ETEP Release Note - Version 2.2.2

This release note describes the new features, system requirements, known issues, and bug fixes in ETEP Version 2.2.2.

New Commands

Added “show version detail” command

A new show interfaces output displays not just the information from ifconfig, but other interface-configuration information that is relevant to the ETEP.

```
admin> show interfaces local.
Comnmand : show interfaces [local | remote | management]
```

If no arguments are specified, then output for all the interfaces will be shown.

```
Interface Name: local, Physical interface:  eth1, Admin:  Up, Link status: Up
Auto-negotiation:  enabled,  Speed:  10Mbps, Duplex:  half, Flow control:  on
Link encapsulation:  Ethernet  HW Address:  00:98:70:0F:FA:BD  VLAN: 0
Tx enable: follow-rx
Operating mode:  Transparent Reassembly mode:  gateway DF-bit ignore: enabled
DHCP relay: disabled
Supported link modes:  10baseT/Half, 10baseT/Full, 100baseT/Half,
                       100baseT/Full, 1000baseT/Full
Advertised link modes: 10baseT/Half, 10baseT/Full, 100baseT/Half,
                       100baseT/Full, 1000baseT/Full
Traffic statistics:
Rx bytes:             4294967295
Rx pkts :               79250590
Tx bytes:             2244121276
Tx pkts :               32006893

Input Errors:
Errors:         0
Dropped:        0
Overruns:       0
Frame:          0
Output Errors:
Errors:         0
Dropped:        0
```

Added support for “clear discards” command

The CLI adds support for a “clear discards” command to clear the discard counters under the “top admin” menu. This command is useful for debugging the specific appliance locally.

ET0005A

The ET0005A is a multi-layer encryption appliance that provides tunnel-less data protection, including Ethernet frame encryption for Layer 2 networks, IP packet encryption for Layer 3 networks, and Layer 4 data payload encryption for IP and MPLS networks. The ET0005A offers full-duplex encryption at speeds up to 5 Mbps using the AES-256 algorithm.

The ET0005A is a cost-effective encryption appliance that is fully interoperable with the Black Box ETEP family of multi-layer encryption appliances. The ET0005A uses the same standardized and FIPS 140-2 validated cryptographic algorithms as the ETEP family. Environmental hardening for enclosed outdoor environments and a very small form factor allow the ET0005A to be deployed at the network edge where encryption was previously difficult to deploy and manage. Furthermore, the ET0005A includes an integrated Ethernet switch, so it can switch traffic locally while encrypting traffic to and from the wide area network. This reduces costs by avoiding the need to deploy additional Ethernet switches.

MPLS encryption

EncrypTight Manager encrypts MPLS at Layer 2 for any encapsulated payload (Ethernet, IP, ATM, FR, Circuit Emulation).

We are solving the problem of encrypting Ethernet frames with the MPLS Ethertype.

The encryption procedure would encrypt the MPLS payload, while leaving the Ethernet header and MPLS stack in the clear. The core (P) LSRs could use the MPLS label stack, but the payload would be encrypted and opaque to the core LSRs. This is fine because the core doesn't look at the payload, and the payload is decrypted on the other side of the network before it reaches the other PE LSR. If configured, the encryption procedure would skip the psuedo-wire control word (first four bytes after the last MPLS label).

The process for MPLS Encryption is:

- Encrypt the MPLS payload, while leaving the Ethernet header and MPLS stack in the clear.
- Insert a new “encrypt” label (reserved label 12) at the bottom of the MPLS stack
  - Clear the BOS bit on the previous BOS label
  - Set the BOS bit on the new label
- Encrypt the original payload, starting after the new “encrypt” label
- Supports zero or one 802.1Q VLAN tags in the outer header.

NOTE

Please note that when incoming MPLS traffic does not contain the BB reserved label, MPLS traffic is dropped, not passed in the clear.

