VETap™

The VETap™, or Virtual Ethernet Tap, is a software module that extends the "listening" capability of a lawful-intercept probe beyond the specific geographical (or virtual) location of the probe.

There are many situations, for cost, network layout, and other, where it is desirable for an interception probe to have remote intelligent surrogates that can extend the listening points of the probe beyond its immediate vicinity. This is provided by the VETap. The VETap can be used in a variety of configurations, such as placed in virtual machines in a cloud-service environment (where physical taps are precluded); placed in network devices to be monitored, such as VoIP SBCs; or even placed alone in a physical device having network interfaces, allowing that device to be a remote surrogate of the probe, or a "subprobe."

KEY FEATURES & BENEFITS

- Supports broadband (e.g., IPv4, IPv6, DHCP-based, RADIUS-based) and VoIP (SIP and RTP) intercepts.
- Does intelligent target-identifier detection under the direction of the probe
- Optionally filters out streaming traffic (e.g., Netflix, YouTube) from intercepts
- Uses TLS for end-to-end encryption with the probe, thus eliminating any cleartext exposure of intercept data, and any need for VPNs
- Uses no CPU and network resources when no intercepts are active in the probe
- Provides the means to implement probe-based solutions in virtual environments such as Amazon Web Services (AWS).
- Available as a container
- Up to 2048 VETaps can act as surrogates to one probe

Basic VETap – Probe Communications

This diagram shows the basic relationship between one or more VETaps and the probe.

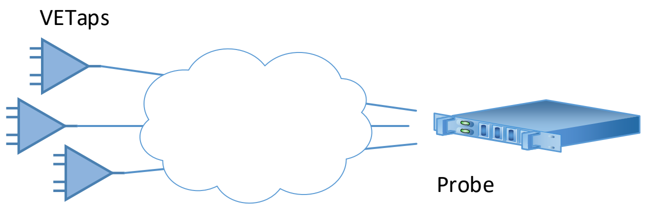

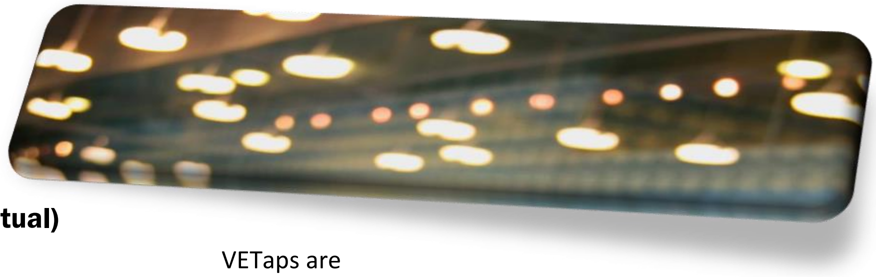

VETaps are typically remote from the probe, and, as indicated, often separated from the probe by a cloud of some type.

VETaps are provisioned (via a configuration file or environment variables) with the IP address of the probe, and thus they are dynamically discovered by the probe as they start or wake up and notify the probe of their existence. There are three types of communications:

1. Commands from the probe
2. Command responses, alarms, and other notifications from the VETap
3. Intercepted traffic from the VETap

Commands from the probe tell the VETap to enable or disable itself, to send statistics to the probe, to send its log, to enable/disable encryption, and to use a supplied set of specific criteria for interception (e.g., an IP address, a phone number, a request to send all DHCP packets to the probe, …).

In addition to responding to requests for statistics and log, the VETap periodically sends a "notify" message to the probe, letting the probe know of the existence of the VETap. If a VETap first makes itself visible to the probe while intercepts are active, the probe immediately activates that VETap and provides it with the intercept identifiers and content-filtering instructions.

Intercepted traffic is sent encapsulated in TCP for reliable delivery. Usually end-end encryption is enabled for this, although cleartext is an option available for troubleshooting.

