 48: Statistics Download

The downloaded statistics are saved as a comma separated value (CSV) file that can be imported into Microsoft Excel or other data visualization application.

The statistics file consists of comma separated values. The first three rows are headings, followed by a separate row for each one second time interval during the selected period.

The three heading rows identify the contents of columns.

Within each row of data, the first two columns show the number of seconds since the last statistics reset or power cycle and the corresponding clock time on the management device that downloaded the file.

For WAN statistics, the direction of packet flow is indicated using the port numbers. For example, "Port 1 - Port 2" indicates traffic flowing from the Port 1 side to the Port 2 side.

For WAN Access statistics, the direction of packet flow is indicated by the port number and whether it is in-bound or out-bound from that port. For example, "Port 1 WAN Access Outbound" indicates packets flowing from the LAN attached to Port 1 out to the WAN.

The statistic name identifies the particular statistic. Interval statistics are followed by "(I)", and cumulative statistics are followed by "(C)".

10 ADMINISTRATION

The Netropy network emulator is administered via a separate window accessed by clicking the *Administration* link at the top of the main page.

10.1 Network Settings

The IP address of the management interface of the Netropy device and other network settings are configured in the Network Settings tab.

The screenshot shows the 'Administration' window with the 'Network Settings' tab selected. The left sidebar contains a tree view with categories: Management (Network Settings, Users), Emulation Network (Bridge/Route, Ethernet Settings), Data Files (Recordings, Packet Captures), Engine Configurations (Save, Restore), Status (Management Network, Engine Log), and Upgrade (Firmware, License Key, Reboot). The main content area is titled 'Management > Network Settings' and contains the following fields:

- Hostname: netropy
- Domain: test.net
- DHCP: Disabled (dropdown)
- IP Address: 10.0.0.10
- IP Netmask: 255.0.0.0
- Default Gateway: 10.0.0.1
- IPv6 Autoconf: Disabled (dropdown)
- DNS Servers: 10.0.0.2 (with two empty input fields below)
- NTP Servers: pool.ntp.org (with two empty input fields below)

At the bottom of the form are 'Apply Changes' and 'Clear Changes' buttons. A note at the very bottom states: 'These fields control the network settings for the management interface. If the IP address or DHCP setting is changed, your browser connection must be reestablished using the new address.'

Figure 49: Network Settings Tab

IPv4 address, netmask, and default gateway, as well as DNS and NTP servers may be configured manually or using DHCP. A default gateway is optional.

All IPv4 addresses are entered in dotted-decimal notation. Entries are checked for validity and consistency before changes are applied.

The device can also be managed using IPv6. When IPv6 Autoconf is enabled, a link-local IPv6 is created and the device attempts to obtain a global IPv6 address and default gateway. The assigned IPv6 address is shown on the Management Network status tab.

When the IP address is changed, connectivity to the device will be lost and must be reestablished using the new address.

Up to 3 DNS servers and up to 3 NTP servers may be configured. DNS or NTP servers configured manually will override any servers set automatically through DHCP.



If you cannot regain connectivity to the device after changing the network settings, use the CONSOLE interface to verify or change the network settings.

10.2 Users

Username and passwords are administered in the Users tab.



Figure 50: User Administration Tab

Initially, the device has a single user, *admin*, with no password. Additional users can be added or deleted by *admin*. Passwords for each user can be set or cleared by the individual user or by *admin*.

Users other than *admin* are unprivileged, and can make configuration changes but cannot make system administrative changes.

Username and passwords created in either the GUI or through the CLI apply to access to both the GUI and CLI.

10.3 Bridge/Route

The Netropy emulator can be installed as either a bridge or router to forward frames between the two ports of each Emulation Engine. By default, each Engine is configured as a bridge, and this mode is recommended for simplicity unless the two ports need to be on separate subnets.

Use the Forwarding Mode drop-down selector on the Bridge/Route tab to choose between bridging and routing. This selection is made separately for each Emulation Engine.

Bridge/Route settings are saved with the Engine configuration and are updated when a saved Engine configuration is restored.

10.3.1 Bridging

In Bridging Mode, the Netropy Engine functions as a bridge between the Ethernet segments connected to the two Ethernet ports. In this mode, it can forward any Ethernet-based frame regardless of network layer protocol.




Figure 51: Bridging

By default, all frames are assumed to be part of the WAN traffic and are subjected to the configured emulation conditions. This includes ARP packets, which on a production network may be processed or filtered prior to traversal of the WAN. Check the *ARP Frames Bypass Emulation* box to have those packets forwarded directly between the two ports with no impairment regardless of the configuration of the Emulation Engine.

10.3.2 Routing

In Routing Mode, the Netropy Engine functions as a router between the Ethernet segments connected to the two ports of the Engine. Configure the IP address and netmask of the two interfaces. If necessary, add static routes to off-link destinations. All addresses are entered in dotted-decimal notation.

Routing Mode supports only the forwarding of IPv4 frames and does not support multicast forwarding.

The screenshot shows the 'Administration' window with the 'Bridge/Route' option selected in the left sidebar. The 'Forwarding Mode' is set to 'Routing'. The configuration is for 'Engine 2'. The 'Port 3 Interface' is set to '10.0.1.1' with a netmask of '255.255.255.0'. The 'Port 4 Interface' is set to '10.0.2.1' with a netmask of '255.255.255.0'. Below these, there is a table for static routes. The table has three columns: 'DESTINATION ADDRESS', 'DESTINATION MASK', and 'GATEWAY'. There are two rows of static routes: one for '10.0.10.0' with mask '255.255.255.0' and gateway '10.0.1.2', and another for '0.0.0.0' with mask '0.0.0.0' and gateway '10.0.2.2'. There are 'Add', 'Delete', 'Apply', and 'Cancel' buttons.

DESTINATION ADDRESS	DESTINATION MASK	GATEWAY
10.0.10.0	255.255.255.0	10.0.1.2
0.0.0.0	0.0.0.0	10.0.2.2

Figure 52: Routing

10.4 Ethernet Settings

The speed, duplex, and flow control settings for the Ethernet ports used for emulation are configured in the Ethernet Settings tab.

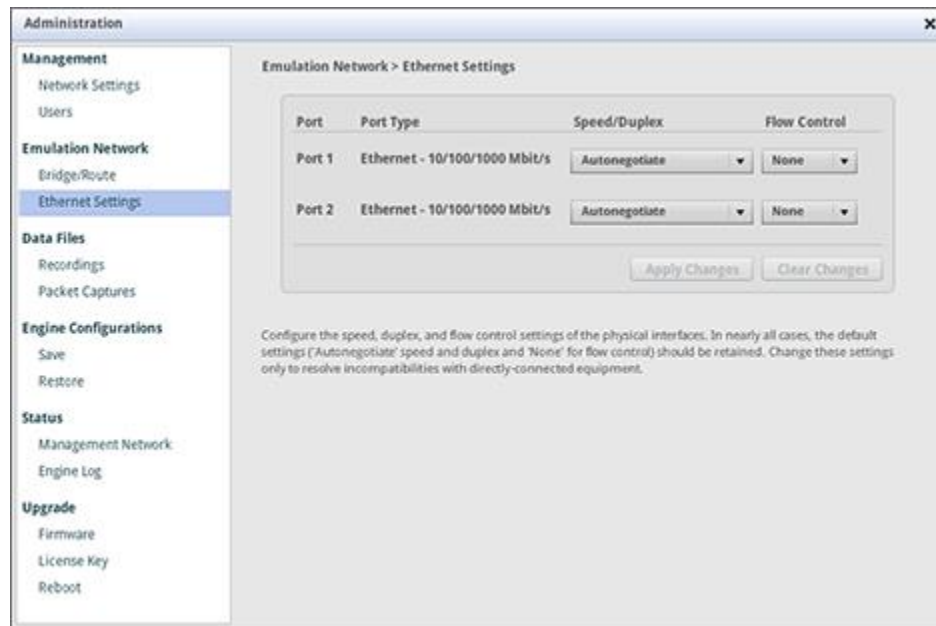


Figure 53: Ethernet Settings Tab

By default, the Ethernet ports are set to auto-negotiate the proper speed and duplex settings, and flow control is turned off. On the 10/100/1000baseT emulation ports, auto-negotiation can be disabled and the ports forced to a particular speed and duplex setting via the drop-down menu. In nearly all cases, the default settings should be used and should only be changed to resolve incompatibilities with directly-connected equipment.

Auto-negotiation cannot be disabled on 1 Gbps SFP ports or any 10 Gbps or higher speed ports, although flow control can be turned on if needed. Auto-negotiation cannot be disabled on the MGMT port.

Jumbo frames of up to 9 KB are supported on all emulation ports.