Remote in-band management

Allows devices to be managed remotely without using external switch ports to connect to the management port

ETEP has a single management interface and single management IP address

- Either out-of-band management or RIBM must be enabled, and both cannot be enabled at the same time
- If RIBM is enabled, then the management IP interface of the ETEP is reachable only via the remote port.
- If RIBM is enabled, then the local port cannot be used to manage the device

RIBM traffic passes through the data plane, so the administrator must set “pass TLS in the clear” or configure a policy to allow it to pass through the data plane

Transparent or non-transparent mode

All policies supported (L2, L3, L4) (except L2 IKE is **not** supported)

Remote IP and Management IP can use the same address

No IPv6 for RIBM

How RIBM affects ETM

With the release of Version 2.2.2 PEP software special policy design considerations must be taken into account for ETM access to ETEP devices for protocols such as FTP and SSH. Before the release of Version 2.2.2, all ETEPs were managed Out of Band. The management port was typically connected to the Local port LAN segment and followed the rules defined by the policy for that LAN segment. In most cases, TLS traffic was always passed in the clear and other management protocols such as FTP were encrypted and decrypted using the configured policy.

With the release of Version 2.2.2 PEP code and the ET0005A, RIBM may change behavior depending on the addressing scheme used for the management IP addresses of the ETEP. If the management network is configured for a different IP subnet than that of the Customers LAN segments, protocols such as FTP and SSH will be blocked since the RIBM IP address does not match the configured policy. In this scenario, protocols such as FTP and SSH will be dropped by the remote ETEP and prevent upgrades, appliance log file retrieval and SSH access to the remote device.

NOTE

Although this may be a corner case in transparent mode, it will always be the case when the ETEP is in Non Transparent mode where the Management IP address will in most cases be the same IP address as the Remote port and not match the customers encryption policy.

Any operation that only requires TLS and XMLRPC will not be affected. For example, refresh status, policy deployments, statistic gathering etc..,

To prevent the potential problem defined above, an EZMesh Policy MUST be configured as the highest priority policy or before any encrypt policy using the wildcard address of 0.0.0.0 and the IP address of the ETM server. If ETM is configured as a cluster, include all /32 addresses of each node in the cluster.

Below is a screen shot of an example Easy Mesh policy that prevents the issue noted above. Also below is a diagram depicting the topology or addressing scheme that would require the Easy Mesh policy.