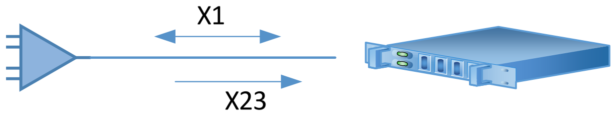

The VETap's configuration information tells it, among other things, which network interfaces of the underlying machine are to be tapped and which is used for the V interface.

VETap Usage Models

The VETap can be used in a wide variety of situations; a few common ones are described below.

Cloud. A typical use of the VETap is deploying it in virtual machines containing services needing interception capabilities. For instance, it can be used in such cloud infrastructure environments as AWS and Google Cloud. By putting the VETap in the virtual-machine image, a new VETap comes alive whenever an incremental virtual machine is instantiated.

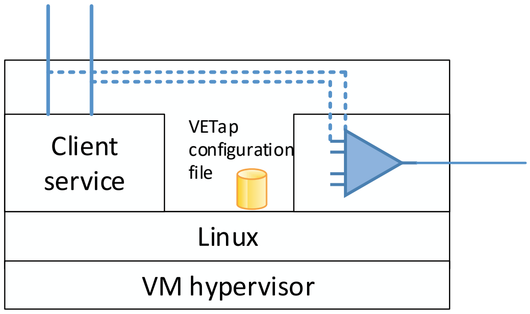

The probe isn't shown in the diagram above. Sometimes, the probe is placed in the same cloud environment. One approach that has been used in AWS is to place the probe in a separate virtual private cloud (AWS terminology), one reason being so that the probe can have a separate set of security rules. The two virtual private clouds can then be connected with an AWS peering connection.

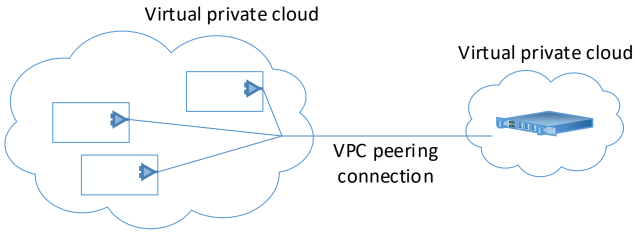

Container. Containers have become an important alternative approach to virtualization. The VETap is available as a Docker container and is also available from the AWS Elastic Container Registry. When used as a container, the VETap gets its configuration information (e.g., probe IP address, license) from environment variables.

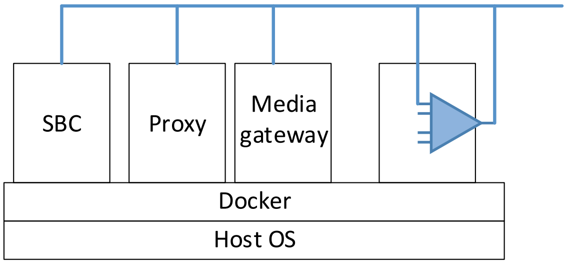

VETap as a VoIP Subprobe. The VETap software can be installed, along with Linux, in a server platform, and the configuration file can tell the VETap to tap certain of the server's network interfaces. Such a physical entity can be called a subprobe, truly a remote surrogate of the probe.

For VoIP intercept, the VETap has two alternate modes of operation, one where the VETap simply returns all SIP packets to the probe, and one where the probe informs the VETap of specific SIP identifiers (e.g., phone numbers). The diagram below depicts the second mode.

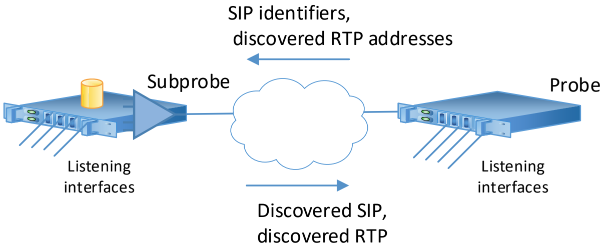

Here the VETap makes the decisions about which SIP packets to send to the probe. If the intercept provisioned in the probe specifies that the call media is also to be intercepted, the probe discovers the media attributes (IP and port addresses) and sends instructions to the VETap(s).