If the speed and duplex setting of an interface is selected manually, the device the port is connected to must also be forced to the same setting.

10.5 Recordings

Importing and deleting Recording files for automated playback of a time series of latency and loss conditions is managed through the Recordings tab. For more details on the use of recordings, see Section 8: Recordings.

10.6 Packet Captures

Importing and deleting PCAP packet capture files for generation of background traffic is managed through the Packet Captures tab. For more details on the use of packet capture files, see Section 6.2.3.2: Packet Replay.

10.7 Save and Restore Engine Configurations

The configuration of a selected Emulation Engine can be saved to a file on the management PC from the Save tab. The stored configuration file can then be loaded into any Emulation Engine from the Restore tab of the same Netropy unit or onto a different Netropy unit regardless of model. The restore operation overwrites the current configuration of the Engine.

Recording and PCAP files are included in the configuration file if they are used in any path in the Engine. This can cause configuration files to be very large.

The bridging or routing settings are saved with the Engine configuration and are updated when a saved Engine configuration is restored.

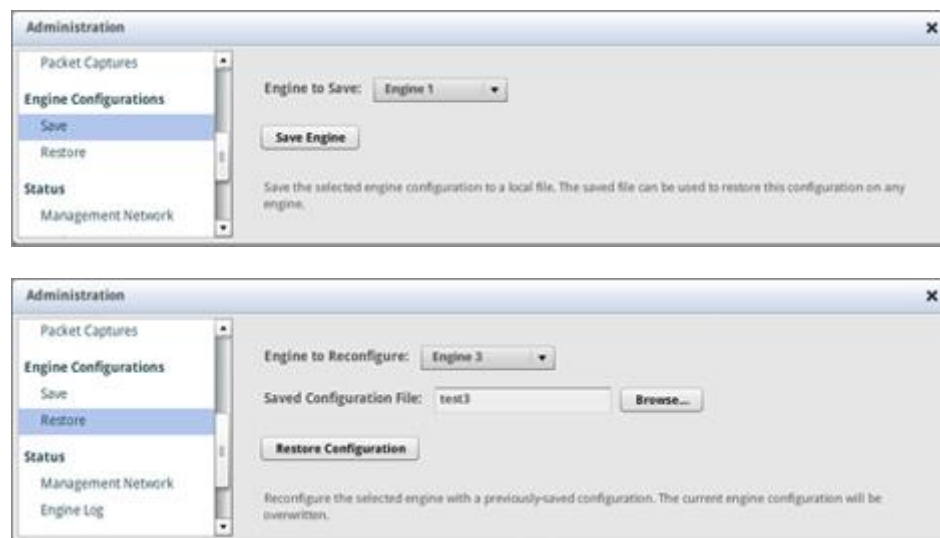


Figure 54: Save and Restore Tabs

10.8 Management Network Status

The Management Network status screen shows the current configuration for the IP address and netmask, default gateway, network domain, and DNS and NTP servers. Press the refresh button to update with the latest status.

10.9 Engine Log

A log file of error messages and warnings is shown on the Engine Log tab. A separate log is maintained for each Engine. Each line includes a sequence number and the time in GMT. Log messages will be generated at most once per second.

If there are any error conditions that could affect the validity of the test results, the LEDs on the Engine tab on the main configuration window remain red until the log message has been marked as read or the log cleared.

Error conditions reported in the log are:

```
timing error exceeded <#>us
```

The engine has detected that the error in emulation timing has exceeded the indicated number of microseconds.

```
<#> frames lost
```

The engine has been overload and was unable to process all received frames, with the indicated number of frames dropped.

```
out of buffers - <#> events
```

The engine ran out of buffer space to receive new frames from the network. The number of failed attempts to allocate buffer space is reported.

```
automatic engine shutdown
```

Emulation was aborted. The current engine configuration requires more resources than supported by the hardware.

```
redundant power supply failure
```

A power supply module has failed and requires replacement. The system is running on a secondary power supply. Applicable only to the Netropy 10G2 with a redundant power supply.

10.10 Firmware

The version of the Netropy firmware currently installed is displayed in the Firmware tab.

To upgrade the firmware, first download the new image to the management PC or a local file server from the Netropy support site. Then use the *browse* button to select the file, and click *upgrade* to install.

The Netropy device will automatically reboot after a successful firmware upgrade.




Figure 55: Firmware Tab



The same procedure can be used to restore an older version of firmware if necessary. However, when downgrading to an older release, some settings may not be preserved.



Firmware updates are available on the Netropy support section of the Apposite website. To obtain access the support site, register your unit at apposite-tech.com/register.html. An active product license or maintenance contract is required. Registered users will also receive email notices of firmware updates.

10.11 License Key

The serial number, license expiration date, and bandwidth limit are displayed in the License Key tab. If an expiration date is not shown, the license does not expire. For information on renewing a license, see apposite-tech.com/renewals.html.

The license key controls the maximum bandwidth that can be configured in either direction for each path.

To add a new license key, use the *browse* button to select the license key file and click *apply key* to install.

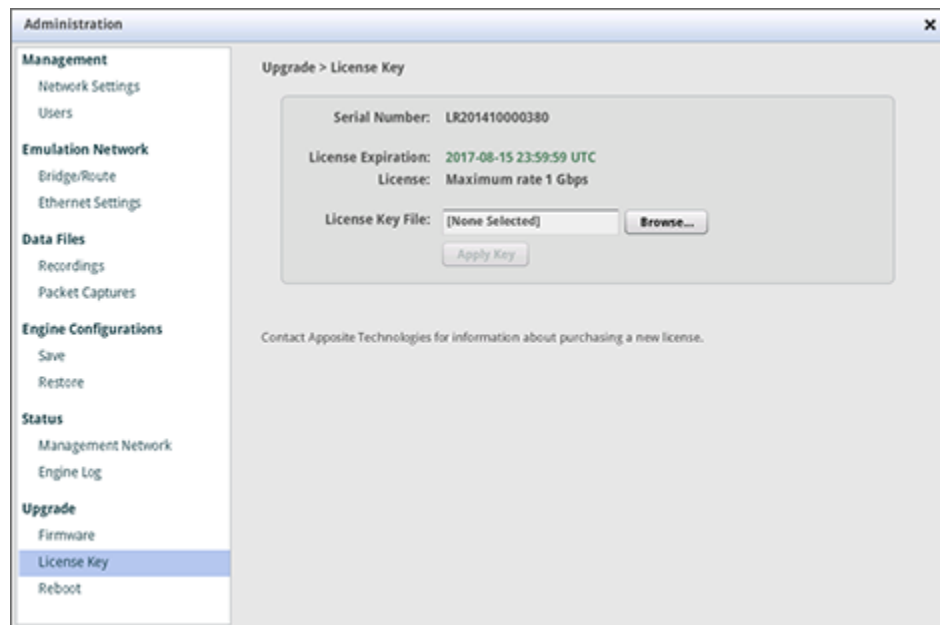


Figure 56: License Key Tab

10.12 Date and Time

System date and time can be set using Network Time Protocol (NTP) (Section 10.1), or using the *clock* command in the CLI (Section 11.3). The date and time are used only for the timestamp on log messages. The time displayed on the graphs and used for statistics download does not use the system time but instead uses the time on the client device being used for management.

11 COMMAND LINE INTERFACE

The Netropy command line interface (CLI) can be accessed via the CONSOLE interface or through a Telnet or SSH connection.

The following features are not available through the CLI and must be completed within the GUI:

- ▶ Addition or removal of paths
- ▶ Modification of packet classification rules
- ▶ Installation of a product license key

The following features are only available through the CLI:

- ▶ ARP and PING commands
- ▶ Manually set system clock
- ▶ Reinitialize the unit to factory settings
- ▶ LDAP authentication
- ▶ Packet capture

To access the CLI, log into the device at the prompt as “admin” or other configured username.

The SSH and Telnet services can be enabled or disabled through the CLI using the `telnet` and `ssh` commands. By default, both are enabled. Multiple simultaneous sessions are allowed.

SSH can be used to execute a single command or to log in for an on-going session similar to Telnet.

11.1 CLI Help

The CLI includes several levels of help. The “help” command by itself displays a list of commands. “help” with a command name displays the syntax and options for the command.

When entering commands, the <Tab> key can be used to complete a partially entered command name. If there are multiple possible completions, pressing the <Tab> key again will display a list of options.

For most commands, typing “?” after a partially entered command provides context-specific help.

11.2 Navigation within the CLI

Engine Command Mode

All emulation configuration commands require an emulation engine to be selected. Use the top-level ‘engine <#>’ command to select an engine and enter engine command mode:

```
[admin@netropy]> engine 1  
[admin@netropy/Engine 1]>
```

While in engine command mode, all engine-level commands are applied to the currently selected engine. Use the ‘exit’ command to exit engine configuration mode and return to the top-level commands, for instance to select a different engine.

```
[admin@netropy/Engine 1]> exit  
[admin@netropy]>
```

Engine commands can also be executed without entering engine command mode by specifying the entire command, including engine number, from the top level.