Figure 1 Topology Requiring Easy Mesh Policy

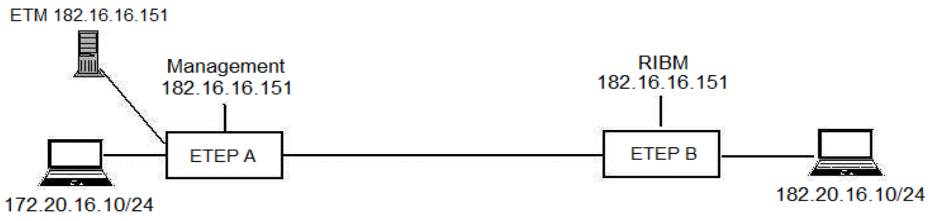

Figure 2 Add Easy Mesh Policy Menu

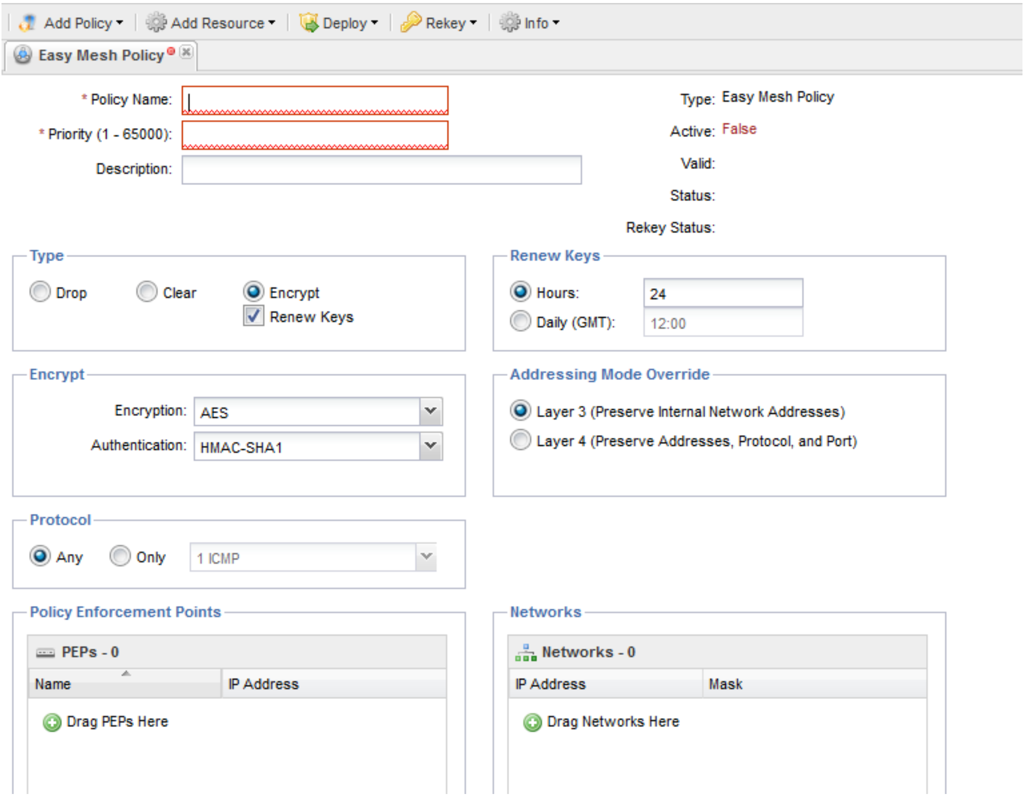

* Policy Name:

* Priority (1 - 65000):

Description:

Type: Easy Mesh Policy

Active: False

Valid:

Status:

Rekey Status:

Type

○ Drop ○ Clear ● Encrypt

☑ Renew Keys

Renew Keys

● Hours: 24

○ Daily (GMT): 12:00

Encrypt

Encryption: AES

Authentication: HMAC-SHA1

Addressing Mode Override

● Layer 3 (Preserve Internal Network Addresses)

○ Layer 4 (Preserve Addresses, Protocol, and Port)

Protocol

● Any ○ Only 1 ICMP

Policy Enforcement Points

PEPs - 0

Name	IP Address
Drag PEPs Here	

Networks

Networks - 0

IP Address	Mask
Drag Networks Here	

GA release of ETVEP

The ETVEP is a virtual appliance for VMWare ESX/ESXi environments that enables sensitive workloads to execute and communicate securely in untrusted networks. The ETVEP provides data confidentiality and integrity for sensitive data in motion in shared environments and prevents one tenant from monitoring the network traffic or attacking the virtual servers of another tenant. Furthermore, the ETVEP allows the data owner or a trusted third party to control the encryption keys without the need to share the encryption keys to the infrastructure provider.

vCPU licensing

For Enterprise customers, ETVEP licenses will be one-time perpetual licenses based on the number of virtual appliances and number of vCPUs allocated to the virtual appliance. Customers will pay a license fee for each ETVEP that they operate, and for each ETVEP, they can choose a ETVEP license that specifies a maximum number of virtual CPUs supported for that appliance.

NOTE

Moving or cloning a ETVEP results in clearing the license of the ETVEP (it goes back to zero CPUs and acts as a wire).

If a licensed ETVEP boots with more vCPUs allocated than are specified in its license, then it will behave as if it has a zero CPU license until the number of allocated vCPUs <= licensed vCPUs.

NOTE

EncrypTight Manager will report the following PEP states as errors:

- **“License is Over-Subscribed”**: if the license is set for less CPUs than currently allocated to the ETVEP. No encryption will happen in this case. Policy configuration changes should still be possible but will not get pushed to the dataplane.
- **“License is Under-Subscribed”**: this means the license is set for more CPUs than currently allocated. The ETEP will operate as normal, but at reduced throughput. This should be colored Orange.
- **“No License”**: ETVEP will act as a wire

Operational Notes

NOTE

The ET0005A does not support flow control on the local and remote ports. Flow control has a separate negotiation which is not part of speed and duplex negotiation. We treat them as one and either enable or disable both.

