



DCB, Inc.
2949 CR 1000 E
Dewey, Illinois
61840

217.897.6600 Tel
800.432.2638 Toll Free
217.897.8023 Fax
www.dcbnet.com

PNE
Network Emulation and Impairment Testing
User Guide

Update October 28, 2021

Copyright © 2018-2021 All Rights Reserved

Table of Contents

Certifications.....	3
FCC Statement.....	3
ROHS.....	3
Introduction.....	4
Hardware.....	4
Connections.....	4
Specifications.....	5
Configuration.....	6
Navigation.....	6
Profile Configuration.....	8
System Configuration.....	8
System Settings.....	9
Network Settings.....	9
Ethernet Configuration.....	9
IPv4 Settings Screen.....	10
Network and Device Status.....	10
Device Software.....	11
Network Emulation and Impairment Tests.....	12
Timing Emulation.....	12
Transmission Rate-Limit.....	12
Latency and Jitter.....	13
Timing Bypass.....	15
Error Emulation.....	15
Drop Errors.....	16
Duplicate Packets.....	17
Out-of-Order Packets.....	17
Corrupted Packets.....	17
Targeting Errors At A Specific Device.....	18
Errors On-Demand.....	19
Monitoring.....	20
Packet Rates.....	20
Average Packet Rates.....	21
Interpreting these displays.....	22
Firmware Updates.....	23
Step by Step Instructions.....	23
Use Examples.....	25

Certifications

FCC Statement

This device complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules.

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instruction manual may cause harmful interference to radio communications.

Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at the user's own expense.

All trademarks and trade names are the properties of their respective owners.

ROHS

Some models of this product are available in RoHS versions.




Introduction

PNE is a network timing and error emulation device. Its purpose is to allow emulation of real-world network conditions in a bench-top test environment. PNE is simply inserted in the data path between devices under test. PNE may then be configured to emulate the expected operating environment. This includes placing rate limits on the link, adding latency and inducing jitter. Network packets can be dropped, duplicated, or sent out-of-order. These are all conditions that happen in a real environment but are difficult to emulate on the bench.

Hardware

The PNE is a small desk-top device containing a touch-sensitive LCD display/input screen, two 100baseT connections, and a 6 VDC power input. It is normally supplied with a 120VAC wall mount power supply. There is also an unused DB15 port on the left side of the device. Please do not connect anything to this interface. PNE is capable of dealing with 802.1q tagged packets, but does not support jumbo-frame packets.



Connections

To power-up PNE, connect the supplied power supply to the DIN connector located on the back right corner of the device. PNE requires 6 VDC at a minimum of 1.5A.

PNE has two Ethernet ports located on the left side of the device. These ports are labeled P1 and P2. From an emulation standpoint, these two ports behave like a dumb Ethernet hub. Packets received on P1 are transmitted out P2. Likewise whatever is received on P2 is transmitted out P1. PNE must be

installed in the path between the devices under test. Usually it is installed as a *bump-in-the-wire* between the devices under test. However, it could also be installed between a pair of routers, switches, or a combination of the two. PNE is capable of dealing with 802.1q tagged packets, but does not support jumbo-frame packets.

There is an unused DB15 port on the left side of the device. Please do not connect anything to this interface.

Specifications

- Simulates IP networks between 1Kbps and 100Mbps
- Two 10/100 MDI/MDIX Ethernet interface
- 6-16 VDC, External AC adapter supplied with unit
- Power On/Off switch
- Firmware update menu option auto connects to DCB web site for updates
- Power requirements: 6-16 VDC, AC adapter supplied with unit
- Size: 7" W 2.5" H x 5" D

Test Configuration per Port

- Simulated network with rate limits from 1Kbps to 100 Mbps
- Packet latency of 0.1ms to 10 seconds
- Packet jitter of 0.1ms to 10 seconds
- All settings may be applied asymmetrically
- Induced packet loss, packet duplication
- Induced out-of-order packet delivery
- Generate errors automatically or on-demand
- Target errors at a specific device and protocol.
- Measuring work traffic rates in real-time.
- Measuring targeted network traffic in real-time.

Capabilities:

- Maximum bridging of 60,000 packet-per-second, bi-directional.
- Added packet latency from 0.1ms to 10 seconds Added packet jitter from 0.1ms to 10 seconds
- All settings may be applied asymmetrically. Induced packet loss, duplication, out-of-order delivery
- Errors generated automatically or on-demand
- Errors may be targeted at a specific device and protocol.
- Measures network traffic rates and targeted network traffic in real-time.

Performance

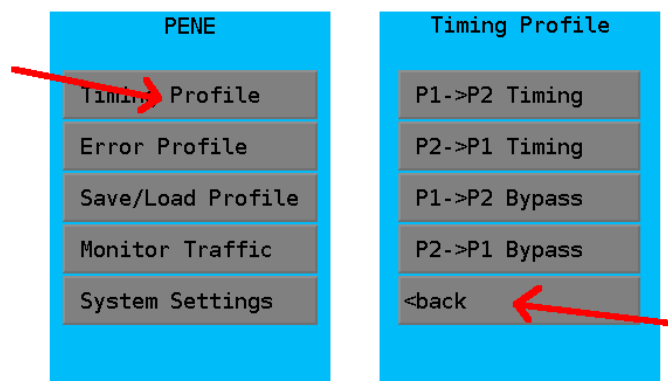
- Limits: Minimum packet jitter of 0.1ms
- Internal time resolution of 0.1ms

- Internal buffering for latency and jitter limited to 8000 packets
- Rate limited buffering limited to 104 packets.