Path Command Mode

Each emulation engine can be configured with up to 15 sets of emulation conditions called *paths*. Use the ‘path <#>’ command from engine command mode to select a path for configuration and enter path command mode.

```
[admin@netropy/Engine 1]> path 2  
[admin@netropy/Engine 1/Path 2]>
```

While in path command mode, all path-level commands such as changes to emulation conditions are applied to the currently selected path. Use the ‘exit’ command to exit path configuration mode and return to engine command mode, for instance, to select a different path.

```
[admin@netropy/Engine 1/Path 2]> exit  
[admin@netropy/Engine 1]>
```

Path commands can also be executed without entering path command mode by specifying the entire command, including the engine and path numbers, from the top level.

11.3 CLI Top Level Commands

Command	Syntax and Description
arp	<pre>arp show</pre> <p>Displays a table of IP addresses and associated MAC addresses for the MGMT interface.</p>
capture	<pre>capture port <port> [raw detail] [arp ip <address>] [snaplen <len>] [filter <filter string>]</pre> <p>Displays frames received by and sent from the specified port until control-C is pressed. See Section 11.3.2 for a detailed description.</p>
clock	<pre>clock show</pre> <p>Displays the current system time and date.</p> <pre>clock set <hh>:<mm>:<ss> <YYYY>-<MM>-<DD></pre> <p>Sets the time and date used for log messages. All fields are required. Clock settings are permanent and survive reboots and power cycles of the device. If an NTP server is configured, it will override the settings from this <code>clock</code> command.</p>
engine	<pre>engine <#> [emulation {on off} show lock unlock backup restore statistics log autostart path ... router ...]</pre> <p>Enters engine command mode, or executes a command for the specified engine. See 11.4 for syntax of the engine subcommands.</p> <pre>engine list</pre> <p>Displays a list of available emulation engines.</p>
help	<pre>{help ?} [<command>]</pre> <p>Displays a list of available commands or syntax of a specified command. A question mark after any command also displays the syntax for that command.</p>
http	<pre>http load <ssl-certificate-url></pre> <p>Installs an SSL certificate. FTP, HTTP, and TFTP services are supported for upload. Uploaded certificate replaces the self-signed certificate supplied in the firmware. Ex: <code>http load ftp://192.168.0.100/certs/certificate.crt</code></p> <pre>http {on off}</pre> <p>Enables or disables unsecured access to the GUI via HTTP. HTTP service is on by default. Secured access to the GUI via HTTPS is always enabled.</p>
init	<pre>init config</pre> <p>Returns the configuration to factory default settings. Takes effect upon reboot unless the configuration is saved prior to reboot.</p>

logout	<p>logout</p> <p>Logout from the command line interface.</p>
mgmt	<p>mgmt show</p> <p>Displays the IP address and other management information.</p> <p>mgmt show ldap</p> <p>Displays the current LDAP configuration.</p> <p>mgmt show ntp associations</p> <p>Displays the synchronization state with the configured NTP servers.</p> <p>mgmt set addr {dhcp addr <addr> netmask <mask>}</p> <p>Sets the IPv4 address and netmask of the MGMT interface either manually or using DHCP. If DHCP is enabled, DHCP sets the default gateway.</p> <p>mgmt set domain <domain></p> <p>Sets the network domain name of the device.</p> <p>mgmt set gw <addr></p> <p>Sets a default gateway for the MGMT interface.</p> <p>mgmt set hostname <name></p> <p>Sets the hostname of the device.</p> <p>mgmt set ipv6 {on off}</p> <p>Turns on or off IPv6 Autoconfiguration to assign an IPv6 address and default gateway for the MGMT interface.</p> <p>mgmt set ldap ...</p> <p>See Section 11.3.1 for LDAP configuration syntax.</p> <p>mgmt set nameserver <addr> [<addr2> [<addr3>]]</p> <p>Sets up to three DNS servers for the device. This command overrides any nameservers set through DHCP and any DNS servers previously configured.</p> <p>mgmt set ntp server <server> [<server2> [<server3>]]</p> <p>Sets up to three network time protocol servers for the device. This command overrides any NTP servers set through DHCP and any NTP servers previously configured.</p> <p>mgmt clear ldap</p> <p>Turns off LDAP authentication and clears the LDAP configuration.</p> <p>mgmt clear nameserver</p> <p>Clears all manually configured DNS servers. Does not change any DNS servers set through DHCP.</p> <p>mgmt clear ntp server</p> <p>Clears all manually configured NTP servers. Does not change any NTP servers set through DHCP.</p>

packet-captures	<pre>packetcaptures list</pre> <p>Displays a list of imported PCAP files, their ID numbers, and size in bytes.</p> <pre>packetcaptures add <packet-capture-url></pre> <p>Imports a PCAP packet capture file. FTP, HTTP, and TFTP services are supported. Ex: packetcaptures add http://192.168.0.100/captures/voipstream.pcap</p> <pre>packetcaptures delete <id></pre> <p>Deletes the specified packet capture file.</p>
password	<pre>password clear [<user-id>]</pre> <p>Clears the password for the specified user. If no user is specified, command applies to the user executing this command. Only <i>admin</i> can specify a user other than himself.</p> <pre>password set [<user-id>]</pre> <p>Prompts for a new password for the specified user. If no user is specified, command applies to the user executing this command. Only <i>admin</i> can specify a user other than himself.</p>
ping	<pre>ping <ip-address> [<size>]</pre> <p>Pings from the device to the IP <code>address</code> with the specified sized packets using the MGMT interface. Use CTRL-C to stop.</p>
reboot	<pre>reboot</pre> <p>Reboots the device. Returns user to the login prompt after reboot.</p>
recordings	<pre>recordings list</pre> <p>Displays a list of imported recording files and their ID numbers.</p> <pre>recordings add <recording-url></pre> <p>Imports a recording file of delay, loss, and bandwidth conditions. FTP, HTTP, and TFTP services are supported. Ex: recordings add http://192.168.0.100/recordings/wireless.txt</p> <pre>recordings delete <id></pre> <p>Deletes the specified recording file.</p>
serialnumber	<pre>serialnumber</pre> <p>Displays the serial number of the unit.</p>
ssh	<pre>ssh [enable disable fingerprint]</pre> <p>Enables or disables SSH service, or displays the fingerprint of the SSH server's public keys. If service is disabled, any sessions in progress are terminated. With no argument, 'ssh' reports current status of the service.</p>
telnet	<pre>telnet [enable disable]</pre> <p>Enables or disables telnet service. If Telnet is disabled, any sessions in progress are terminated. With no argument, the command reports current status of Telnet service.</p>

upgrade	<pre>upgrade <upgrade-image-url></pre> <p>Upgrades the Netropy firmware. FTP, HTTP, and TFTP services are supported. Ex: upgrade ftp://server/netropy-image</p>
user	<pre>user [add <user-id> delete <user-id> list]</pre> <p>Adds, deletes, or lists usernames.</p>
version	<pre>version</pre> <p>Displays the operating firmware version.</p>

11.3.1 LDAP Management Commands

CLI commands for LDAP authentication are listed in the table below:

show	<pre>mgmt show ldap</pre> <p>Displays the current LDAP configuration.</p>
set	<pre>mgmt set ldap server <server> [port <#>]</pre> <p>Identifies the LDAP server by IP address or host name. Optionally specifies a non-standard port to connect to. The default port is 389.</p> <p>Ex.: <code>mgmt set ldap server ldapserver.example.com</code></p> <pre>mgmt set ldap basedn <search-base-DN></pre> <p>Identifies the Distinguished Name of the search base in the remote LDAP database. If there are embedded spaces in the base DN string, the string must be enclosed in quotation marks.</p> <p>Ex.: <code>mgmt set ldap basedn dc=example,dc=com</code></p> <pre>mgmt set ldap filter attribute <string> [<LDAP-search-string>]</pre> <p>Specifies the attribute that contains the username and optionally an additional LDAP search string. The LDAP search string must be compatible with the formal definition found in RFC 4515. If there are embedded spaces in the search string, the string must be enclosed in quotation marks.</p> <p>Ex.: <code>mgmt set ldap filter attribute uid (&(gidNumber=20)(class=Expert))</code></p> <p>The resulting search string will be <code>(&(uid=user) (&(gidNumber=20) (class=Expert)))</code> where "user" is replaced by the login name.</p> <pre>mgmt set ldap security {disable enable}</pre> <p>Enable or disable Transport Level Security. TLS is enabled by default.</p> <pre>mgmt set ldap bind dn <bind-DN> [password <password>]</pre> <p>Specifies a Distinguished Name and password with which to bind to the LDAP server before performing a search operation. If not configured, an anonymous bind will be used.</p> <p>Ex. <code>mgmt set ldap bind dn cn=user,dc=example,dc=com</code></p> <pre>mgmt set ldap {on off}</pre> <p>Turns LDAP authentication on or off.</p>
clear	<pre>mgmt clear ldap</pre> <p>Turns off LDAP authentication and clears the LDAP configuration.</p>

11.3.2 Capture Command

The 'capture' command displays packets received or transmitted through the emulation ports.