Known Issues

ET0005A: Remote ports wired together always comes up with 10Mbps/Half duplex [ETEP-1250]

Auto-negotiation is “off” (disabled) by default on the remote port. To turn “on” (enable) auto-negotiation on the remote port, configure autoneg off and then autoneg on. Remote ports if wired together with default settings (autoneg “on”) always comes up with 10Mbps/Half duplex setting, instead of 1Gbps/Full Duplex.

NOTE

To enable auto-negotiation, auto-negotiation must first be set to “off” and then set to “on”. This serves as a “reset” of the ports to the expected condition of “on” (enabled).

Workaround:

Although all ports *should* be configured for auto-negotiation “on”, the remote ports default to the “off” (disabled) condition. To ensure the state change occurs, auto-negotiation must first be set to “off” (disable) and then to “on” (enable) on the remote ports. This ensures the auto-negotiation on these ports will be “on” (enabled).

1 Set auto-negotiation on remote ports to “off” (disable).

2 Set auto-negotiation on remote ports to “on” (enable).

ET0005A: PEPd crashed with readDataplaneMsgComplete timeout [ETEP-1249]

Issue occurs when ET0005A's policy mode is switched from Layer3 to Layer2. This is an intermittent problem.

Workaround:

Reboot the appliance to clear the failure.

ET0005A: Switch ports don't come up with autoneg 'off' and speed set to 1Gbps [ETEP-1241]

The switch ports will not establish a 1G link when auto-negotiation is “off” (disabled). This applies to all switch ports. If auto-negotiation is “off” (disabled), all switch ports will establish a link at 10M and 100M, but not at 1G. If auto-negotiation is “on” (enabled), a 1G link can be established as expected.

ET0005A: By default switchport 2 and 3 are being configured for autoneg “off” [ETEP-1240]

The switch ports 1 and 4 default to auto-negotiation “on” (enabled). The switch ports 2 and 3 default to auto-negotiation “off” (disabled). In order to enable auto-negotiation on switch ports 2 and 3, the commands to disable and then re-enable auto-negotiation must be applied. By default all ports should be configured for autoneg 'on'.

NOTE

To enable auto-negotiation on switch ports 2 and 3, auto-negotiation must first be set to “off” and then set to “on”. This serves as a “reset” of the ports to the expected condition of “on” (enabled).

Workaround:

Although all ports *should* be configured for auto-negotiation “on”, ports 2 and 3 default to the “off” (disabled) condition. To ensure the state change occurs, auto-negotiation must first be set to “off” (disable) and then to “on” (enable) on ports 2 and 3. This ensures the auto-negotiation on these ports will be “on” (enabled).

1 Set auto-negotiation on switch ports 2 and 3 to “off” (disable).

2 Set auto-negotiation on switch ports 2 and 3 to “on” (enable).

Upgrades fail if no subdirectory is given [ETEP-1235]

Upgrades will fail if trying to scp/ftp from root directory because the underlying code doesn't properly handle an empty string for the local download directory.

Workaround:

To install the build under a subdirectory.

Customers can no longer ping their local gateway if ETEP is gateway [ETEP-1229]

You cannot ping the ETEP's local port if an IP address has been assigned. This is different behavior from earlier releases.

Using ETVEP and ET0005A with VIP mode enabled, customers can no longer ping their local gateway if ETEP is a gateway.

This is because the local interface no longer has an IP address. In order to allow the local interface to have an IP, we need a set of rules to implement static MAC-NAT and to handle ARPing properly since by default both Linux and the dataplane will respond to ARP requests.

Workaround:

Here's the full set of rules as they stand today to enable this feature (and they must be applied in this order to work):

to force packets destined for Linux to be routed instead of bridged:

- ebtables -t broute -A BROUTING -p ipv4 -i eth1 --ip-dst <ip-of-inband-mgmt-port> -j redirect --redirect-target DROP
- ebtables -t broute -A BROUTING -p ipv6 -i eth1 --ip6-dst <ip6-of-inband-mgmt-port> -j redirect --redirect-target DROP
- ebtables -t broute -A BROUTING -p ipv4 -i eth1 --ip-dst <ip-of-local-port> -j redirect --redirect-target DROP
- ebtables -t broute -A BROUTING -p ipv6 -i eth1 --ip6-dst <ip6-of-local-port> -j redirect --redirect-target DROP

to force ARPing from only one interface:

