

SecurActive

SPV - User Guide Documentation

Release 2.9

by the Securactive Documentation Team

July 02, 2012

CONTENTS

1	Release notes	1
1.1	What's new in 2.5	1
1.2	What's new in 2.6	2
1.3	What's new in 2.7	3
1.4	What's new in 2.8	3
1.5	What's new in 2.9	4
2	Main terms and concepts	5
2.1	General Conventions	5
2.2	Zones & Fallbacks	5
2.3	Application concept	6
2.4	IP Merging	7
2.5	Concept of Conversation	8
2.6	Source / Destination Matrix	10
2.7	Data Aggregation	11
3	Deployment	13
3.1	How to integrate Performance Vision in your network?	13
3.2	How to capture traffic?	14
3.3	Supported Protocols	15
3.4	Port-mirroring and duplicated packets	17
3.5	Distributed Architecture	19
3.6	Virtual Performance Vision	22
4	Configuration	25
4.1	Hardware	25
4.2	License and upgrade installation	25
4.3	System	26
4.4	SPV Functional Configuration	29
5	Interpreting the results	43
5.1	Business Critical Application Dashboard	43
5.2	Business Critical Networks Dashboard	44
5.3	VoIP Module	45
5.4	Application dashboards	49
5.5	SPV Comparison tables	53
5.6	TCP Errors / Events	55
5.7	Packet level analysis	56
5.8	Interpretation Guidelines	59
6	Frequently Asked Questions	77
6.1	Firefox freezes randomly on some pages	77
6.2	Aggregate level changes when browsing from tables to charts	77
6.3	How can SRT be greater than DTT ?	77

6.4	How can we have 0 packets and no traffic at all on a conversation?	78
6.5	What is this timeout column (in Analysis/TCP Error)?	78
6.6	Why are some DNS request names missing?	78
6.7	Some TCP conversations are reported twice, what's wrong?	78
6.8	Pcap files generated by tcpdump are (mostly) empty	78
6.9	How to do complex searches on domain names?	78
6.10	What about Open Source?	79
6.11	Standard TCP Session	79
7	Known issues	81
7.1	Configuration	81
7.2	Interface	81
7.3	Various	81
7.4	Sniffer	81
7.5	Upgrading	81
7.6	Metrics	82
8	Glossary	83
9	Appendix	87
9.1	Virtual Appliance Step-by-Step	87
Index		113

RELEASE NOTES

1.1 What's new in 2.5

1.1.1 Installation notes

- Service Pack update must be installed before migrating from 2.x to 2.5. If the Service Pack is not installed, the 2.5 upgrade will not start.
- Migration must be done from a 2.x version. If you currently have 1.x version, please update first to version 2.0 or 2.3. Then, install the Service Pack, then install the 2.5 update.

1.1.2 New Features

- Autopcap for Business Critical Applications: available in Network conversation, DNS and VoIP depending on configuration. It works for both local and distributed environments.
- New Metric: DTT Client added to the several screens where the DTT Server was already present.
- New Protocols: LLMNR (Link Local Multicast Name Resolution), mDNS (Multicast DNS), NDNS (Net-BIOS Name Service / WINS).
- Distributed poller management.

1.1.3 Changes

Network sniffing

- Automatically detects and listens again to network interfaces that come back up after a downtime period.
- At startup, automatically adjust and fine-tune deduplication parameters for the best balance between processing power required and deduplication efficiency.

Reporting

- User / Password / TLS security support.
- User can customize "From" field when sending a report.
- Reports stored as Pdf files on the probe and available through ftp.

GUI

- For Business Critical Networks, the Retransmission Rate threshold can now be < 1%.
- Configuration area reorganized to be clearer.
- In the Configuration area, deletion buttons have been made more intuitive.
- Animation when running a request (to avoid overloading the probe by launching several times the same request).
- The timeframe selection in the “Watch last” filter is now more intuitive.
- When a filter is set to some value, it will be highlighted to be more visible.
- In Non IP traffic screen, data can be filtered by MAC address.
- Bookmarked pages now have their own specific title instead of a generic name.
- In DNS screens the filter on request types are now sorted alphabetically.
- New Screens
 - DNS Performance Graph, with DNS response times and number of packets over time.
 - TOP DNS Servers: DNS traffic and average response time sorted by servers.
 - TOP DNS Clients: DNS traffic and average response time sorted by clients.
 - DNS Overview
- New filters
 - Synthesis per DNS request types and DNS responses codes.