This feature is intended to be used for examining and troubleshooting network connectivity. At high packet rates, the capture command will not capture all packets and may cause packet loss or timing errors for forwarded traffic.

This feature is only available through the CLI.

The capture command is unrelated to the packetcapture feature used to import and replay PCAP files as emulated background traffic.

Capture Command Syntax:

```
capture port <port> [raw|detail] [arp|ip <address>] [snaplen <len>]  
[filter <tcpdump filter string>]
```

Displays frames received by and sent from the specified port until control-C is pressed. Only one running capture can be active per port. Options are described below:

- ▶ **None:** a timestamp and summary description are displayed for each frame.
- ▶ **detail:** the Ethernet header and additional details about the contents of the frame are displayed.
- ▶ **raw:** binary PCAP data is output. Using this 'raw' option, it is possible to save a packet capture file on a PC by running the capture command via ssh. The resulting file can then be read by a packet analyzer such as Wireshark. For example, a Linux command could be:

```
ssh admin@netropy capture port 1 raw > port_1_packets.pcap
```

It is also possible to display packets in real-time in Wireshark with the command (on Linux):

```
ssh admin@netropy capture port 1 raw | wireshark -k -i -
```

- ▶ **arp|ip <address>:** adds a high speed pre-filter for lossless capture of a subset of packets matching a specific IP address or all ARP packets from a high packet rate stream.
- ▶ **snaplen:** limits frame capture to the first 'len' bytes of the frame.
- ▶ **filter:** controls which frames will be captured. Tcpdump-style filter options are accepted. For example, "filter ip host 10.0.0.1" captures all frames with IP source or destination address of 10.0.0.1.

11.4 CLI Engine Mode Commands

All emulation configuration commands require an emulation engine to be selected. Use the top-level 'engine <#>' command to select an engine and enter engine command mode.

While in engine command mode, all engine-level commands are applied to the currently selected engine. Use the 'exit' command to exit engine configuration mode and return to the top-level commands, for instance, to select a different engine.

On Netropy models with a single Emulation Engine, the engine number is 1. On models with multiple Emulation Engines, a list of available engine numbers is shown with the top level 'engine list' command.

Command	Syntax and Description
autostart	<pre>[engine <#>] autostart {on off}</pre> <p>Sets whether emulation is on or off for each engine when the system boots up. Autostart is off by default.</p>
backup	<pre>[engine <#>] backup</pre> <p>Displays the XML configuration of the selected engine for use in saving the engine configuration. Either copy and paste the terminal output or run a command that stores the output to a file.</p> <p>Ex: <code>ssh admin@netropy engine 1 backup > netropy-engine-1.xml</code></p>
emulation	<pre>[engine <#>] emulation {on off}</pre> <p>Turns emulation on or off for the specified engine.</p>
exit	<pre>exit</pre> <p>Exits engine configuration mode and returns to the top-level commands.</p>
linkfault mirror	<pre>[engine <#>] linkfaultmirror {on off}</pre> <p>Turns on or off the link fault mirror feature, which forces the Ethernet link down on both ports if the link partner on either port is down. Enables faster link outage detection in some high-availability networks. Only available on Netropy 10G1, 10G2, and 10G4 models.</p>
lock	<pre>[engine <#>] lock</pre> <p>Locks the engine configuration from changes by any other user except admin.</p>

log	<pre>[engine <#>] log read [unread all]</pre> <p>Displays engine log messages. By default only previously unread log messages are displayed.</p> <p>Log messages are displayed in the format: <sequence #> <time> <message></p> <p>Ex: 1 Nov 19 05:20:09 timing error exceeded 20us</p> <pre>[engine <#>] log clear</pre> <p>Clears all messages for this engine from the log.</p>
path	<pre>[engine <#>] path <#> [show set...]</pre> <p>Enters path command mode, or executes a command for the specified path. See Section 11.5 for path command mode syntax.</p>
restore	<pre>[engine <#>] restore <saved-engine-url></pre> <p>Loads the specified XML engine configuration file and restores the configuration to the selected engine. FTP, HTTP, and TFTP services are supported.</p>
router	<pre>[engine <#>] router {on off}</pre> <p>Turns routing on or off for the engine. If routing is off, the engine functions as a transparent bridge and forwards all traffic between the two interfaces of the engine. Routing is off by default.</p> <pre>[engine <#>] router show</pre> <p>Displays the current routing status and the ARP table for the engine.</p> <pre>[engine <#>] router interface list</pre> <p>Displays the interface name, MAC address, IP address, and netmask for both ports of the engine.</p> <pre>[engine <#>] router interface set <name> addr <addr> netmask <mask></pre> <p>Sets the IP address and netmask for the named interface. Interface names are PORT_<#> and shown using the <code>router interface list</code> command.</p> <pre>[engine <#>] router route add <addr>[/<prefix-length>] gw <gw-addr></pre> <p>Adds a static route to the specified address range via the specified next-hop gateway. For a default route, use a prefix length of 0. Each engine may be configured with up to 20 static routes. There can only be one gateway for each unique destination.</p> <p>Ex: [admin@netropy]> engine 1 router route add 10.0.1.0/24 gw 10.0.0.1</p> <pre>[engine <#>] router route list</pre> <p>Displays the list of static routes for the engine.</p> <pre>[engine <#>] router route delete <addr>[/<prefix-length>]</pre> <p>Deletes the specified static route.</p> <pre>[engine <#>] router route flush</pre> <p>Clears all static routes.</p>
show	<pre>[engine <#>] show</pre> <p>Displays the emulation on/off state, port numbers and connectivity, configured path numbers and names, and whether the engine is locked.</p>

statistics	<pre>[engine <#>] statistics [total interval cumulative] [full]</pre> <p>Reports statistics for all paths in the selected engine. Statistics are displayed as a single line of current totals, or one line for each one second sample using either cumulative or non-cumulative interval counters.</p> <p>By default, only the overall rate, frames, bytes, and drops are displayed for each direction of each path. To display all statistics, use the 'full' option.</p> <p>For more flexibility in selecting specific statistics and time ranges, use the statistics download feature in the GUI.</p> <pre>[engine <#>] statistics reset</pre> <p>Resets all statistics for the selected engine.</p>
unlock	<pre>[engine <#>] unlock</pre> <p>Unlocks the specified engine.</p>

11.5 CLI Path Mode Commands

Comm and	Syntax and Description
show	<pre>[engine <#> path <#>] show</pre> <p>Displays the current configuration of the specified path.</p>
set	<pre>[engine <#> path <#>] set...</pre> <p>Configures emulation properties for the specified path. See Section 11.5.1 for command syntax to set emulation properties.</p>

11.5.1 CLI Set Emulation Parameters

Commands to set the emulation parameters are listed below and entered in the form:

[engine <#> path <#>] set...

The syntax for each command is listed in the table below.

- ▶ Emulation parameters that are percentages may be entered as either a number between 0 and 100 with a % sign (i.e. 27.1%) or as a number between 0 and 1 (i.e. "0.271" or ".271").
- ▶ Attributes of the WAN portion of the path (delay, loss, duplication, reordering) must specify a direction from one port to another, i.e. "port 1 to port 2".
- ▶ Attributes of the interface between LAN and WAN (bandwidth, background utilization, queuing, and frame overhead) must specify the port number and whether the condition is outbound or inbound between WAN and LAN. By default, inbound rate control is disabled, and inbound attributes are not used.