Configuration

Navigation

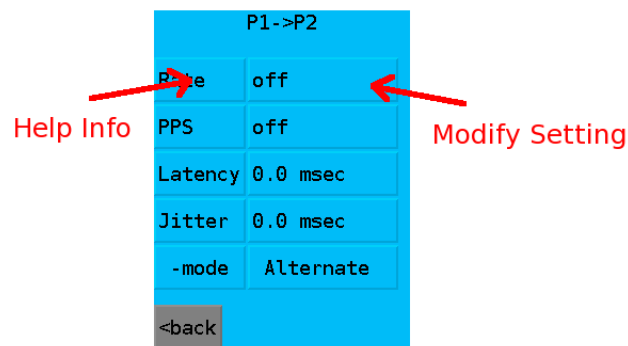
All configuration and control of PNE is performed through the LCD touch screen. To navigate the menu tree, simply press the desired menu button area. Sub-menus display a **<back** button that will return you to the previous menu.



Main Menu Display

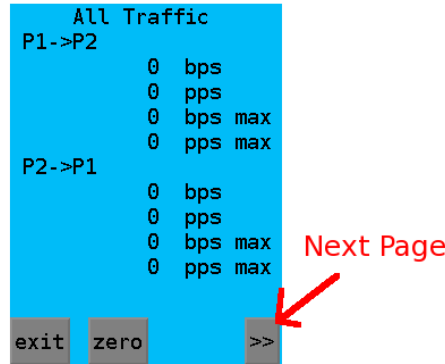
Test Menu Display

Configuration screens contain settings that can be modified. To modify a setting, press on the item button. Doing so will bring up an edit box or a pick list. All configuration items have context sensitive help. To access the help information, press on the configuration item's label.



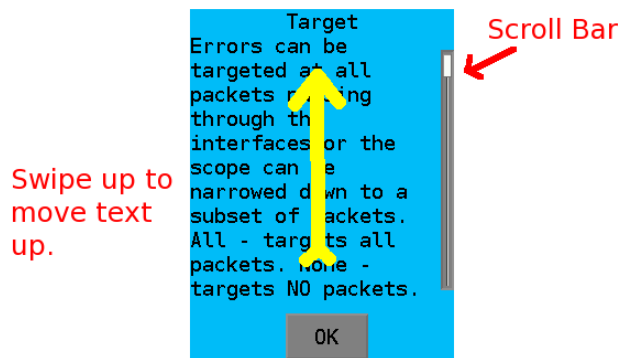
Configuration Display

Configuration and Status screens may have more information than can be displayed on a single screen. If a **>>** button is displayed on the bottom right corner of the display, pressing it will cycle through the pages.



Display Screen with *Next* Button

Some help and message screens may have more information than can be displayed on a single page. This is indicated by a scroll-bar on the right side of the display. To scroll the display, touch and drag your finger in the direction you want the page to move. Then lift your finger. The page won't move until you remove your finger from the display.



Display Screen with Scroll Bar

PNE maintains two types of configuration. First, the PNE Profile which contains configuration items related to PNE's timing and error emulation. These are the items that are commonly modified during the course of performing an emulation. The System Configuration, contain the configuration items related to the PNE device itself. These includes items such as the screen-saver time, the optional IP configuration, and the Ethernet Port link and duplex. These items are rarely modified once preferences are configured.

Profile Configuration

These are configuration items related to PNE's timing and error emulation that are commonly modified during the course of performing an emulation. Profile modification doesn't immediately take effect. For emulation purposes, it may be desirable to make a number of modification and then apply them all at the same time. An *apply* button will appear anytime there are profile modifications that have not been applied.

P1->P2	
Rate	off
PPS	off
Latency	0.0 msec
Jitter	0.0 msec
-mode	Alternate
<back	

P1->P2	
Rate	8.00Mbps
PPS	off
Latency	0.0 msec
Jitter	0.0 msec
-mode	Alternate
<back	
apply	

Profile Configuration Screen Profile Screen with *apply* button

The Profile is not automatically saved to non-volatile memory. To save the profile, you must navigate to the “Save/Load” profile menu and explicitly save the profile. A saved profile will automatically load on power-up. The Profile may be cleared without effecting the System Settings.

PENE
Timing Profile
Error Profile
Save/Load Profile
Monitor Traffic
System Settings

Profile
Save Profile
Load Profile
Default Profile
<back

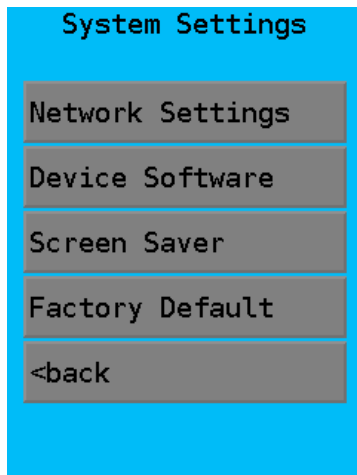
System Configuration

The System Configuration, contain the configuration items related to the PNE device itself. These


includes items such as the screen-saver time, the optional IP configuration, and the Ethernet Port link speed and duplex. These items are rarely modified once preferences are configured. The System Configuration is automatically applied when you exit a system-related configuration screen. It is also automatically stored to non-volatile memory and will become the power-up default.