Pulsar

- “vpn” command has been renamed as “support”.

1.1.4 Major bug fixes

- Display of some charts could fail in some cases (long zone names added to long application names)
- Configuration was not correctly flushed in some cases.
- It was possible to define two applications on the same ports for the same IP or subnet which was leading to approximate metrics for these applications.
- Oracle parser could stop working in some cases.
- Potential deadlock under intensive usage with the implication of several different parsers at once.
- Fix an issue with Flash player and Internet Explorer that forbids drill-down into graphics.

1.2 What’s new in 2.6

1.2.1 New Features

- [GUI] User manual is now accessible from the GUI.
- [GUI] Advanced filters on client/server pages.
- [GUI] IP/subnet filter in “matrix” page.
- [GUI] Improved time frame selection with “last five used” history.
- [Pulsar] Pulsar now displays license information on the `poller` command.

1.2.2 Changes

- [GUI] “Top” screen reorganisation. We now have Tops for clients, servers, applications and ports.
- [GUI] ICMP messages regarding different connections are no more merged.

1.2.3 Major bug fixes

- [Metrics] TCP keepalives do not interrupt a data-flow any more.
- [Pulsar] Fix pulsar `process` command.
- [GUI] Fix filters on unilateral flows or retransmission.
- [Reports] Fix missing columns in some reports.

1.3 What’s new in 2.7

1.3.1 New Features

- [Config.] POSIX regular expressions are available in web patterns.
- [Reports] Can now reorder pages in a report.
- [GUI] DNS resolution requests can now be done and undone with a button (column by column) and no longer through field mouse-over.

1.3.2 Changes

- [GUI] Replace in/out by srv/clt in all pages.
- [Metrics] Deduplication is now performed independently for every interfaces/vlans if these are not aggregated.
- [Config.] Search and zone edition is now faster.

1.3.3 Major bug fixes

- [Metrics] SIP connection were not properly tracked in some cases.
- [Pulsar] Fix pulsar `analyzer ifaces` and `help` commands.
- [GUI] Fix empty unfolded line bug in grouping tables.
- [System] Restart processes when they consumes too much memory.

1.4 What’s new in 2.8

1.4.1 New Features

- [Alerts] Business Critical Applications metrics are available through SNMP. The values can be queried through SNMP (Performance Vision MIB).
- [GUI] Find the company vendor name behind a MAC address for non IP traffic.
- [Metrics] Added a new metric: “0-Window event” in TCP Events.
- [GUI] JavaScript performance improvements

1.4.2 Changes

- [PCAP] AutoPcap files are now kept for 72 hours instead of 48 hours.
- [Export] All data views can now be exported directly as a PDF page (new “Export as PDF” icon).
- [GUI] Updated TCP conversation workflow for an improved usability.

1.5 What’s new in 2.9

1.5.1 New Features

- [Alerts] Business Critical Networks metrics are available through SNMP. The values can be queried through SNMP (Performance Vision MIB).
- [Metrics] Implementation of a new heuristic to find out clients from servers without ‘SYN’ packets.
- [Metrics] Support for HTTP chunked transfer encoding.

1.5.2 Changes

- [Reports] Queried time interval in reports has been simplified.
- [Reports] Email recipients are now optional as reports are now also stored on the probe and available through ftp.
- [Reports] Reports edition now displays time intervals of each individual pages.
- [PCAP] The former limitation on storage size of manual PCAP files (20 GB) has been removed. User can now freely manage size of captures depending on available storage capacity.
- [GUI] Time selection improvement.
- [GUI] In “Monitoring” information displayed by “top” screens has been harmonized.
- [Metrics] DTT will timeout after 1 second with no data transfer. If no more data is received during this period, we considered that last packet received was the one to take into account for the DTT.