Command	Syntax and Description
bg	<pre>set bg none port <#> {out in}</pre> <p>Turns off background traffic generation for the specified port and direction.</p> <pre>set bg random <percent> burst <bytes> port <#> {out in}</pre> <p>Creates random background traffic with the specified link utilization rate and burst size for the specified port and direction.</p> <pre>set bg replay pcap {id <id> [count <#>] [scale <factor>] [priority <0-7>]}... port <#> {out in}</pre> <p>Replays imported packet capture files as background traffic for the specified port and direction. Use 'packetcaptures list' from the top command level to view a list of available PCAP files and their ID numbers. Multiple PCAP files can be run simultaneously</p> <p>Each file optionally be can be modified with a count to replicate multiple copies of the same stream, a scaling factor to modify the playback speed, and a priority level to use for priority and round robin queuing. See Section 6.2.3.2 for details on PCAP replay parameters.</p>
bw	<pre>set bw fixed <value> [bps Kbps Mbps Gbps] port <#> {out in}</pre> <p>Sets a fixed value for the outbound or inbound bandwidth on the specified port. The entered value is in bps if no units are specified.</p> <pre>set bw recorded port <#> {out in}</pre> <p>Uses the bandwidth values from the recording file previously selected with the playback command.</p>

corruption	<pre>set corruption none port <#> to port <#></pre> <p>Turns off data corruption in the specified direction.</p> <pre>set corruption ber <rate> port <#> to port <#></pre> <p>Sets data corruption as a bit error rate in exponential or floating point notation.</p>
delay	<pre>set delay none port <#> to port <#></pre> <p>Disables latency emulation in the specified direction.</p> <pre>set delay constant <ms> port <#> to port <#></pre> <p>Sets latency emulation to a constant value in milliseconds.</p> <pre>set delay uniform <min> <max> [reordering {enabled disabled}] port <#> to port <#></pre> <p>Sets latency emulation to a uniform distribution between the specified minimum and maximum values in milliseconds. Enables or disables reordering of packets with different delays (default is no reordering).</p> <pre>set delay normal <mean> <stddev> [reordering {enabled disabled}] port <#> to port <#></pre> <p>Sets latency emulation to a normal distribution with mean and standard deviation specified in milliseconds. Enables or disables reordering of packets with different delays (default is no reordering).</p> <pre>set delay exponential <min> <max> [reordering {enabled disabled}] port <#> to port <#></pre> <p>Sets latency emulation to an exponential distribution with the specified minimum and mean values in milliseconds. Enables or disables reordering of packets with different delays (default is no reordering).</p> <pre>set delay accumulate-and-burst <count> [extra-delay <ms>] [timeout <ms>] port <#> to port <#></pre> <p>Packets are held until either the packet count or an optional timeout is reached, then are optionally delayed by an additional 'extra delay,' and transmitted in a burst. If a timeout is not specified, the default 10,000 ms will be used.</p> <pre>set delay recorded port <#> to port <#></pre> <p>Uses delay values from the recording file previously selected with the <code>playback</code> command.</p>
duplication	<pre>set duplication none port <#> to port <#></pre> <p>Turns off packet duplication in the specified direction.</p> <pre>set duplication random <percent> port <#> to port <#></pre> <p>Specifies the probability for random packet duplication set in increments of 0.0001%.</p>
inbound-rate-control	<pre>set inbound-rate-control {enable disable} port <#></pre> <p>Turns on or off controls on bandwidth, background utilization, queue management, and framing overhead for traffic arriving in-bound from the WAN to the specified port number. Turned off by default.</p>

link-type	<pre>set link-type {single per-ip-src per-ip-dst} port <#> [out]</pre> <p>Switches paths between single-link mode and per-client emulation mode. Set to single-link mode by default. Can only be changed while emulation is off. See Section 5.2 for more details.</p>
loss	<pre>set loss none port <#> to port <#></pre> <p>Disables packet loss emulation in the specified direction.</p> <pre>set loss random <percent> port <#> to port <#></pre> <p>Sets loss emulation as a random packet loss rate set in increments of 0.0001%.</p> <pre>set loss burst <percent> min <min-burst-packets> max <max-burst-packets> port <#> to port <#></pre> <p>Sets burst packet loss emulation as a burst probability and minimum and maximum burst size in packets. The percentage can be set in increments of 0.0001%.</p> <pre>set loss periodic <period> [burst <count>] port <#> to port <#></pre> <p>Sets periodic loss emulation as a loss period in packets. An optional burst count specifies the number of packets to be lost in a row in each cycle, and is set to one packet by default.</p> <pre>set loss ber <rate (1e-18 - 9.999999e-1)> port <#> to port <#></pre> <p>Sets packet loss emulation as a bit error rate in exponential or floating point. Ex. set loss ber 1.0e-8 port 1 to port 2</p> <pre>set loss gilbert-elliott good-state loss <pct> change <pct> bad-state loss <pct> change <pct> port <#> to port <#></pre> <p>Sets packet loss separately for a good state and bad state and the probability of transitioning between the two states.</p> <pre>set loss recorded port <#> to port <#></pre> <p>Uses packet loss rates from the recording file previously selected with the playback command.</p>
mtu	<pre>set mtu-limit fixed <length> [send-icmp {enabled disabled}] [fragmentation {standard never ignore-df}] port <#> {out in}</pre> <p>Limits the MTU to the specified length in bytes. Sending of ICMP Destination Unreachable Fragmentation Needed messages can be enabled or disabled. Select 'fragmentation standard' to have fragmentation follow the Don't Fragment (DF) bit setting, 'fragmentation never' to drop all packets larger than the MTU length, or 'fragmentation ignore-df' to fragment regardless of the DF setting. If unspecified, send-icmp is enabled and fragmentation set to 'standard.'</p> <pre>set mtu-limit none port <#> {out in}</pre> <p>Disables MTU limitations on the specified port.</p>
name	<pre>set name <path-name></pre> <p>Sets the name of the path.</p>

network-outage	<p><code>set network-outage none port <#> to port <#></code></p> <p>Disables network outage emulation in the specified direction.</p> <p><code>set network-outage periodic interval <min>[-<max>] duration <min>[-<max>] port <#> to port <#></code></p> <p>Sets the duration and interval of a periodic network outage in seconds, set as single values or a range of values separated with a hyphen. Values can be set between 0.001 seconds and 604,800 seconds in increments of 0.001 seconds.</p>
overhead	<p><code>set overhead fixed <bytes> port <#> {out in}</code></p> <p>Sets the framing overhead in bytes for each packet for the specified port and direction.</p>
playback	<p><code>set playback none port <#> to port <#></code></p> <p>Turns off use of recordings in the specified direction.</p> <p><code>set playback recording <id> port <#> to port <#></code></p> <p>Selects the recoding to use for delay, loss, and/or bandwidth conditions in the specified direction. Use the <code>recordings</code> command to display a list of available recording IDs. <code>playback</code> can only be run while emulation is off. After the recording has been selected, use the <code>set delay</code>, <code>set loss</code>, and <code>set bandwidth</code> commands to use the recorded values of delay and loss.</p>
queue	<p><code>set queue default port <#> {out in}</code></p> <p>Returns the maximum queue depth setting to the default value for specified port and direction.</p> <p><code>set queue droptail <value> {packets KB} port <#> {out in}</code></p> <p>Sets the queue management algorithm to Drop Tail with the maximum queue depth specified in packets or kilobytes.</p> <p><code>set queue red depth <value> min <min> max <max> {packets KB} port <#> {out in}</code></p> <p>Sets the queue management algorithm to RED. Specifies the maximum queue depth and RED minimum and maximum thresholds in either packets or kilobytes.</p>
queuing-strategy	<p><code>set queuing-strategy fifo port <#> {out in}</code></p> <p>Returns the queuing strategy to the default FIFO, with no prioritization.</p> <p><code>set queuing-strategy priority {ip-precedence vlan-pcp} port <#> {out in}</code></p> <p>Sets the queuing strategy to prioritize packets based on either IP precedence or VLAN PCP bits.</p> <p><code>set queuing-strategy round-robin {ip-precedence vlan-pcp} port<#> {out in}</code></p> <p>Sets the queuing strategy to round-robin packets between queues, with packets assigned to different queues based on either IP Precedence or VLAN PCP bits.</p>

reordering	<pre>set reordering none port <#> to port <#></pre> <p>Turns off packet reordering for the configured engine/path in the specified direction.</p> <pre>set reordering random <percent> min <min-offset-pkts> max <max- offset-pkts> [timeout <ms>] port <#> to port <#></pre> <p>Specifies the probability for random packet reordering and the minimum and maximum offset for reordered packets. The percentage can be set in increments of 0.0001%. An optional timeout value specifies the maximum amount of time to wait for the offset packets to arrive. If a timeout value is not specified, the default of 10,000 ms will be used.</p>
------------	--

12 SECURITY

12.1 Users and Passwords

Initially, the Netropy system has a single configured user, *admin*, with no password. Additional users can be added or deleted by the *admin* user through the Administration window of the GUI or through the CLI. Passwords for each user can be set or cleared by the individual user or by *admin*. Usernames and passwords created in either the GUI or through the CLI apply to access to both the GUI and CLI. LDAP server can also be used for user authentication.

12.2 Engine Locking

Users can lock individual Emulation Engines to prevent any changes made to the configuration by other users. The user that locked the engine can continue changing the configuration of a locked engine, as can the *admin* user. Locked engines can only be unlocked by the user that locked the engine or by the *admin* user.