System Settings

Navigate to the *System Settings* screen from the *main menu* screen *System Settings* button. From this screen there are options for network settings, updating the device software, screen saver, and resetting to factory default.



System Settings



Network Settings

Network Settings

From this navigation screen, select the ethernet port to configure, IPv4 parameters, and check network status.

Ethernet Configuration

Configure each ethernet port for allowable modes to allow or disable Auto Negotiate, 100Mbps or 10 Mbps, and Full or Half duplex.

Ethernet	
Mode	Auto Negotiate
100Mb-FD	allow
100Mb-HD	allow
10Mb-FD	allow
10Mb-HD	allow
exit	

Ethernet Configuration

IPv4 Settings Screen

It may be useful to assign an IP address to the PNE. Having an IP address allows the PNE to automatically download updated firmware, and respond to typical IP tests such as PING. However, if those features aren't needed, PNE works without an IP address assignment.

From this screen, select an IPv4 mode by clicking on the *Mode* value field. Options are *DHCP*, *static IP*, or *Disable IPv4*. If Static is selected, enter the IP configuration values. If disabled or DHCP, the IPv4 Settings are blanked.

IPv4 Settings	
Mode	static
IP	192.168.1.223
Mask	255.255.255.0
GW	192.168.1.1
DNS	8.8.8.8
exit	

IP Settings Screen

IP Config. Mode
Select DHCP Mode
Select Static Mode
Disable IPv4


Mode Pick-List Screen

IPv4 Settings	
Mode	Disabled
exit	

Disabled IPv4 Screen

Network and Device Status


This screen displays the status of all Network and IPv4 configuration values.




Network Status Display

Device Software

From the Device Software screen, the Software Info display can be selected, an automatic software update initiated, or the system rebooted.



Device Software Menu



Device Info Display

Network Emulation and Impairment Tests

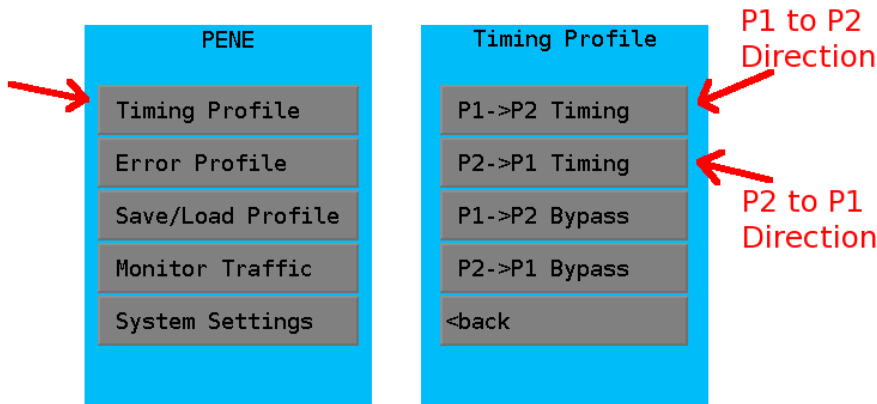
Timing Emulation

A core feature of PNE is its ability to manipulate packet timing as packets flow through the device. At the lowest level, PNE is simply moving packets from P1 to P2 and likewise from P2 to P1. In this regard, it functions very similar to a store-and-forward switch. However, while it is storing a packet, PNE can be configured to delay the forwarding action. By delaying the forwarding action, it is possible to emulate the timing characteristics of a real-world link. These characteristics are transmission rate, latency, and jitter.

Transmission Rate-Limit

A network connection has a maximum speed in which it can operate. For example 100M Ethernet has a maximum speed of 100M bps (bits-per-second). A DSL link may run at a much lower rate, for example, it may only run at 1M bps. In addition, the speed of the link may not be symmetrical. Again, as an example, a DSL link may have a download rate of 1M bps and an upload rate of only 500K bps.

To configure the transmission rate limit, navigate to the *Timing Profile* submenu. From this menu, the timing for packets flowing in P1 and out P2 and the packets flowing in P2 and out P1 may be independently configured.



Timing Profile Configuration Timing Profile Details

Rate limiting is off by default. To set a rate, press on the current setting and enter a new value. Transmission rate limiting can be set in the range from 0.001M to 100M. **The value 0 disables rate limiting.**

P1->P2

Rate	off
PPS	off
Latency	0.0 msec
Jitter	0.0 msec
-mode	Alternate
<back	

Rate Limit Mbps

1

7	8	9	clr
4	5	6	del
1	2	3	esc
0	.	enter	

Enter the new rate and press enter.

Timing Profile Configuration

Rate Limit Configuration

PNE uses two different methods for emulating the transmission rate. For rates 40Mbps and below, the rate is simulated in software on ingress to the device. This means the rate is simulated before applying latency and jitter. Rates below 40 Mbps are available down to 1Kbps.

To handle rates above 40Mbps, PNE needs the assistance of a hardware rate limiter. The hardware rate limiter emulates the behavior of a rate limited link by clocking the data bit-by-bit at the selected rate. However, due to hardware constraints, it is applied at packet egress. When emulating a low speed link, rate limiting on egress may effectively negate jitter, defeating the jitter configuration describe in the next section. At high rates, this smoothing effect is minimized. In addition, the hardware rate limiter has limited granularity. For rates above 40Mbps, the actual rate will be rounded to the nearest multiple of 250K and will vary slightly based on packet size.

