FortiTAP Deployment Guide

Thursday, November 19, 2015

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK

Email: techdocs@fortinet.com

Thursday, November 19, 2015

FortiTAP Deployment Guide

TABLE OF CONTENTS

Change Log

Date	Change Description
Nov 11, 2015	Initial document release. This guide is based on Rev 2 of Deploying FortiTap Solutions Guide (09/30/2014).
Nov 19, 2015	Minor corrections. Clarified the deployment 'before-and-after' diagrams. Added a link to the datasheet (which lists the loss values for each FortiTap module).

Introduction

This document provides background on using network taps versus SPAN ports, and discusses considerations for deploying FortiTaps into different network architectures.

Passive network taps provide a simple and powerful way to monitor optical networks. Fortinet's FortiTap product line provides passive network tap modules supporting 1G/10G, 40G, and 100G optical link connectivity, as well as a chassis for flexible rack deployment that allows for any combination of mixed module types.

A critical aspect of effective network security is the ability to perform forensic analysis and reporting on network traffic flows. While a great deal of information can be gathered from inline network devices, this is often at a resource cost and/or performance penalty. This is particularly true as networks start using 10G/40G/100G links, where the cost of using SPAN ports and other traffic logging/analysis techniques becomes prohibitive.

FortiTap allows near-line deployment of Fortinet or third-party security solutions, without the requirement for such solutions to be inline to the network traffic flows. Devices such as FortiGate can be deployed in a one-armed sniffer mode to observe and inspect network traffic and report on detected security events and policy violations. Devices such as FortiSandbox can observe network traffic and monitor and report on potential zero-day malware.

Note: FortiTap is compatible with any Fortinet or third-party network security device.

As a passive optical tap, with no firmware or power requirements, it can be deployed into networks with high reliability requirements.

How this Guide is Organized

This guide contains the following sections:

Background On Network Taps - overview of FortiTap, comparison with SPAN.

Deploying FortiTap Into Networks - considerations for new and existing networks, effect of the split ratios.

Additional Information - references for connecting FortiTap to other Fortinet devices.

Background On Network Taps

In many cases, vendors will denote the use of a network tap as a TAP, and apply an acronym such as Test Access Point, Traffic Analysis Port, or similar 'backronym' to somehow equalize the comparison with SPAN (Switched Port ANalyzer) ports on switching devices. The reality is that the concept of a network tap goes back to the early days of networking when vampire taps were used on coaxial-based network links such as 10Base5 (Thicknet) Ethernet, to allow additional nodes to operate on the link.

Optical network taps have a degree of analogy to those old vampire taps, in that by using optical mirror technology, they allow an additional network node to receive a copy of traffic passing between the two terminating nodes of an optical link.

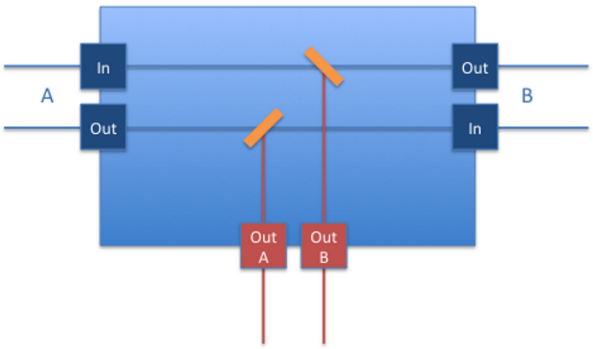

In the figure above, the individual fibers of an optical network link are shown as A and B (ports in blue). As traffic on each fiber passes through the tap, mirrors are used to optically split part of the signal to the tap monitor output (ports in red).