Another variant is shown in the following diagram.

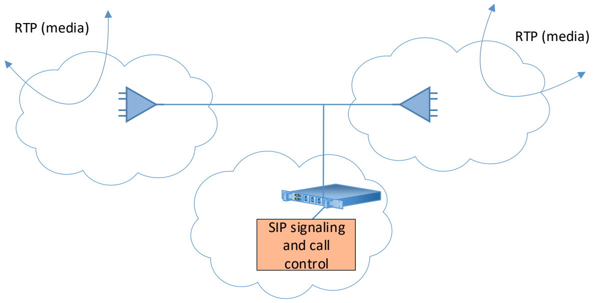

This might occur when there is a central part of the network where the SIP signaling is processed, and a probe could be placed there and listen to all the SIP traffic, but the media don't come to the central location and instead are handled in remotely dispersed gateways. Here the VETap subprobes are used at these remote sites to capture the appropriate media streams under the control of the probe.

VETap as a Data-Intercept Subprobe. For data intercepts, how the probe uses the VETaps depends on the nature of active intercepts in the probe. For instance, if one or more intercepts are provisioned as dynamic IP addresses to be assigned by a DHCP server (e.g., intercept identifier is a MAC address that will appear in a DHCP request), the probe requests the VETap(s) to send all DHCP protocol packets to the probe so that the probe can do the discovery.

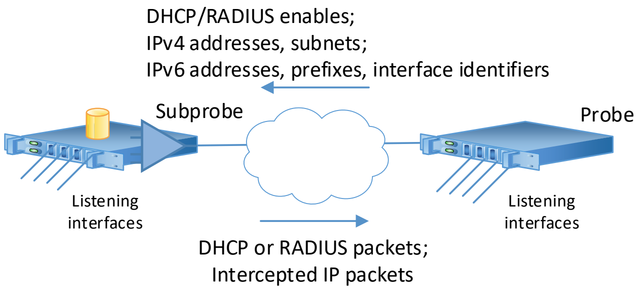

As the probe discovers dynamic IP address assignments to intercept targets, or if static IP addresses are being used for targeting, the probe provides the IP addresses (or subnets, etc.) to the VETaps.

Because data intercepts involve much higher bandwidths than VoIP intercepts, whether a cloud (e.g., public Internet) can appear between the VETaps and probe depends on the maximum traffic rates expected.

VETap as an LI API. Another use of the VETap is as a software package that is built into a manufacturer's network-equipment product to provide required LI capabilities. For instance, in the U.S., the CALEA statute requires that network-equipment manufacturers whose products provide services that are covered by the statute provide means in that equipment to facilitate lawful interception. Some manufactures satisfy this requirement by building in proprietary interfaces for LI mediation systems, but the VETap provides an alternative where the interfaces have already been implemented (by the probe).

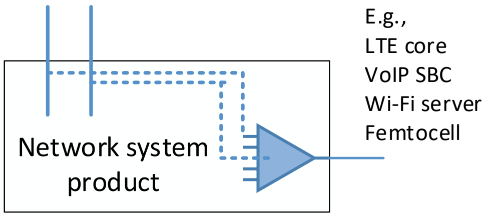

VETap Options

On whatever machine it is installed, the VETap expects configuration information, which may come from a file or, in the case of a container, from environment variables. One crucial thing in this file is the IP address and transport port of the associated probe. One can also specify a gateway address that is used as the route to reach the probe.

The VETap can run in two basic modes: (1) "raw," where all the filtering is done in the VETap process above the Linux kernel, and (2) "kernel," where the filtering is pushed into the kernel using the kernel's Berkeley packet filter. The raw mode is adequate for subprobes and the kernel mode is suggested when there is other software running on the platform. One can switch between the two modes by simply changing the configuration file.