1.5.3 Major bug Fixes

- [Metrics] Retransmission rate is now computed regardless of empty packets.
- [Metrics] The de-duplication process is no more fooled by varying ethernet padding.
- [GUI] There were occasionally some empty lines in grouping tables.
- [Reports] Scheduling of report dates when set across two days (ex: from 23:00 to 01:00).
- [Reports] For reports, some client email applications were not displaying the PDF attached file.
- [PCAP] better autopcap performances when lots of files are generated.

MAIN TERMS AND CONCEPTS

2.1 General Conventions

2.1.1 Byte metric unit

All byte metric values are given in Byte as KiB, MiB, GiB, etc as recommended by the INTERNATIONAL ELECTROTECHNICAL COMMISSION (IEC) in 2000 when using power of 2^{10} multiple. This means that MiB and KiB mean that the values are in Binary and equal 1024 raised to the power of 2 and 1024 raised to the power of 1, respectively. This notation was designed to distinguish 103 bytes (referred as KB) and 1024 bytes (referred as KiB).

In other words, you would say:

- in decimal notation: 1000 k (kilo) and 10002 M (mega)
- in : 1024 Ki (kibi) and 10242 Mi (mebi)

For more information about binary prefix please refer to [Wikipedia page](http://en.wikipedia.org/wiki/Binary_prefix) (http://en.wikipedia.org/wiki/Binary_prefix).

2.2 Zones & Fallbacks

2.2.1 Principles

A *zone* is a **virtual container** in which groups of IP subnets can be kept and organized. Zones are used to map the network and to present data in accordance with the context. A zone contained inside another zone is called a *child zone* of that zone. Together, the Zones form a **hierarchy**, or a **tree structure**.

The following zones are created by **default**:

- All contains all possible IP addresses.
- Private contains:
 - **RFC 1918** (<http://tools.ietf.org/html/rfc1918.html>) (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16).
 - **RFC 4193** (<http://tools.ietf.org/html/rfc4193.html>) (FC00::/7).
 - **RFC 4291** (<http://tools.ietf.org/html/rfc4291.html>) (224.0.0.0, 239.255.255.255, FF00::/8).
 - **RFC 3927** (<http://tools.ietf.org/html/rfc3927.html>) (169.254.0.0/16). This RFC describes for IPv4 a standard method of automatically configuring network interface addresses. In IPv6, link-local addresses are required and are automatically chosen with the FE80::/10 prefix (**RFC 4291** (<http://tools.ietf.org/html/rfc4291.html>)).
- Internet is the *Fallback* of All, which means all the IP addresses which are not part of Private or of any other zones defined by the administrator.