Network Taps Vs. SPAN Ports

The ability to configure a SPAN port is a basic function of all managed switches. Generally, the switch is configured to copy the ingress traffic, egress traffic, or both directions of a monitored port to the egress of the SPAN port. There are a significant number of differences to using a SPAN port versus a network tap:

- SPAN ports are configured on managed switches, which are active network components. Traffic copied to a span port is subject to potential performance issues associated with the switch. Subsequently, the reliability network monitoring solutions connected to SPAN ports can be affected by the switch's performance as well. Network taps are passive components and operate independently of any active networking components, thus maintaining high resiliency of any tap-attached network monitoring solutions.
- The traffic monitored by a SPAN port is actually a copy of forwarded traffic, and can be distorted by the forwarding characteristics of the switch. For example, if monitoring traffic in both directions, if the aggregate of this traffic exceeds the egress performance of the SPAN port, the switch could drop monitored packets. Using a SPAN port can skew monitored characteristics such as latency, jitter, and even basic packet ordering. A network tap provide a true copy of the network traffic to monitoring devices, since the traffic signal is split rather than copied to the monitoring ports. However, unlike a SPAN port, and tap does not integrate both directional flows into a single egress port. The traffic from both monitored egress ports must by integrated by the externally-connected monitoring device.

- SPAN ports do have the advantage of being able to copy packets between ports of mixed speed and/or media types (10G to 1G, multimode fiber to copper, etc.), as well as the ability to copy packets to SPAN ports on remote switches (remote SPAN – RSPAN). Passive network taps require identical speed/media characteristics for the output monitored traffic, and the output is always local to the monitored link.
- As SPAN ports are configured on active network switches, filters can be applied to the monitored network traffic. Using network taps, filters on monitored traffic must be applied externally by the network monitoring solution.
- For network taps, the traffic is optically split using a mirror, such that the signals egressing the tap (both to the active devices as well as the network monitoring devices) is attenuated by the tap. Therefore, signal strength and link distance is a concern when deploying a network tap. As traffic is actively copied to SPAN ports and SPAN ports are switch-based, there are no signal-attenuation issues to consider.
- SPAN ports can be configured to allow injection of packets into the SPAN port to the active devices. For example, this capability allows IPS devices in inject TCP reset (RST) packets to be received by the active network devices. Passive network taps cannot support packet injection.
- SPAN ports are configured onto network switches, and are therefore subject to potential configuration error. The deployment of network taps does not require any device configuration.

Generally, SPAN ports are switch-centric and subject to being affected by switch performance and configuration. Effects may include the potential distortion of actively copied traffic. Network taps are link-centric, and independent of the performance and configuration of the active network devices. However, they are subject to geographic architecture factors and signal attenuation. However, network taps provide for true monitoring of traffic on the links they serve.

FortiTap 124A Chassis

The FortiTap 124A chassis is a 1RU enclosure that can support up to 24 1G/10G FortiTap single-width modules, 12 40G/100G FortiTap double-width modules, or a combination of both widths (ex. 12 single-width plus 6 double-width).

FortiTap124A Chassis.

The chassis does not have power or ventilation requirements, and is intended to provide rack organization support.

FortiTap 15S/15M/17S/17M Modules

The FortiTap 15S/15M/17S/17M modules support 1G/10G links. These are single-width modules: up to 24 of them are supported within a single FortiTap124A chassis.

FortiTap 1G/10G Module (for FortiTap 15S/15M/17S/17M)

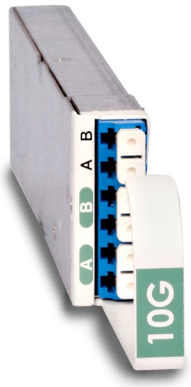

The module nomenclature of 1nX is read as:

- 1 – indicates support for 1G/10G links
  - As the module is optically passive, the same module can support either a 1G or 10G link
- n – is either a 5 or 7, and is indicative of signal attenuation:
  - 5 = 50/50 split, in which 50% of the signal is sent to the egress port, and 50% is sent to the monitor port
  - 7 = 70/30 split, in which 70% of the signal is sent to the egress port, and 30% is sent to the monitor port
- X – is either S or M, and is indicative of the media type supported:
  - S = single-mode fiber
  - M = multi-mode fiber

FortiTap 45M/47M/105M/107M

The FortiTap 45M/47M modules support 40G links, and the FortiTap 105M/107M modules support 100G links. These are double-wide modules, with up to 12 of them supported within a single FortiTap 124A chassis.