Latency and Jitter

Latency is the time it takes a network packet to travel from the sender to the receiver. In a bench-test environment, this is often very fast, usually a fraction of a millisecond. However, in a real-world environment, the network packet may be traveling a long distance, traversing multiple routers or bridges. Latency may range from 10s of milliseconds to 100s of milliseconds.

Actual latency is rarely a fixed amount of time. There is a fixed component since it will take a fixed amount of time for a packet to physically traverse the network. However there is also a variance to the latency. This variance is called jitter. It is a result of varying network congestion and processing limitations as a packet makes it way from sender to receiver.

PNE allows a fixed latency plus a variable amount jitter to be configured. As packets flow through the device, each packet will be delayed by the latency time. In addition, some jitter time will be added to the base latency. There are four different algorithms used for applying jitter. They are *alternate*, *random*, *sweep* and *random delay*.

P1->P2	
Rate	off
PPS	off
Latency	0.0 msec
Jitter	0.0 msec
-mode	Alternate
<back	

Latency and Jitter Configuration

Alternate mode: The full amount of jitter is applied to every other packet, basically going from min (0) to max with each alternating packet.

Random mode: A pseudo-random time between 0 and the selected jitter time is applied to each alternating packet*. The resulting jitter should be approximately half of the selected jitter time.

Sweep mode: The delay added to each packet sweeps back and from no delay to the selected time. For example, if jitter was set to 100ms, the delay would start at 0 and progressively grow to 100ms. Once it reaches 100ms, it would start to decay back to 0. When it reaches 0, the process repeats. The time it takes to sweep is not related to clock time but instead is based on packet flow. Each packet steps the delay by 0.1ms.

Random delay mode: This mode is a variation of Random mode. However, instead of applying jitter to alternating packets, the random jitter is applied to random packets, with a probability of 1/1000.

*Note: The implementation of Random and Sweep jitter was modified with the release of firmware v1.02, thus will yield differing results. Previous firmware applied the jitter time to every packet instead of alternating packets, resulting in a flattening effect.

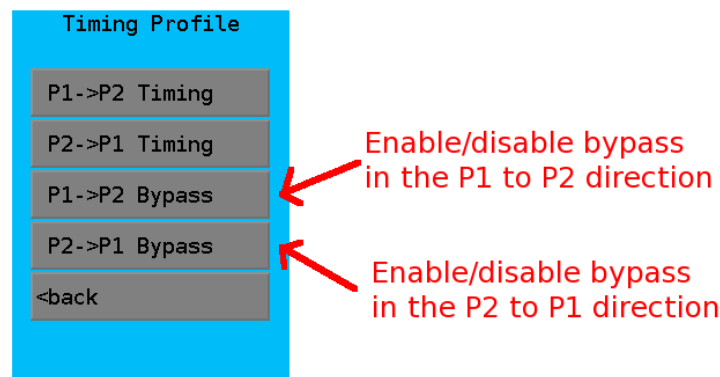
Which one to use? Consider the real-world network you are emulating and try to match it. Using DCB's companion product, the ETTA, to characterize the real network the best way to understand it. Other methods are simply observing the delay and jitter by running ping tests over time. For new equipment analysis, it may help you understand the equipment responses by running tests using all three modes

When applying jitter, make sure to take in account the data rate. The PNE device can only generate a jitter effect if there is sufficient bandwidth headroom to support the tested data rate including the jitter. If there is not sufficient headroom, the PNE will buffer the data which will flatten out the jitter effect. Given enough time, the PNE buffer will overflow resulting in packet loss.

Timing Bypass

PNE operates at layer-2. Doing so greatly simplifies the test environment. However, it also means that PNE will apply rate limiting and packet delay to protocols that are not designed or intended to be delayed. One such example is the ARP protocol. ARP does not cross router boundaries. When emulating the timing of a routed network, it may be unreasonable to apply the same timing to ARP packets. Some devices have fixed ARP time-outs that are quite short and will fail when transiting an emulated high latency link.

PNE can be configured to bypass all non-IP traffic, immediately moving the packet from the receiving port to the transmit port. The packet is not subject to any additional delay, except for the minor store-and-forward delay imposed by PNE. Also, these packets will be excluded from error emulation. Non-IP traffic includes any packet that does not carry an IPv4 or IPv6 protocol ID in the Ethernet header.



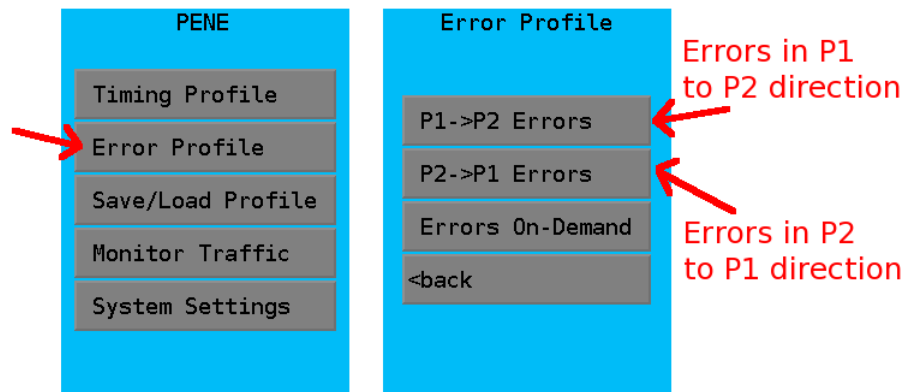
Timing Profile Bypass Configuration

Please keep in mind that when emulating a bridged environment, non-IP traffic **should be** subjected to timing delays and error induction. This is often overlooked when bench-testing and a cause of field failures when ARP, STP, and other layer-2 protocols behave unexpectedly in bridged environment with high latency. Examples of this type of network are WiFi or other wireless bridges.