Locking is specific to each individual engine. Different users can lock different engines, or a single user can lock multiple engines.

12.3 Recovering from a Lost Admin Password

If the *admin* password has been lost, boot to the recovery firmware to gain access to the device and reset the configuration to factory defaults:

1. Connect to the serial console port.
2. Power cycle the unit. Type “*recovery*” at the boot prompt to load the recovery image.

```
boot: recovery
```

3. At the prompt, log in as “*admin*”.

```
netropy login: admin
```

4. Reset the configuration.

```
[admin@netropy]> init config
```

5. Reboot or power cycle the unit and return to the regular operating firmware with no configuration. The management interface will also be reset to the default of 10.0.0.10.

12.4 SSL

The Netropy GUI is accessible via HTTP or HTTPS. HTTPS allows administration of the GUI using SSL security. Netropy includes a non-unique, self-signed certificate. Use of this self-signed certificate may generate an error in the browser that the signing certificate authority is unknown and not trusted. Either ignore this error or install your own certificate.

To install a new certificate, use the “http” command from the CLI:

```
http load <url>
```

For example:

```
netropy> http load http://192.168.0.100/certs/certificate.crt
```

By default, the Netropy GUI is accessible via unsecured HTTP. However, HTTP service can be disabled using the “http off” command from the CLI. HTTPS service is always enabled.

12.5 SSH

The Netropy CLI is accessible over the network via SSH. To verify the identity of the Netropy SSH server, use the “ssh fingerprint” command to display the fingerprints of the SSH server’s public keys.

12.6 LDAP

An LDAP server can be used for user authentication. When LDAP authentication is enabled, it is used in addition to local authentication for GUI and CLI access.

LDAP can only be configured through the command line interface. See Section 11.3.1 for LDAP command syntax.

To authenticate a user via LDAP, the following steps are performed:

1. A connection to the LDAP server is initiated. If LDAP security is enabled, a TLS session is negotiated.

2. An initial bind operation is performed on the connection. If a 'bind DN' has been configured, then the configured 'bind DN' and password are used for the bind. Otherwise an anonymous bind is attempted.
3. A search is performed to find the database entry that corresponds to the user's username. The search parameters consist of a configurable 'base DN' and a search filter. The 'base DN' specifies the root of the subtree that will be searched. The default search filter requires an exact match between the username and the 'uid' attribute. An alternate attribute for this comparison can be configured. For a more specific search, a full LDAP search string can be ANDed to this filter.
4. If a matching entry is returned by the server, then a new bind is attempted for authentication. This bind uses the DN found in the returned entry and the password supplied by the user. The "Simple" authentication method is requested.
5. If the bind is successful, then authentication is successful.

Notes:

Only one matching entry will be returned by the LDAP search even if there are multiple database entries that match the search filter.

For locally-defined users, the local authentication is always attempted before LDAP authentication.

Only the user *admin* has administrative privileges.

13 APPOSITE SUPPORT

13.1 Registration

For access to the Netropy support site with firmware upgrades, documentation, and other support materials, register your unit at: <http://www.apposite-tech.com/register.html>. License key download will automatically register you.

Registered users will receive email notification whenever new firmware versions are released.

13.2 Customer Support

If you experience any problem with the Netropy hardware, consult the *Hardware Guide* for your model. If you have any questions about the firmware not answered in this *User's Guide*, please check the Apposite Technologies website at <http://www.apposite-tech.com> for updated firmware and documentation. If your question is not answered, please contact Apposite Support.

Netropy units with a product license receive support and maintenance for the term of the license, including hardware warranty, firmware updates, and customer support. The license expiration date for limited term licenses is shown in the Netropy GUI on the Administration window License Key tab.

For units with an unlimited term license (purchased prior to August 2017), one year of maintenance is included with purchase, after which a yearly maintenance renewal is required for ongoing support. License and maintenance renewals are available from Apposite or your Netropy reseller.

If you believe the firmware is not functioning properly, please upgrade to the latest firmware release. If the problem persists, please contact Apposite Support at:

support@apposite-tech.com
1.310.477.9955 ext. 2

When contacting Apposite Support, please include the following information:

- ▶ Netropy model
- ▶ Serial number

- ▶ Your e-mail address and phone number
- ▶ Installed firmware version
- ▶ A detailed description of the problem



Do not attempt to fix any hardware problem yourself. Netropy Network Emulators contain no user serviceable parts. Opening the chassis voids the warranty.

APPENDIX A:

APPOSITE END USER LICENSE AGREEMENT

THIS END USER LICENSE AGREEMENT ("AGREEMENT") IS A LEGAL AGREEMENT BETWEEN YOU AND APPOSITE TECHNOLOGIES LLC ("APPOSITE"). BY DOWNLOADING, INSTALLING, COPYING, ACCESSING OR USING THIS SOFTWARE, YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON OR COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS, ARE NOT AT LEAST 18 YEARS OF AGE, OR DO NOT HAVE AUTHORITY TO ENTER THIS AGREEMENT, DO NOT DOWNLOAD, INSTALL, COPY, ACCESS OR USE THIS SOFTWARE, AND PROMPTLY RETURN THIS SOFTWARE AND PROOF OF ENTITLEMENT TO THE PARTY FROM WHOM YOU ACQUIRED THEM.

1. DEFINITIONS.

- a. "Documentation" means published user manuals and other explanatory materials provided by Apposite relating to the Software or the Hardware.
- b. "Hardware" means the Apposite-branded hardware product on which the Software is designed to run.
- c. "Software" means the computer software (regardless of format or media in which distributed) designed to run on the Hardware.
- d. "Modifications" means any applicable upgrades, modified versions, updates, bug fixes, patches, or other modifications to the Software made generally available by Apposite to its customers.
- e. "Product" means the Apposite-branded product comprised of the Hardware, the Software, and the Documentation.
- f. "Legacy Product" means any Product purchased prior to August 1, 2017.
- g. "License Period" means for any Product that is not a Legacy Product, the license period identified on the relevant order form or receipt upon your purchase of the Product or the Software, including any purchased extension thereto.
- h. "Warranty Period" means (a) for any Legacy Product, the period during which you maintain a continuous, active subscription to Apposite maintenance services, or (b) for any Product that is not a Legacy Product, the License Period, unless such Product has a perpetual license, in which case the Warranty Period shall be five (5) years from the date of purchase of the Product from Apposite unless extended by purchase of maintenance services.

2. LICENSE GRANT.

Conditioned upon compliance with the terms and conditions of this Agreement, Apposite hereby grants you the following licenses (collectively, the "License"):

- a. With respect to any Product that is not a Legacy Product, for the relevant License Period, a nonexclusive, nonsublicensable, nontransferable license to use the Software and Modifications solely as embedded in the Hardware, and to use the Documentation solely in conjunction with the Software, Modifications, and Hardware; and
- b. With respect to any Legacy Product, a perpetual, nonexclusive, nonsublicensable, nontransferable license to use the Software solely as embedded in the Hardware, and to use the Documentation solely in conjunction with such Software and Hardware, and to use any Modifications released by Apposite and installed on the Hardware during the Warranty Period.

This license is not a sale. The Software may contain open source or third-party software, subject to separate license terms. Acknowledgements, licensing terms and disclaimers for such materials are contained in the Documentation, and your use of such material is governed by their respective terms and not by the License.