FortiTap 40G Module (for FortiTap 45M/45S/47M/47S)

FortiTap 100G Module (for 105M/105S/107M/107S)

The module nomenclature of (1)NnX is read as:

- (1)N – indicates support for 40G or 100G links:
  - 4 = 40G link
  - 10 = 100G link
- n – is either a 5 or 7, and is indicative of signal attenuation:
  - 5 = 50/50 split, in which 50% of the signal is sent to the egress port, and 50% is sent to the monitor port
  - 7 = 70/30 split, in which 70% of the signal is sent to the egress port, and 30% is sent to the monitor port
- X – is either S or M, and is indicative of the media type supported:
  - S = single-mode fiber
  - M = multi-mode fiber

Deploying FortiTap Into Networks

FortiTap modules offer the capability to reliably support large-scale network monitoring solutions while maintaining maximum network reliability and performance. However, a number of factors need to be considering in planning for their deployment within network architectures:

- Is this a deployment into a new or existing network?
- What are the monitored link speeds, and optical cable media/type requirements?
- What are the optical cable lengths, and distances to active and monitoring equipment?

The need for careful physical plant planning is required to ensure a successful FortiTap deployment. The following sub-sections describe these factors.

New Vs. Existing Network Deployments

The most optimal deployment of FortiTap modules occurs when planning a new network architecture. This allows for the best synchronization with the physical media (Layer 1) distribution plan.

The FortiTap 124A chassis is a 1RU device, supporting up to 24 single-width or 12 double-width modules, or a combination of both widths. The chassis can be positioned near a fiber distribution panel. Note that most commercially available fiber patch cords have a 1-meter minimum length, so patch cords lengths are an important consideration in positioning the chassis with a rack.

In new network deployments, using the FortiTap chassis and modules as a physical layer top of rack device allows fiber to fan out to active and monitoring network devices as required.

When deploying FortiTap modules into existing networks, it may be necessary to position the FortiTap chassis and modules closer to active network equipment, with particular consideration to active devices that aggregate multiple network links to be monitored. This strategy allows for minimal disruption when disconnecting and re-establishing links to existing active devices.

Effect Of Split Ratios

For new or established networks, the most significant consideration is to verify that the signal attenuation imposed by the installations of the FortiTap module still allows sufficient signal strength between the active network devices and to the monitoring device. This requires consideration of:

- Link speed
- Optical media (single-mode vs multi-mode)
- Overall link length between active devices, and between each active device and the monitoring device
- The amount of signal attenuation imposed by the FortiTap module

The following table summarizes the maximum distances attainable, based on fiber type and link speed (without attenuation):

Mode	Fiber Type (um)	1G	10G	40G	100G
Multi-Mode Fiber	OM1 (62.5)	220M	33M	n/a	n/a
	OM2 (50)	550M	82M	n/a	n/a
	OM3 (50)	550M	300M	100M	100M
	OM4 (50)	1000M	400M	150M	150M
Single-Mode Fiber	ITU-T G.652.D (9)	5KM	10KM	10KM	10KM

Each FortiTap module is available with a 50/50 split ratio or a 70/30 split ratio. The split ratio is a measure of how the loss is distributed between the network link and the monitor link. Each FortiTap module is rated with a signal loss value. This value depends on the module speed, fiber type, and split ratio.

On single-mode fiber, where maximum distances are measured in kilometers, the signal loss is reasonably sustainable, since the resulting link distances are still normally above a kilometer. However, using optical taps on multi-mode fiber link can have significant consequences on the resulting link distances.