- ebtables -t nat -A PREROUTING -p arp --arp-op Request --arp-ip-dst <ip-of-inband-mgmt-port> -j arpreply --arpreply-mac <mac-of-br1> --arpreply-target DROP
- ebtables -t nat -A PREROUTING -p arp --arp-op Request --arp-ip-dst <ip-of-local-port> -j arpreply --arpreply-mac <mac-of-br1> --arpreply-target DROP

to force incoming ARP replies to be accepted instead of transformed by the later MAC-NAT rule:

- ebtables -t nat -A PREROUTING -p arp --arp-op Reply --arp-mac-dst <mac-of-br1> -j ACCEPT

to force incoming packets actually destined for Linux to be routed to Linux instead of the dataplane since the dataplane can also generate ARP requests:

- ebtables -t nat -A PREROUTING -d <mac-of-tap1+1> -i eth1 -p ipv4 --ip-dst <ip-of-local-port> -j dnat --to-dst <mac-of-br1> --dnat-target ACCEPT

to force all other frames to the dataplane:

- ebtables -t nat -A PREROUTING -d <mac-of-br1> -i eth1 -j dnat --to-dst <mac-of-tap1+1> --dnat-target ACCEPT

to force all outgoing frames from the dataplane to look like they came from br1:

- ebtables -t nat -A POSTROUTING -s <mac-of-tap1+1> -o eth1 -j snat --to-src <mac-of-br1> --snat-target ACCEPT

Layer2 IKE is not negotiating SA entries when policy is pushed from ETM [ETEP-1213]

First attempt at negotiating keys with IKE may fail and thus, could take 1-2 minutes to establish a proper connection.

Workaround:

Wait for up-to 2 minutes on initial start

Policy Entry [ETEP-995]

When entering a Layer 2 Point to Point policy the group ID becomes part of the SPI. The Group ID must be a value from 0 and 10. If a value other than this is entered the appliance will go into a failure state.

Workaround:

ETM:

Cancel the policy deploy and reboot the appliance to clear the failure. Correct the Group ID value and then deploy the policy.

CLI:

Reboot the appliance to clear the failure. Reenter the policy with the correct the Group ID.

Restore-filesystem command [ETEP-987].

Using the restore-filesystem command in the three conditions below will cause the command to not take effect. This is due to the fact there is nothing to restore. The situations are:

1. Command executed on a new ETEP
2. File System Reset
3. Going in/out of FIPS Mode

SFP Transceivers

The following SFP transceivers have been tested and approved for use with the ET1000A and ET1000A VSE.

Table 1 Approved SFPs

Interface	SFP
Fiber - single mode	Fiberxon FMT-3012C-LG Agilent AFCT-5701-LZ Agilent AFCT-5710-LZ
Fiber - multi mode	Fiberxon FTM-8012C-SLG Fiberxon FTM-8012C-L Agilent HFBR 5710L
Copper Avago	ABCU-5700RZ

NOTE

When using copper SFPs, auto-negotiation must be enabled on the ETEP and on the device to which it is connecting.

Related Documentation

The following documents describe how to install and configure the ETEP as part of the EncrypTight Manager system. They are available on the product media.

- *ETEP Installation Guide*

  Describes how to install the ETEP and perform initial setup. It also includes maintenance and troubleshooting information.
- *ETEP CLI User Guide*

  Describes how to configure and manage the ETEP from the command line.
- *EncrypTight Manager User Guide*

  Describes how to use the EncrypTight Manager software to configure and manage appliances, and create and deploy security policies.
- *EncrypTight Manager Installation Guide*

  Provides an overview of how to use EncrypTight Manager to configure appliances and create and deploy policies.

Contacting Black Box Technical Support

Contact our FREE technical support, 24 hours a day, 7 days a week:

- Phone: 724-746-5500
- Fax: 724-746-0746
- e-mail: info@blackbox.com
- Web site: www.blackbox.com