The configuration information contains numerous other options, such as the names of the network interfaces to tap, the time interval at which the VETap should notify the probe of its presence, some TCP keep-alive parameters, and controls for troubleshooting and performance measurement.

Performance and Capability

VETap performance is a function of a wide range of things, such as the performance of the underlying system, the incoming traffic rates, the types of intercepts being performed, and the available bandwidth to the probe. In raw mode on a reasonable machine, the VETap can examine about 600K packets per second, which is roughly 4 Gbps of typical Internet traffic or 15,000 concurrent voice calls. In kernel mode, the rates are typically much higher.

Assuming the underlying Linux system is properly tuned for high TCP performance, the VETap can deliver 1 Gb/s of intercepted traffic to the probe. As mentioned earlier, when the public Internet sits between the VETap and the probe, the Internet reduces the realizable maximum transfer speed to 50-100 Mb/s.

There is a special version of the VETap that runs in a system containing 10G, 40G, and 100G ASIC-based network modules. In this system the VETap is capable of watching 200 Gbps of incoming traffic.

When used in configurations other than the dedicated subprobe, the added load put on the host system by the VETap is an important question. For both raw and kernel mode, the load is zero whenever there are no active provisioned intercepts. For raw mode, the load is primarily a function of the packets per second (pps) being examined by the VETap, and not affected significantly by the type(s) or number of active intercepts. On a current-generation 4-core, 8-CPU machine, for a typical Internet mix at 1 Gb/s actual bandwidth (about 140K pps), the load is 5% of one CPU, or <1% of the system's processing capability. For a typical VoIP media mix at 1 Gb/s actual bandwidth (about 500K pps), the load is about 6% of one CPU.

When operated in kernel mode, the added load in both scenarios above is under 1%.

Generally, the only significant load placed on the system by an actual intercept is for higher-speed data intercepts. For instance, the added overhead of the VETap intercepting and sending the traffic of a subscriber using 25 Mb/s of bandwidth is well under 0.5% of one CPU. Use of the features to filter out traffic from specified streaming services improves both CPU and network usage.

The VETap has no built-in limitations on the number of intercept identifiers it can handle.

Probe User Interface Interaction

Generally the VETaps act as logical extensions of the probe and aren't visible at the user interface. For instance, if one provisions a VoIP intercept on the probe and VETaps are present, the probe will automatically reflect this intercept to the currently discovered VETaps.

There is a page in the probe user interface where the probe user can interact with VETaps, as shown below. This page shows a list of currently discovered VETaps. By selecting one of these VETaps, statistics and a log from the VETap can be downloaded from that VETap and displayed.

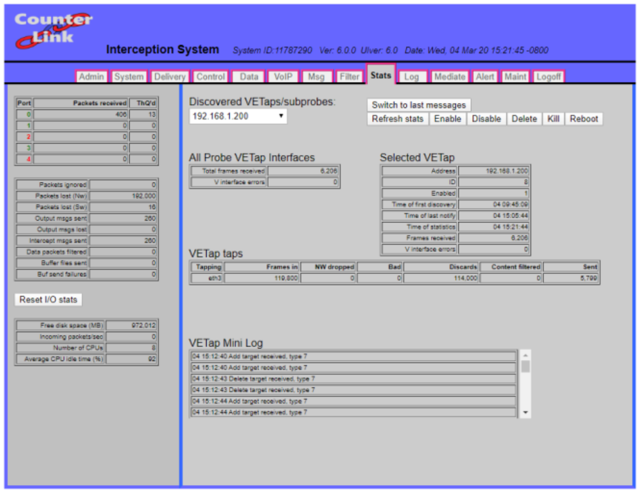

Interception System   System ID:11787290 Ver: 6.0.0 UIver: 6.0 Date: Wed, 04 Mar 20 15:21:45 -0800