Calculating the effect on overall link distance depends on a number of factors:

- The power loss budget, which is the difference between the output of the optical transmitter, and the sensitivity of the receiver. This is usually measured in decibels (db)
- The propagation or cable attenuation loss, based on the fiber type, and is usually measured in decibels per 1000 meters (db/km)
- Connector losses, based on the number of times a link passes through a physical connector between two fiber lengths, and is measured decibels (db) on a per connectorbasis
- The loss imposed by the FortiTap optical splitter, in decibels (db). The signal loss values for each module are listed in the FortiTap datasheet, available at the following location:

  http://www.fortinet.com/sites/default/files/productdatasheets/FortiTap-124A.pdf

These factors allow the installer to calculate the exact available link length between the active network components, and between each active network component and the monitoring device.

For detailed network design, you should calculate the distances using the factors described above. For initial planning, you can use the following rules-of-thumb to estimate the overall effect of deploying FortiTap.

A general conservative rule of thumb in planning is that a tap with a 50/50 split ratio will impact all link distances by 80%, leaving 20% of the original distance remaining. This is true between the active networking components, as well as from each active component to the monitoring device. For example, a 1G link on single-mode fiber can expect to support a 1km distance (20% of the 5km limit), but a 1G link on 62.5 (OM1) multi-mode fiber can expect to support only about 44m (20% of the 220m limit).

A similar conservative rule of thumb can be used on taps with a 70/30 split. Generally, one can expect to support up to 60% of the limit distance on the link between the active networking components. The difficulty is that much more of the signal attenuation impacts the monitoring link, so it is important to locate the monitoring device (or perhaps a signal amplifier/repeater such as a switch) very close to the tap itself.

In summary, a 50/50 split ratio tap incurs a much higher distance penalty on the active link, but it is more forgiving on the placement of monitoring equipment. A 70/30 split-ratio tap offers improved distance on the active link, but requires strict placement of the monitoring equipment relative to the tap, or the use of an amplifier or repeater on the monitoring link.

FortiTap Example

In the following example, our goal is to monitor the link between a FortiSwitch and a FortiGate, using a FortiSandbox in sniffer mode:

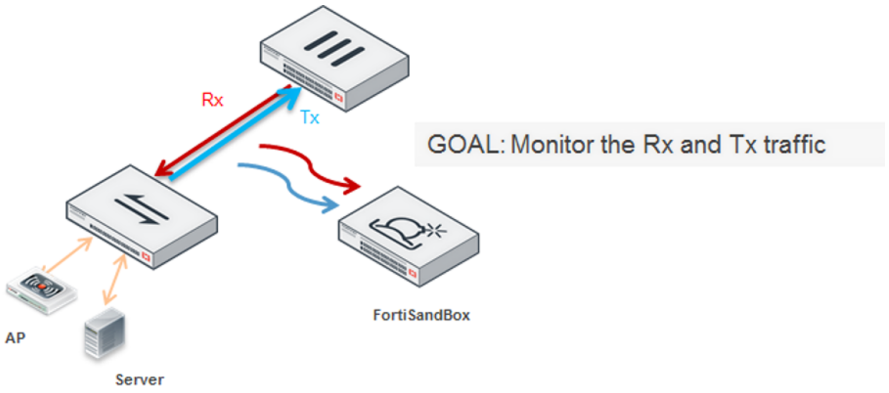

Instead of one connection between FortiSwitch and FortiGate, the FortiSwitch and FortiGate are each connected to a module on the FortiTap. The module transparently preserves the connection between FortiSwitch and FortiGate. In addition, the monitor port on the module transmits a copy of the Rx and Tx traffic, using both of its fibers to transmit.