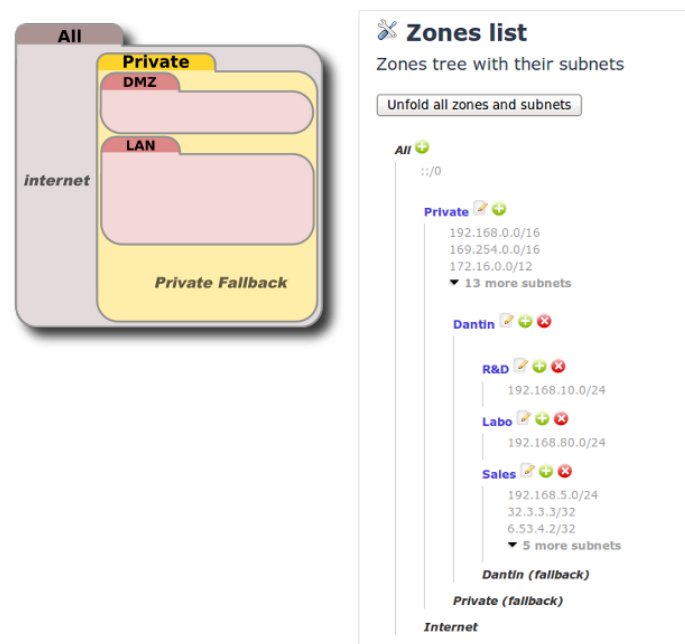


Figure 2.1: Zone set schema and zone tree displayed in SPV configuration

2.2.2 Zone fallback

A *zone fallback* or *fallback* only is the set of IP addresses which belong to a *zone* but to none of its children zones. They are automatically created for zones that contain child zones. Fallbacks are thus not configurable. At first, we didn't have the *DMZ* and *LAN* zones (see illustration). The IP addresses 192.168.0.42 and 172.16.30.45 were thus put into the *Private* zone, since they both match the definition of this zone. Then, we created both the *DMZ* and *LAN* zones. Now, the classification of those IP addresses is as follow:

- the IP address 192.168.0.42 is both part of the *Private* zone (since it is defined by the subnet 192.168.0.0/16) and the *LAN* zone (which is defined by 192.168.0.0/24). Since the latter one has a more accurate definition to store the IP, it falls into the *LAN* zone.
- the IP address 172.16.30.45, however, is part of the *Private* zone, but none of its children. Thus, to distinguish from the previous case, it is stored in a special zone, called *Private (fallback)*, which means the IP is part of *Private*, but not more.

2.3 Application concept

The main objective of *application* is to easily categorize network usage. Through this concept, which is a key notion of Performance Vision, the administrator can group similar network usages into categories that will make sense for his network context. Additionally, by configuring Applications, reports on network traffic are made clearer and are readable by any user regardless of their understanding of the underlying infrastructure (IP addresses and *subnet*, or ports used by each application).

An *application* is a set of network services which together correspond to a business application. For example, an application named *ERP* could be configured to match network traffic on port TCP/80 on a server *Zone* containing the specific server 192.168.20.4/32.

2.3.1 Application definition

An *application* can be defined using one or several of these elements:

- Application Port Range*: single port or on a defined protocol (UDP or TCP).

- *Application Signature*: designates a pattern contained in the payload of a packet which is used to recognize an application. There are two types of Applicative Signature:
 - *Signature Web application*: pattern matching on urls from HTTP requests.
 - *Signature Dynamic port*: connection tracking on supported protocol (such as FTP, Bittorrent...).
- Server Zone: zone in which the servers are located (see *Types of conversations* (page 8) for details on client/server identification).
- Client Zone: zone in which the clients of the application are located.

All packets matching the specified criteria will be identified as the particular application:

Any Port Range **OR** Applicative Signature **AND** Any Server zone **AND** Any Client zone

For more information about the configuration of applications, refer to the *Configuration* (page 25) section.

2.3.2 Examples

An application which is run on a server which IP is 192.168.1.4 with *MSSQL* will be defined as follows:

- Port Range: 1433/TCP.
- Server zone: a *MSSQL Server* zone which contains the subnet: 192.168.1.4/32.

An HTTP application running on a server along with several other applications will be defined as follows:

- Applicative signature: a Web application on *intranet.securactive.lan*.

2.4 IP Merging

In order to maximize usage of the available disk space some information are removed to allow better aggregation. This is the case for IP data of foreign host on aggregation level 3 and 4.

2.4.1 Principle

Upon data consolidation of third aggregation level, all IP tagged on the *Internet* zone will be removed in favor of a *merged* identifier. In consequence, these IP will appear as merged in all tables where IP values are displayed if the IP was belonging to Internet Zone and your observation period is such that the third or the fourth aggregation level is used. This will happen with big observation period (> 8 hours) and also on old data (> 1 week old).

2.4.2 Example

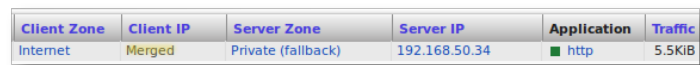
Let's say a user has accessed to the *Internet* zone with the same application, for example a web browser using HTTP on port 80 to access to different web sites for a period of time. Originally, you will see for that period :

Client Zone	Client IP	Server Zone	Server IP	Application	Traffic
Internet	86.71.197.86	Private (fallback)	192.168.50.34	■ http	535Bytes

Figure 2.2: TCP conversation before degradation

Once data has been aggregated, if you query the same period of back in time, you will have :

Merged, for the *Client IP* means that the two conversations to the different Internet clients have been merged into one single entry. This is only done when the Zone is *Internet* and matches the same server / application couple. So, you still know that this server was accessed from the *Internet* zone with the `http` application on the port 80.