Error Emulation

Transmission errors are a fact-of-life in a real-world environment. Packets may be dropped, due to network congestion or interference on a wireless link. Wireless links are prone to packet duplication, and multipath routing may cause packets to be delivered out-of-order.

Similar to timing, errors are configured asymmetrically. The specific rates are independently set for each direction of packet flow.



Error Profile Configuration Error Profile Direction

PNE has a very simple but flexible mechanism for automatically generating various errors. For each of the error types an error ratio is configured.

P1->P2 Drop		P1->P2 Dup		P1->P2 Order		P1->P2 Corrupt	
Drop	0	Dup	0	Reorder	0	Corrupt	0
per	100	per	100	per	100	per	100
		Latency	10.0 msec	MaxWait	100.0 msec	offset	24
<back apply		<back apply		<back apply		<back apply	

Dropped Packets

Duplicate Packets

Out-of-Order Packets

Corrupted Packets

Drop Errors

For example, P1 to P2 drop may be configured to drop 1 packet per every 100 total packets, yielding a 1% drop rate. Likewise, drop rate may be configured to drop 10 packets every 1000 total packets. This is still a 1% drop rate, but with very different implications. Instead of a single packet lost, a block of 10 packets are lost.

To generate a drop error, a packet is simply dropped instead of being transferred from the input port to output port. The *drop* value configures the number of packets to drop in a block. It must be smaller than the *per total* value. The per total value may be set anywhere from 1 to 100,000,000.

Duplicate Packets

A duplication error is generated by sending the same packet twice. The *dup* value configures the number of packets, within a block, to duplicate. It is not the number of times that a packet is duplicated. For example, if the *dup* value is set to 10, a block of 10 packets is duplicated one time. The *per total* value controls how frequently to generate the duplicate packets. So a dup of 10 per 100 means a block of 10 packets will be duplicated once every 100 total packets. The amount of time between sending the first instance of a packet and sending the second instance of a packet is specified as the *latency*.

Out-of-Order Packets

The behavior of out-of-order packets is a little more complex than the other error types. The concept is simple, hold on to a packet, wait until another packet comes along, then deliver the two packets in reverse order. A block of packets can also be reordered. Along the same idea, two or more packets can be placed in the reorder buffer. Then as additional packets come along, the packets in the reorder buffer are merged back into the packet stream. To really mix this up, the packets are merged back in reverse order from their original receipt.

To explain this a little better, let's take the packet sequence of p1, p2, p3, ... p9. If reorder is set to 3, packets p1, p2, and p3 are placed in the reorder buffer. As the remaining packets arrive, the transmit order will be as follows: p4, p3, p5, p2, p6, p1, p8, p9.

To complicate things, we have to deal with the possibility that “another” packet may not come along in a reasonable amount of time. With streaming type protocols this is an unlikely event. However, lock-step type protocols are prone to this behavior. It's not that another packet will never come along, but that it may take a long period of time before this occurs. The system under test will be handling this as a dropped packet condition instead of an out-of-order condition. The *MaxWait* configuration item addresses this problem. It specifies how long a packet should sit in the reorder queue before being kicked out based on time.

Corrupted Packets

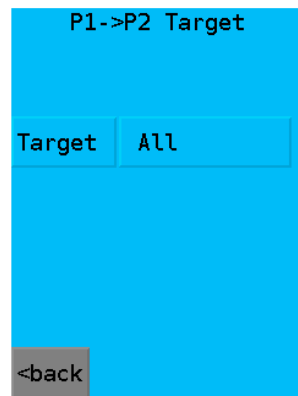
This error will flip the low-order bit of a single byte within the packet. The user may select which byte within the Ethernet frame as the target. The intended purpose is to generate a checksum data in a higher level protocol. For example, corrupting the byte at offset 24 will cause an IP header checksum error in a standard Ethernet frame carrying an IPv4 packet.

Please note that this feature cannot be used to generate an Ethernet packet CRC error. The PNE hardware recomputes the CRC on packet egress, which will correct the Ethernet CRC of the modified packet.

Targeting Errors At A Specific Device

Errors may be directed at a specific device and/or protocol instead of being applied to all packets that traverse the ports. This allows testing to be focused on a particular application and insure that errors are generated against packets related to that specific application. Targeting can be based on Ethernet MAC addresses or by IPv4 addresses. When using IPv4 addresses, its possible to drill down on a specific TCP or UDP protocol.

Like all other settings, targeting is configured independently for each packet direction, P1 to P2 and P2 to P1. By default all packets are targeted. The options are to target *all* packets, *no* packets, target by Ethernet *MAC address*, or to target by *IPv4* addresses.