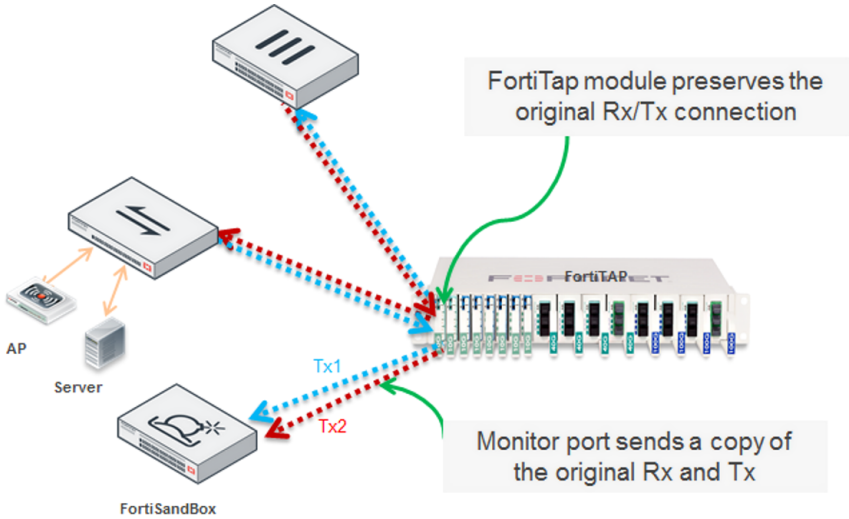

If you need to inspect traffic from both directions of the original link, Tx1 and Tx2 must connect to separate ports on FortiSandbox. When using a 40G or 100G module, this connection requires a Y-cable, which is described in the following section.

Using Y cables

The monitor port is transmitting on both fibers to the network security device. In some scenarios (where the device needs to inspect traffic from both of the monitor port outputs), you must connect the monitor port to the Rx connectors on two ports of the network security device.

The FortiTap 1G/10G modules use LC connectors. The LC connector can be manually split in half, and each half is connected to the Rx side of the ports on the network security device.

The 40G and 100G FortiTap modules use MPO connectors. You must use a Y-cable as follows, to split the monitor port output into separate A & B outputs:

1. Connect one end of the Y-cable into the monitor port of the module.
2. Connect the two "Y" ends of the cable into two ports on the network security device.

The following diagram illustrates the use of the Y-cable:

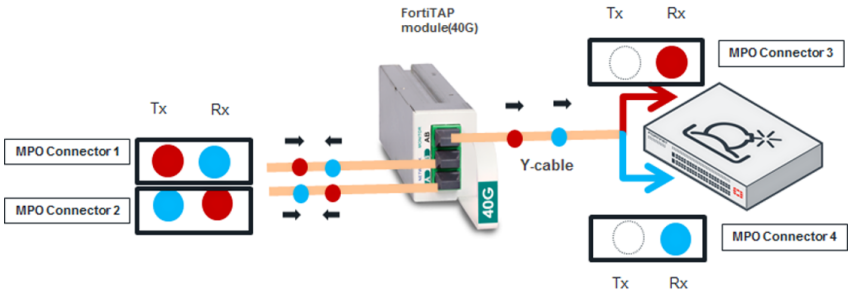

You will need one of the following Y-cables, available from Fortinet : FTP-40YM (40G) and FTP-100YM (100G)

Additional Information

FortiTap is compatible with any network security device. Depending on the capabilities of the device, FortiTap can be deployed in a simulated inline mode or in sniffer-mode.

Refer to the following links for information about FortiLink deployment with other Fortinet products.

FortiGate One-armed Sniffer

http://help.fortinet.com/fos50hlp/52data/index.htm#FortiOS/fortigate-system-administration-52/Interfaces/interfaces.htm#One-arme

Also available as a pdf file:

http://docs.fortinet.com/d/fortigate-system-administration-52

FortiSandbox deployment options

http://help.fortinet.com/fsandbox/olh/2-1-1/index.htm#FortiSandbox-211-Admin/200_Deployment/200_Deployment.htm

Also available as a pdf file:

http://docs.fortinet.com/d/fortisandbox-2.1.1-administration-guide