Port	Packets received	ThQ'd
0	406	13
1	0	0
2	0	0
3	0	0
4	0	0

Packets ignored	0
Packets lost (Nw)	192,000
Packets lost (Sw)	16
Output msgs sent	260
Output msgs lost	0
Intercept msgs sent	260
Data packets filtered	0
Buffer files sent	0
Buf send failures	0

Free disk space (MB)	972,012
Incoming packets/sec	0
Number of CPUs	8
Average CPU idle time (%)	92

Discovered VETaps/subprobes:
192.168.1.200

All Probe VETap Interfaces

Total frames received	6,206
V interface errors	0

Selected VETap

Address	192.168.1.200
ID	8
Enabled	1
Time of first discovery	04 09:45:09
Time of last notify	04 15:05:44
Time of statistics	04 15:21:44
Frames received	6,206
V interface errors	0

VETap taps

Tapping	Frames in	NW dropped	Bad	Discards	Content filtered	Sent
eth3	119,800	0	0	114,000	0	5,799

VETap Mini Log

```
04 15:12:40 Add target received, type 7
04 15:12:40 Add target received, type 7
04 15:12:43 Delete target received, type 7
04 15:12:43 Delete target received, type 7
04 15:12:44 Add target received, type 7
04 15:12:44 Add target received, type 7
```

Another capability of the probe is POI (point of interception) mapping, where specific intercepts can be directly to named subsets of VETaps rather than to all. This has numerous uses, such as when the set of VETaps span legal jurisdictions, when the location in the network of the intercept target is known, and when duplicate identifiers could exist, such as private IP addresses.

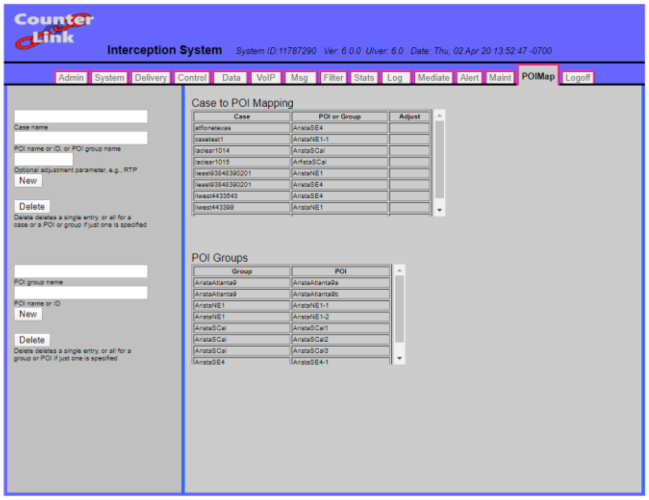

Interception System   System ID:11787290 Ver: 6.0.0 UIver: 6.0 Date: Thu, 02 Apr 20 13:52:47 -0700

Case name
POI name or ID, or POI group name
Optional adjustment parameter, e.g., RTP

Delete deletes a single entry, or all for a case or a POI or group if just one is specified

POI group name
POI name or ID

Delete deletes a single entry, or all for a group or POI if just one is specified

Case to POI Mapping

Case	POI or Group	Adjust
atflonetexas	AristaSE4	
casetest1	AristaNE1-1	
laclear1014	AristaSCal	
laclear1015	ArfistaSCal	
lieast93848390201	AristaNE1	
lieast93848390201	AristaSE4	
liwest4433543	AristaSE4	
liwest443399	AristaNE1	

POI Groups

Group	POI
AristaAtlanta9	AristaAtlanta9a
AristaAtlanta9	AristaAtlanta9b
AristaNE1	AristaNE1-1
AristaNE1	AristaNE1-2
AristaSCal	AristaSCal1
AristaSCal	AristaSCal2
AristaSCal	AristaSCal3
AristaSE4	AristaSE4-1

The probe can also produce a “satellite report,” which is a file listing the currently connected VETaps.