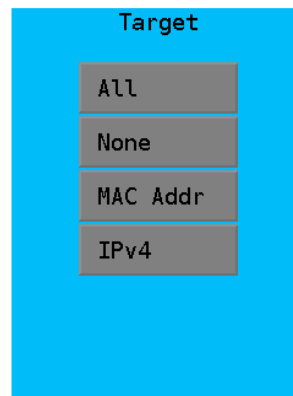


P1->P2 Target

Target	All
--------	-----

<back

Target Selection

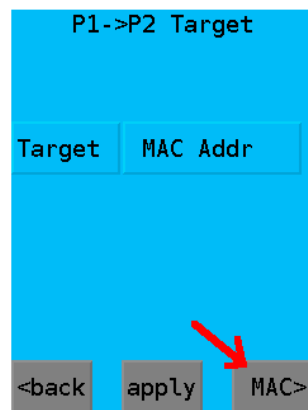


Target

All
None
MAC Addr
IPv4

Target Type Selection

MAC Addr: Packets are targeted by source and destination Ethernet MAC address.



P1->P2 Target

Target	MAC Addr
--------	----------

<back apply MAC>



P1->P2 Target

Src	any
Dst	any

<back apply

Target MAC
Addresses

IPv4 Address: Packets are targeted by source and destination IPv4 address. When this mode is selected, it is also possible narrow the scope by also targeting protocol and port number.

P1->P2 Target

Target	IPv4
--------	------

<back
apply
IPv4>

P1->P2 Target

SrcIP	any
DstIP	any
Proto	All

<back
apply

Target Type
Target IP and Protocol

IPv4 Targets

Target by protocol

If the Protocol field is set to TCP or UDP, a specific application can be targeted by setting the port or a range of ports used by the protocol. For example, lets say we want to target a web server attached to P2. HTTP uses TCP port 80. The P1->P2 target would be set to TCP, and the Dst Low and High would be set to port 80. We would leave the Src Low and High set for the full range of 0 – 65535 since the web browser will be using an ephemeral port as it's source. This example target will match any packet moving from P1 to P2 with a destination of TCP port 80.

Protocol

All

TCP

UDP

P1->P2 Target

SrcIP	any
DstIP	any
Proto	TCP

<back
apply
port>

P1->P2 Target

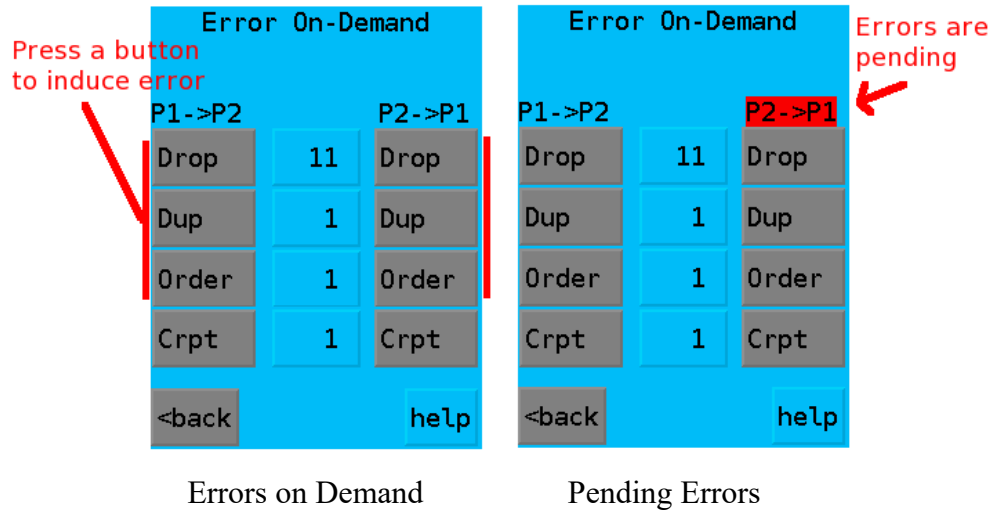
Src Low	0
High	65535
Dst Low	80
High	80

<back
apply

Protocol Targeting
IP Address Targeting
IP Port Targeting

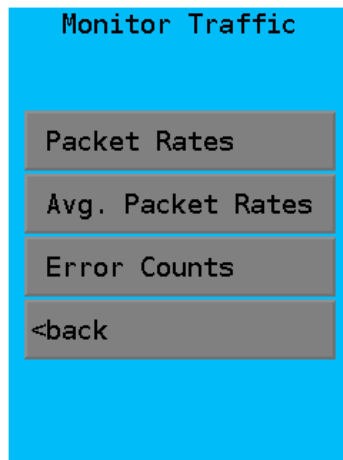
Errors On-Demand

PNE can also be used to generate errors on-demand. Each button press will generate the requested number of errors in either the P1 to P2 direction or the P2 to P1 direction. The same targeting rules, described above, apply to errors on-demand. The title above the buttons will highlight in red while waiting for enough target packets to satisfy the demand request. Keep in mind, if no packets meet the target requirements, the errors won't be generated.



Monitoring

When emulating a network, it is helpful to get some feedback. PNE gathers statistics on traffic flow. This include separate statistics on total traffic flow and targeted traffic flow. Both current and long-term average rates are gathered. In addition, a count of induced errors are kept.