Client Zone	Client IP	Server Zone	Server IP	Application	Traffic
Internet	Merged	Private (fallback)	192.168.50.34	■ http	5.5KiB

Figure 2.3: TCP conversation after degradation

2.5 Concept of Conversation

2.5.1 Objective & Definition

The objective of a *conversation* is to group a set of data exchanges between two hosts for a single *application* into one basic entity to be able to report on network traffic in a more user friendly way.

A *flow* is a group of data exchanges between two hosts for one *application* over the *aggregation period*. A *conversation* is a group of flows over the observation period. The observation period is defined by a starting time and an ending time provided by the user. A conversation is defined by the following criteria:

- The *device identifier* that received the packets
- The VLAN tag that might be present in the packets
- Source or client IP address (please refer to the chapter *Types of conversations* (page 8)).
- Destination or server IP address
- Application (please refer to the chapter *Application concept* (page 6))

2.5.2 Types of conversations

Performance Vision offers two ways to analyse network *conversation*. From an user's perspective, network conversations can be seen in two different ways, which correspond to two different needs: Client/Server or Source/Destination. This chapter explains how those views differ, which kind of information they provide, and how they can be used.

Source / destination

In a source/destination *conversation*, all flows between two hosts will be classified following the concepts of source and destination. This means that the flows will group data exchanges from a source IP address to a destination IP address regardless of their function of client or server.

For instance, a traffic from A to B for an application will be broken down in two conversations: a conversation from A to B and a conversation from B to A.

Src/Dst conversations correspond to a view of network flows for traffic analysis. When analysing data for traffic analysis purposes, an administrator wants to view flows without considering the role of each host, that is to say, **disregarding if the host is a client or a server**.

For example, traffic from A to B takes into account all traffic coming from a host in A to a host in B, whichever the role they played (client or server). The above graphs take into account the communications from A to B, only in one direction.

Client / server

In a client/server *conversation*, all flows between two hosts will be classified following the concepts of client and server. This means that the flows will group data exchanges to (and from) a client IP address from (and to) a server IP address.

For instance, a traffic from A to B for an application (provided both A and B can be a server for a single application) will be broken down in two conversations: a conversation for client A & server B (with

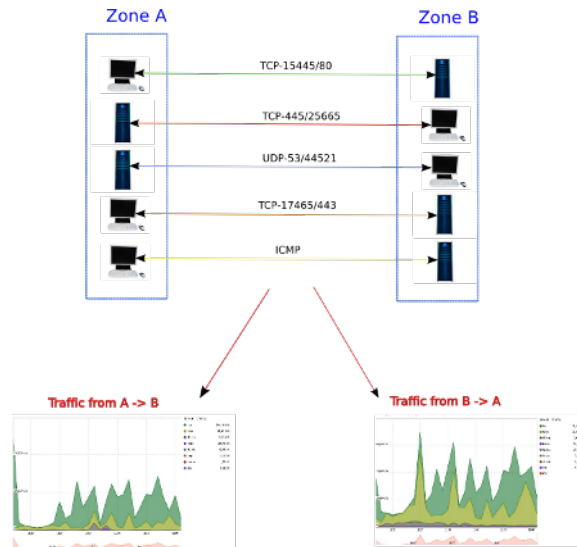


Figure 2.4: Source/Destination treatment

traffic from A to B and from B to A) and a conversation from client B to server A (with traffic from A to B and from B to A)

ClT/Srv corresponds to a view of network flows for performance analysis. When analysing data for performance analysis purposes, an administrator wants to view flows **in function of the role of each host, client or server**. Indeed, the role of a host has an impact on the metrics displayed and the clients & servers cannot be mixed.