3. **OWNERSHIP.** Apposite and its licensors retain ownership of all intellectual property rights in and to the Software and Documentation, including any copies, improvements, enhancements, derivative works and modifications thereof. The Software and Documentation are protected by United States and international copyright and other intellectual property laws and international trade provisions. Your rights to use the Software and Documentation are limited to those expressly granted by this Agreement. No other rights with respect to the Software or any related intellectual property rights are granted or implied.
4. **RESTRICTIONS.** Except as otherwise expressly provided under this Agreement, you shall not directly or authorize any third party to:
 - (i) copy the Software, except as necessary for archival or backup purposes only;
 - (ii) transfer, assign, sublicense, rent, lease, lend, or otherwise transfer your license rights to any other person or entity;
 - (iii) install or use the Software on any computing device other than the Hardware;
 - (iv) modify or adapt the Software or Documentation or create derivative works based upon the Software or Documentation;
 - (v) reverse engineer, disassemble, decompile, decrypt, or otherwise attempt to derive the source code of the Software, except to the extent otherwise expressly permitted under applicable law notwithstanding this restriction; or
 - (vi) remove, alter, cover or obfuscate any of the trademarks, trade names, logos, patent or copyright notices or markings, or add any other notices or markings to or on the Software, Documentation, or Hardware without the express written authorization of Apposite.
5. **SUPPORT.** Apposite shall have no obligation to support the Software except as provided in Apposite's Customer Support Terms and Conditions, which is incorporated by reference into this Agreement.
6. **LIMITED WARRANTY.** Apposite warrants that, for the duration of the Warranty Period, the Hardware will be free of any defects in materials and workmanship under normal use and the Software will perform substantially in accordance with the Documentation. This limited warranty extends only to the original user of the Product. This limited warranty is void if failure of the Hardware or Software to conform to the warranty has resulted from improper installation, testing, misuse, neglect, accident, fire or other hazard, a virus, or any breach of this Agreement. Apposite and its suppliers' entire liability and your sole and exclusive remedy shall be, at Apposite's sole discretion, to: (i) repair the Software or Hardware; (ii) provide replacement Hardware or Software; or (iii) refund the purchase price and terminate this Agreement. This limited warranty applies only if the Product is returned to Apposite, freight and insurance prepaid, in accordance with Apposite's Return Material Authorization (RMA) procedures. Any repaired or replaced Software or Hardware will be subject to this same limited warranty for the remainder of the original Warranty Period or thirty (30) days, whichever is greater. THE FOREGOING LIMITED WARRANTY IS IN LIEU OF ALL OTHER WARRANTIES OR CONDITIONS, EXPRESS OR IMPLIED, AND APPOSITE DISCLAIMS ANY AND ALL IMPLIED WARRANTIES OR CONDITIONS, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF TITLE, NONINFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. THIS WARRANTY GIVES YOU SPECIFIC LEGAL RIGHTS, AND YOU MAY ALSO HAVE OTHER RIGHTS WHICH VARY BY JURISDICTION.

- 7. LIMITATION OF LIABILITY.** REGARDLESS WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE, IN NO EVENT WILL APPOSITE OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION, ANY LOST PROFITS, BUSINESS INTERRUPTION, LOSS OF INFORMATION, OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE THE PRODUCT AND EVEN IF APPOSITE OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL APPOSITE OR ITS SUPPLIERS' LIABILITY, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY, OR OTHERWISE, EXCEED THE PRICE PAID FOR THE PRODUCT. SOME STATES DO NOT ALLOW THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. YOU HEREBY WAIVE CALIFORNIA CIVIL CODE SECTION 1542, WHICH SAYS "A GENERAL RELEASE DOES NOT EXTEND TO CLAIMS WHICH THE CREDITOR DOES NOT KNOW OR SUSPECT TO EXIST IN HIS FAVOR AT THE TIME OF EXECUTING THE RELEASE, WHICH, IF KNOWN BY HIM MUST HAVE MATERIALLY AFFECTED HIS SETTLEMENT WITH THE DEBTOR."
- 8. TERM AND TERMINATION.** This Agreement takes effect upon your use of the Software and remains effective until terminated. You may terminate it at any time by destroying all copies of the Software and Documentation in your possession. The licenses granted by Apposite herein will terminate immediately if you fail to comply with any term or condition of this Agreement. You agree on termination of the licenses to cease all use of the Software and Documentation. In addition, the provisions of the sections "U.S. Government Restricted Rights" and "Export" shall survive termination of this Agreement.
- 9. EXPORT.** The Product, including the underlying technology, is subject to U.S. export control laws, and may be subject to export or import regulations in other countries. You may not export or import the Product and the underlying technology, directly or indirectly, in violation of these laws. You agree to comply strictly with all such laws and regulations and acknowledge that you have the responsibility to obtain such licenses to export, re-export, or import as may be required.
- 10. U.S. GOVERNMENT RESTRICTED RIGHTS.** All Software and Documentation provided pursuant to this Agreement qualifies as "commercial computer software" and "commercial computer software documentation" pursuant to FAR 12.212 and DFAR 227.7202. The Software and Documentation are being licensed to U.S. Government end users only as commercial items and with only those rights as are granted to all other end users pursuant to the terms and conditions of this agreement.
- 11. GENERAL PROVISIONS.** This Agreement will be governed by and construed in accordance with the laws of the State of California without reference to its conflicts of law provisions. Any dispute arising out of or relating to this Agreement shall be subject to the exclusive jurisdiction and venue of the courts sitting in Los Angeles, California. This Agreement shall not be governed by either the Uniform Computer Information Transactions Act (UCITA) or the United Nations Convention on Contracts for the International Sale of Goods (CISG), the application of which is expressly excluded. If for any reason a court of competent jurisdiction finds any provision of this Agreement, or any portion thereof, to be unenforceable, the remainder of this Agreement shall continue in full force and effect. This Agreement constitutes the entire agreement between the parties with respect to the Product and supersedes all prior or contemporaneous understandings regarding such subject matter, whether written or oral, and supersedes any conflicting or additional terms contained in any purchase order or elsewhere, all of which terms are excluded. No amendment to or modification of this Agreement are binding unless in writing and signed by Apposite.

APPENDIX B: THIRD PARTY LICENSES

Apache

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

You must give any other recipients of the Work or Derivative Works a copy of this License; and

You must cause any modified files to carry prominent notices stating that You changed the files; and

You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

Click License

Portions of this software are subject to the license below. The relevant source files are clearly marked; they refer to this file using the phrase "the Click LICENSE file". This license is an MIT license, plus a clause (taken from the W3C license) requiring prior written permission to use our names in publicity. The AUTHORS file lists the people who have contributed to this software.

(c) 1999-2007 Massachusetts Institute of Technology

(c) 2000-2007 Mazu Networks, Inc.

(c) 2001-2007 International Computer Science Institute

(c) 2004-2007 Regents of the University of California

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the Software without specific, written prior permission. Title to copyright in this Software and any associated documentation will at all times remain with copyright holders.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Degrafa

Beta 3.1

Copyright (c) 2008 The Degrafa Team : <http://www.Degrafa.com/team>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Flex-chart

Copyright (c) 2007 Brendan Meutzner

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Flexlib

Copyright (c) 2007 FlexLib Contributors

See <http://code.google.com/p/flexlib/wiki/ProjectContributors>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

GNU Software

Netropy Software incorporates open source program files distributed under the GNU General Public License (GPL), version 2, a copy of which is included below. For a period of three years, upon request, Apposite will send you a machine-readable copy of the source code of these program files.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991. Copyright (C) 1989, 1991 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new

free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this

License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute

so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Jonpy

Copyright (c) 2002 Jon Ribbens

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Libedit

Copyright (c) 1992, 1993 The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Christos Zoulas of Cornell University.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Libxml2

Except where otherwise noted in the source code (e.g. the files hash.c, list.c and the trio files, which are covered by a similar licence but with different Copyright notices) all the files are:

Copyright (C) 1998-2012 Daniel Veillard. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Libxslt2

Licence for libxslt except libexslt

Copyright (C) 2001-2002 Daniel Veillard. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE DANIEL VEILLARD BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Daniel Veillard shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from him.

LightTPD

Copyright (c) 2004, Jan Kneschke, incremental

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the 'incremental' nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

LILLO

Linux LOader (LILLO) program code, documentation, and auxiliary programs:

Copyright 1992-1998 Werner Almesberger

Copyright 1999-2007 John Coffman

Copyright 2009-2013 Joachim Wiedorn

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the names of the authors nor the names of other contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Ncurses

Copyright (c) 1998-2014, 2015 Free Software Foundation, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, distribute with modifications, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESSOR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE ABOVE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name(s) of the above copyright holders shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization.

NTP

Copyright Notice

jpg "Clone me," says Dolly sheepishly.

Last update: 17-Jan-2015 00:16 UTC

The following copyright notice applies to all files collectively called the Network Time Protocol Version 4 Distribution. Unless specifically declared otherwise in an individual file, this entire notice applies as if the text was explicitly included in the file.

Copyright (c) University of Delaware 1992-2015

Permission to use, copy, modify, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appears in all copies and that both the copyright notice and this permission notice appear in supporting documentation, and that the name University of Delaware not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. The University of Delaware makes no representations about the suitability this software for any purpose. It is provided "as is" without express or implied warranty

Content starting in 2011 from Harlan Stenn, Danny Mayer, and Martin Burnicki is:

Copyright (c) Network Time Foundation 2011-2015 All Rights Reserved

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following

disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE

ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

OpenLDAP Public License

Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions in source form must retain copyright statements and notices,
2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

OpenSSH License

(1) Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland. All rights reserved.

(2) Cryptographic attack detector for ssh: Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina. All rights reserved.

(3) Ssh-keyscan: Copyright 1995, 1996 by David Mazieres <dm@lcs.mit.edu>

(4) Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto.