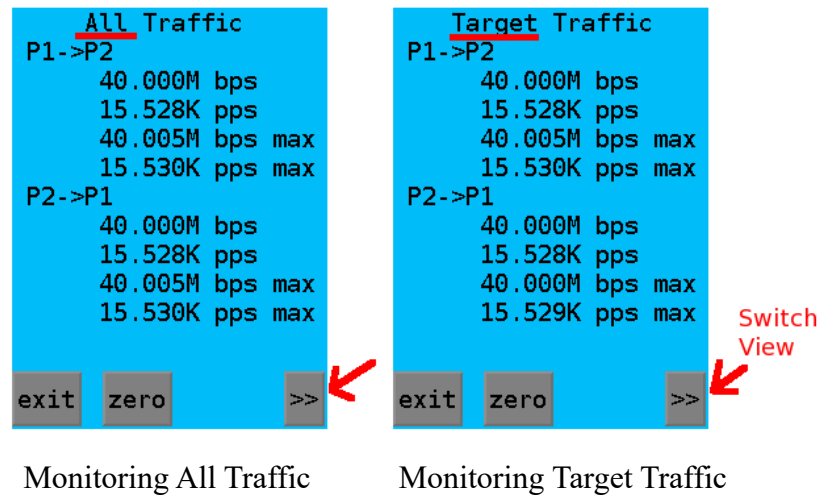
Traffic Monitor Selection

Packet Rates

Packet Rates show the traffic flow over the last 1 second sampling interval. The *max* fields keep track

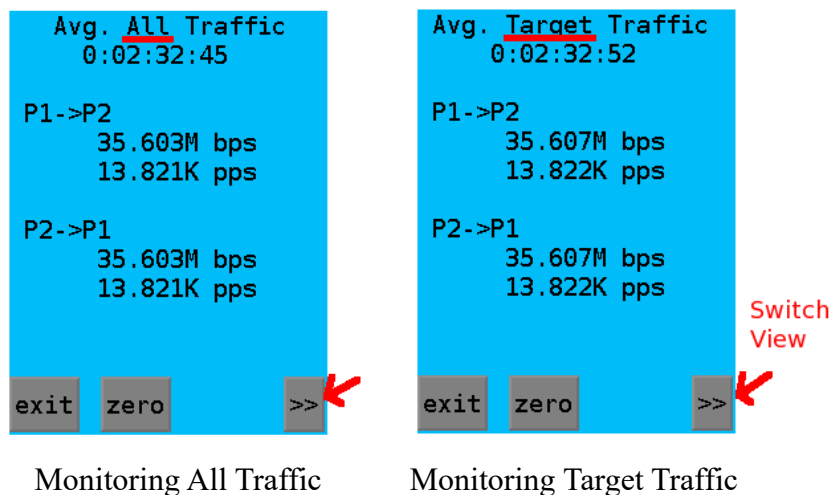
of the highest level since start-up or the last time the *zero* button was pressed. This is useful in capturing the short-term bursts.

Pressing the **>>** button will toggle between all traffic and targeted traffic. This information can be used to verify the effects of rate limiting and latency. It can also be used to measure the load a device or application generates under various condition.



Average Packet Rates

Average Packet Rates show the flow of traffic since start-up or the last time the *zero* button was pressed. Please note that pressing the *zero* button will reset all statistics.



Induced Errors		Buffer Errors	
P1->P2		P1->P2	
drop:	11	pps:	0
dup:	1	rate:	0
order:	2	tx_q:	0
P2->P1		P2->P1	
drop:	157	pps:	0
dup:	3	rate:	0
order:	1	tx_q:	0
exit	zero	exit	zero
>>		>>	

Switch
View

Monitoring All Traffic

Monitoring Target Traffic

Interpreting these displays

All values are for packets transversing a single direction (Port 1 to Port 2, or Port 2 to Port 1).

Drop: Dropped packets

Dup: Duplicate packets

Order: Out of order packets

Pps: packets dropped due to ingress rate exceeding the configured pps limit.


Rate: Packets dropped due to ingress queue overflow. This error will be a result of the ingress rate exceeding the configured rate limit. For example, the rate limit is set to 10Mbps and the network is trying to push 20 Mbps.

Tx_q: Packets dropped due to egress queue overflow. This may be caused by hardware rate limiting (rate limit > 40 Mbps) or by a speed/duplex mismatch between ethernet ports. For example, one port is running 100Mbps and the other 10Mbps.

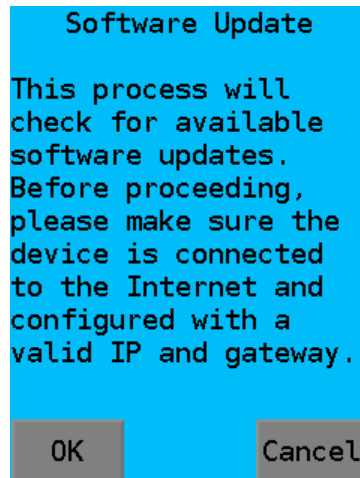
Firmware Updates

PNE has a built-in easy-to-use update feature. Whenever new firmware is released, it's made available for download via the Internet from a DCB server. Update requires an Internet connection and an appropriate static or DHCP IPv4 address assigned to the PNE.

You can also display the current software version on the “Device Info” screen, accessible from the “System Settings” screen.



System Settings Screen




Software Update Screen

From the main screen, press “Settings”.
From the Settings screen, press “System”.
On the System Settings screen, press “Update Software”.