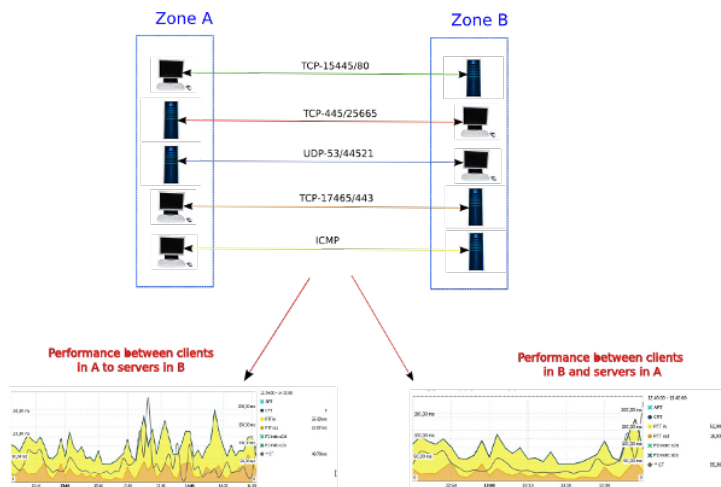


Figure 2.5: Client/Server treatment

For example, clt/srv graphs, as the ones shown above, will be generated taking into account the communications:

- from clients in A to servers in B
- from servers in B to clients in A

In short, the traffic displayed in client/server conversations will take into consideration the data transfer in both directions.

Note: Notice that the appliance can only distinguish reliably clients from servers when IP protocol in use is TCP, when the connection establishment was successfully received by the probe, and when the connection states is sufficiently active not to be timeouted. In all other cases the probe assumes that the lower port is used on the server's side.

Where are both being used?

Src/Dst will be used for all views of oriented traffic i.e. where the reports need to show the amount of data from one zone to another zone. Hereunder (in the first and second lines of the table) you can see that the data exchange between the two hosts has been split up in two conversations from A to B and from B to A.

Source IP	Destination Zone	Destination IP	Application	Traffic	Payload	Packets
192.168.80.22	VLAN_Sales (fallback)	192.168.20.208	Smtp-ssl	3.5MiB	3.0MiB	8389
192.168.20.208	VLAN_Labo (fallback)	192.168.80.22	Smtp-ssl	1.2MiB	799.8KiB	7907
192.168.20.237	VLAN_R&D	192.168.10.9	ssh	1.1MiB	771.3KiB	5294
192.168.80.22	VLAN_Sales (fallback)	192.168.20.217	Smtp-ssl	864.9KiB	824.8KiB	752
192.168.80.6	VLAN_R&D	192.168.10.6	Web Intranet Sec...	353.6KiB	521.5KiB	497
192.168.10.9	VLAN_Sales (fallback)	192.168.20.237	ssh	372.1KiB	41.0KiB	5128
204.14.234.36	VLAN_Sales (fallback)	192.168.20.213	Salesforce	352.7KiB	337.0KiB	298
192.168.20.217	VLAN_Labo (fallback)	192.168.80.22	Smtp-ssl	243.0KiB	216.7KiB	498

Figure 2.6: Source/Destination conversations

Client Zone	Client IP	Server Zone	Server IP	Application	Traffic	Packets	Handshake	Transactions
VLAN_R&D	192.168.10.5	Internet	128.237.157.136	ircu	1.7KiB	18	0	1
VLAN_Sales (fallback)	192.168.20.217	Internet	174.36.30.4	http	3.6KiB	22	4	7
VLAN_R&D	192.168.10.10	Internet	209.85.137.125	NC tcp	2.6KiB	34	4	1
VLAN_R&D	192.168.10.8	Internet	208.71.169.36	ircu	1.6KiB	17	0	2
VLAN_Sales (fallback)	192.168.20.202	Internet	91.121.2.221	vpn	16.9KiB	184	4	42
VLAN_R&D	192.168.10.4	Mother2	88.191.105.6	Srv_Mother2	16.8KiB	180	4	33
VLAN_R&D	192.168.10.6	Internet	140.211.115.34	http	12.5KiB	91	9	1
VLAN_R&D	192.168.10.4	Internet	193.48.186.4	ssh	608.6KiB	919	4	11

Figure 2.7: Client/Server conversations

On the other hand, client/server conversations will be used for all views reporting performance. Hereunder you can see (in the first line of the table) that a client/server conversation takes into account the traffic in both directions.