(5) One component of the ssh source code is under a 3-clause BSD license held by the University of California. Copyright (c) 1983, 1990, 1992, 1993, 1995 The Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

(6) Components of the software are provided under a standard 2-term BSD license with the following names as copyright holders: Markus Friedl, Theo de Raadt, Niels Provos, Dug Song, Aaron Campbell, Damien Miller, Kevin Steves, Daniel Kouril, Wesley Griffin, Per Allansson, Nils Nordman, Simon Wilkinson.

(7) Portable OpenSSH additionally includes code from the following copyright holders, also under the 2-term BSD license: Ben Lindstrom, Tim Rice, Andre Lucas, Chris Adams, Corinna Vinschen, Cray Inc., Denis Parker, Gert Doering, Jakob Schlyter, Jason Downs, Juha Yrjölä, Michael Stone, Networks Associates Technology, Inc., Solar Designer, Todd C. Miller, Wayne Schroeder, William Jones, Darren Tucker, Sun Microsystems, The SCO Group, Daniel Walsh. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

8) Portable OpenSSH contains the following additional licenses:

a) md5crypt.c, md5crypt.h: Poul-Henning Kamp

b) snprintf replacement: Copyright Patrick Powell 1995

c) Compatibility code (openbsd-compat)

Apart from the previously mentioned licenses, various pieces of code in the openbsd-compat/ subdirectory are licensed as follows: Some code is licensed under a 3-term BSD license, to the following copyright holders: Todd C. Miller, Theo de Raadt, Damien Miller, Eric P. Allman, The Regents of the University of California, Constantin S. Svintsoff. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Some code is licensed under an ISC-style license, to the following copyright holders: Internet Software Consortium, Todd C. Miller, Reyk Floeter, Chad Mynhier. Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

Some code is licensed under a MIT-style license to the following copyright holders: Free Software Foundation, Inc. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, distribute with modifications, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE ABOVE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name(s) of the above copyright holders shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization.

OpenSSL License

Copyright (c) 1998-2008 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)".

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license[including the GNU Public License.]

PAM

Redistribution and use in source and binary forms of Linux-PAM, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain any existing copyright notice, and this entire permission notice in its entirety, including the disclaimer of warranties.
2. Redistributions in binary form must reproduce all prior and current copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of any author may not be used to endorse or promote products derived from this software without their specific prior written permission.

ALTERNATIVELY, this product may be distributed under the terms of the GNU General Public License, in which case the provisions of the GNU GPL are required INSTEAD OF the above restrictions. (This clause is necessary due to a potential conflict between the GNU GPL and the restrictions contained in a BSD-style copyright.)

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.

IN NO EVENT SHALL THE AUTHOR(S) BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PureMVC

Copyright(c) 2006-08 Futurescale, Inc., Some rights reserved.

Reuse governed by Creative Commons Attribution 3.0 United States License
View this license and the associated Creative Commons deed online at:
<http://creativecommons.org/licenses/by/3.0/us/>
You must leave all copyright and license statements intact in your final product.

License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED.

BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. TO THE EXTENT THIS LICENSE MAY BE CONSIDERED TO BE A CONTRACT, THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. Definitions

1. "Collective Work" means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with one or more other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this License.

2. "Derivative Work" means a work based upon the Work or upon the Work and other pre-existing works, such as a translation, musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work will not be considered a Derivative Work for the purpose of this License. For the avoidance of doubt, where the Work is a musical composition or sound recording, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered a Derivative Work for the purpose of this License.

3. "Licensor" means the individual, individuals, entity or entities that offers the Work under the terms of this License.

4. "Original Author" means the individual, individuals, entity or entities who created the Work.

5. "Work" means the copyrightable work of authorship offered under the terms of this License.

6. "You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.

2. Fair Use Rights. Nothing in this license is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.

3. License Grant. Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:

1. to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;

2. to create and reproduce Derivative Works provided that any such Derivative Work, including any translation in any medium, takes reasonable steps to clearly label, demarcate or otherwise identify that changes were made to the original Work. For example, a translation could be marked "The original work was translated from English to Spanish," or a modification could indicate "The original work has been modified.";;

3. to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works;

4. to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission Derivative Works.

5. For the avoidance of doubt, where the Work is a musical composition:

1. Performance Royalties Under Blanket Licenses. Licensor waives the exclusive right to collect, whether individually or, in the event that Licensor is a member of a performance rights society (e.g. ASCAP, BMI, SESAC), via that society, royalties for the public performance or public digital performance (e.g. webcast) of the Work.

2. Mechanical Rights and Statutory Royalties. Licensor waives the exclusive right to collect, whether individually or via a music rights agency or designated agent (e.g. Harry Fox Agency), royalties for any phonorecord You create from the Work ("cover version") and distribute, subject to the compulsory license created by 17 USC Section 115 of the US Copyright Act (or the equivalent in other jurisdictions).

6. Webcasting Rights and Statutory Royalties. For the avoidance of doubt, where the Work is a sound recording, Licensor waives the exclusive right to collect, whether individually or via a performance-rights society (e.g. SoundExchange), royalties for the public digital performance (e.g. webcast) of the Work, subject to the compulsory license created by 17 USC Section 114 of the US Copyright Act (or the equivalent in other jurisdictions).

The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensor are hereby reserved.

4. Restrictions. The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

1. You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this License, and You must include a copy of, or the Uniform Resource Identifier for, this License with every copy or phonorecord of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that restrict the terms of this License or the ability of a recipient of the Work to exercise the rights granted to that recipient under the terms of the License. You may not sublicense the Work. You must keep intact all

notices that refer to this License and to the disclaimer of warranties. When You distribute, publicly display, publicly perform, or publicly digitally perform the Work, You may not impose any technological measures on the Work that restrict the ability of a recipient of the Work from You to exercise the rights granted to that recipient under the terms of the License. This Section 4(a) applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this License. If You create a Collective Work, upon notice from any Licensor You must, to the extent practicable, remove from the Collective Work any credit as required by Section 4(b), as requested. If You create a Derivative Work, upon notice from any Licensor You must, to the extent practicable, remove from the Derivative Work any credit as required by Section 4(b), as requested.

2. If You distribute, publicly display, publicly perform, or publicly digitally perform the Work (as defined in Section 1 above) or any Derivative Works (as defined in Section 1 above) or Collective Works (as defined in Section 1 above), You must, unless a request has been made pursuant to Section 4(a), keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or (ii) if the Original Author and/or Licensor designate another party or parties (e.g. a sponsor institute, publishing entity, journal) for attribution ("Attribution Parties") in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; the title of the Work if supplied; to the extent reasonably practicable, the Uniform Resource Identifier, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work; and, consistent with Section 3(b) in the case of a Derivative Work, a credit identifying the use of the Work in the Derivative Work (e.g., "French translation of the Work by Original Author," or "Screenplay based on original Work by Original Author"). The credit required by this Section 4(b) may be implemented in any reasonable manner; provided, however, that in the case of a Derivative Work or Collective Work, at a minimum such credit will appear, if a credit for all contributing authors of the Derivative Work or Collective Work appears, then as part of these credits and in a manner at least as prominent as the credits for the other contributing authors. For the avoidance of doubt, You may only use the credit required by this Section for the purpose of attribution in the manner set out above and, by exercising Your rights under this License, You may not implicitly or explicitly assert or imply any connection with, sponsorship or endorsement by the Original Author, Licensor and/or Attribution Parties, as appropriate, of You or Your use of the Work, without the separate, express prior written permission of the Original Author, Licensor and/or Attribution Parties.

5. Representations, Warranties and Disclaimer

UNLESS OTHERWISE MUTUALLY AGREED TO BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND ONLY TO THE EXTENT OF ANY RIGHTS HELD IN THE LICENSED WORK BY THE LICENSOR. THE LICENSOR MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MARKETABILITY, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. Limitation on Liability. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination

1. This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Derivative Works (as defined in Section 1 above) or Collective Works (as defined in Section 1 above) from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

2. Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. Miscellaneous

1. Each time You distribute or publicly digitally perform the Work (as defined in Section 1 above) or a Collective Work (as defined in Section 1 above), the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.

2. Each time You distribute or publicly digitally perform a Derivative Work, Licensor offers to the recipient a license to the original Work on the same terms and conditions as the license granted to You under this License.

3. If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

4. No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

5. This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

Python-lxml

Copyright (c) 2004 Infracore. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of Infrae nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL INFRAE OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Python-pam
(c) 2007 Chris AtLee <chris@atlee.ca>

Licensed under the MIT license:

<http://www.opensource.org/licenses/mit-license.php>

Pyparsing
Copyright (c) 2003-2013 Paul T. McGuire

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Apposite® Technologies LLC

17835 Ventura Blvd., Suite 211

Los Angeles, CA 91316 USA

1.310.477.9955

info@apposite-tech.com

www.apposite-tech.com