On the Software Update screen, press “OK” and the “Check for Software Update” screen is displayed.


To obtain the latest firmware, there is no need to change any values in this screen. Press “start”.

If you wish to use any other version of software, enter the version number. The HTTP Port should remain configured to port 80. Press “Start”



Software Update Configuration

The unit will display the “Checking..” screen and connect to the firmware server to download the firmware image. This screen may only display for a second or two.



Firmware Download Complete

After downloading the image, the unit will check that an update is needed. If so, it will load the new firmware and restart.

During the download process if you decide not to complete the update, press the “abort” key.

Once the image is downloaded, you can cancel the operation or install the software. If you accept the install, PNE will update itself and restart.

Previous versions of PNE firmware download files are also available from the update site. If you want to downgrade to a prior version or obtain a custom version, enter the requested version number prior to pressing “start”.

If the requested firmware version isn’t available or the update server can’t be reached, an error screen is displayed, along with a return path to the System Settings screen.

Use Examples

Using the PNE Network Emulator To Test E&M and TDM Replacement Products November, 2018

Many of our customers are faced with the “disappearing E&M line blues” as telephone companies withdraw their leased E&M line offerings from the marketplace. They are replacing those with digital lines, typically T1 or packet switched networks (basically ethernet). In some cases TDM networks are being replaced with packet switched networks.

We used PNE to analyze a few TDM and E&M replacement products to help those customers, and in some cases the equipment manufacturers, better characterize the performance of new circuit transport via real world ethernet connections.

The test setup is simple. Two CPE end products are connected via the PNE ethernet ports. Traffic is passed through those CPE products and the ethernet network between them. The tests are straight forward. TDM traffic was generated using a T-Berd, E&M traffic was generated using modems. The results are sometimes surprising, sometimes as expected.

TDM (T1) Product Tests.

Test 1: Dropped packets for T1 transport emulation: With jitter buffers and jitter tolerance set at manufacturer’s recommended values, we configured the PNE to drop one packet per each thousand packets. MFRA’s product always generated an error.

Test 2: Duplicate packets for T1 transport emulation: Same equipment as above, and we duplicated one packet per hundred. It always generated an error. We then lowered the error rate to one packet duplication per thousand. Same thing... the hit always generated an error.

Test 3: Out of order packets for T1 transport emulation: One packet out of order per thousand packets. No error was generated, but when we raised it to any number greater than one out of order packet per thousand, there was an error generated.

Test 4: Packet jitter: Jitter didn’t cause a problem up to 20 msec. But, since the jitter buffer setting was 24, we know there will be problems with jitter values near 24 msec or greater.

E&M Product Tests

We used a 202T FSK modem over an emulated E&M channel. RFC 5087 – Time Division Multiplexing over IP (TDMoIP) provides various ways for a manufacturer to handle packet switched network real-world problems. It allows for either configurable or dynamically adjusted jitter buffers, and requires a lost packet processing method.

Test 1: Out of Order Packets: We discovered that the CPE E&M tunnel device does not tolerate any out of order packets.

Test 2: Duplicate Packets: It was more tolerant of duplicate packets. The link worked with ten packets duplicated per hundred packets transmitted.

Test 3: Packet Jitter: The unit under test automatically adjusts its jitter buffer to compensate for network jitter. When the jitter value is changed very quickly, a single error was generated, and then error free operation resumed. When we bumped the jitter from 30 msec up to 600 msec, then back down to 30 msec, again there was a single error with each jump, but it's obvious that the jitter buffer is elastic and grows and shrinks as necessary.

And a Product Verification

While this one is a bit self-serving, we'll mention one more test suite we recently performed. We tested two of our AVA-E low bit-rate VOIP boxes. The AVA-E configuration was typical..

Interfaces configured to FXS to FXS–PLAR the P25 3600 bps voice CODEC. Jitter buffer is 40 msec.

For the test, timing was set for 3 msec latency, 100 msec jitter, and random jitter. We tested for typical Internet problems such as drop one packet per thousand, duplicate one packet per thousand, and move three packets per second out of order.

We found no perceptible impact on the voice traffic. Of course, we attribute this to good design work. But, buffering in the network interface engine and CODEC choice are the primary reasons that network imperfections didn't affect the end to end traffic in this robust product.



PNE Testing Two AVA-E VOIP Units

Conclusions

PNE enables you to learn things about your ethernet products that aren't readily apparent. With a low entry cost, it's convenient to characterize how YOUR installation will operate on a less than perfect real world network. This appliance provides a solid methodology for comparing alternate products as well as learning how your installation will work in the real world. By measuring the actual network path characteristics using the companion product, the ETTA ethernet transport analyzer, and configuring the PNE appropriately, you can test your proposed equipment installation on the bench using the PNE with repeatable measurements.

About the PNE

The PNE (Packet Ethernet Network Emulator) is available from Data Comm for Business, Inc (DCB) either directly or through resellers such as Graybar or Anixter. At only \$995, and with a small desktop footprint, there should be one on each technician's test bench.

Read the PNE data sheet at <https://www.dcbnet.com/datasheet/pneds.html>

More information about testing is in the PNE manual at <https://www.dcbnet.com/manuals/pne.pdf>

A companion product ETTA, the ethernet performance test set, is used to characterize an existing network. Read the ETTA data sheet at <https://www.dcbnet.com/datasheet/ettads.html>