In summary, you will find that:

- Client/Server appears when we are speaking about Performance;
- Source/Destination appears for Usage purpose.

2.6 Source / Destination Matrix

2.6.1 Principles of Source/Destination Matrix

The Src/Dst Matrix provides a representation of the volume of traffic exchanged from zone to zone. The result is a matrix in which every cell represents the traffic from one zone to another zone. This report provides a very synthetic view of the mapping of the traffic which are observed.

The Src/Dst Matrix will show a mapping of all flows as follows:

- *blue* cells represent the internal traffic within a single zone,
- *green* to orange cells represent the traffic from one zone to another zone,
- the *intensity* of the color represents the relative volume of traffic observed in each cell.

You can filter the flows taken into account in while forming the matrix by defining:

- the *observation period*
- the source *zone*
- the destination zone
- the *application*.

The matrix is presented as follows:

The traffic data displayed in each cell must be read as from the 'Line zone' to the 'Column zone'. You can expand a parent zone to show its children zones by clicking on the + symbol.

You can use the filters to display the traffic from one specific zone to another specific zone:




Figure 2.8: Source/Destination Matrix

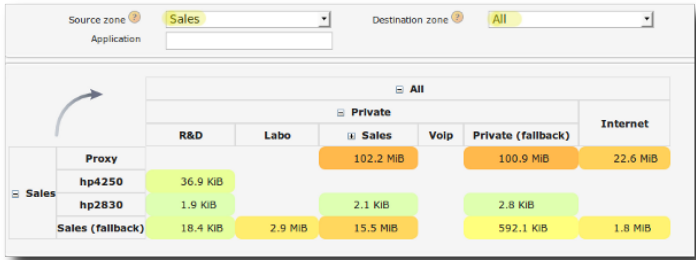


Figure 2.9: Filtered matrix

2.6.2 Top down analysis

The Src/Dst matrix can be the starting point for a fine tuned analysis of traffic: bandwidth and conversation. In each cell, there are two buttons:

- one to display the bandwidth graph from zone A to zone B
- one to display the conversations from zone A to zone B.




Figure 2.10: Cell detailed view

The first link will open the conversation table and will display all the traffic between the two zones, whereas the second one will display a bandwidth chart from the source zone on the left and the destination zone on the top.

2.7 Data Aggregation

2.7.1 Rationale

By nature, the operations of statistical analysis performed require the storage of large amounts of data. Furthermore, that data must be stored over extended lengths of time so as to expose overall trends. In order to minimize storage space while still making it possible to reveal trends over weeks or months, Performance Vision automatically summarizes the collected data over standard lengths of time. The process of creating these summaries is called aggregation.

2.7.2 Process

Aggregation occurs automatically. Whenever your probe displays a chart or a table, this is based on already aggregated data. In order to display this aggregated data, Performance Vision first decides on an aggregation granularity depending on the length of the time period you requested and how far back into the past it goes.

Aggregation granularity	Storage duration	Request length for tables	Request length for graphs
2 minutes	48 hours	60 minutes	120 minutes
15 minutes	7 days	8 hours	16 hours
2 hours	2 months	2 days	5,25 days
1 day	1 year	359 days	359 days

For example, with graphs, if you want a data granularity of two minutes, you can request a period length up to 120 minutes anywhere during the last two days.

Another example: with tables, if you want a data granularity of two hours, you can request a period length up to two days anywhere during the last two months.

Note that because the larger aggregate levels summarize more data at once, they take up less disk space, and can be kept in storage much longer without filling out the hard drive. This strikes a good balance between data granularity and duration of retention: performance data for the last two days is available with the best granularity, and long-lasting global trends can be exposed from as far back as one year (albeit with less detail), all from the same interface.

Aggregated data is computed, in a nutshell, by identifying network conversations where the same server and the same client talked using the same application, and grouping them together. The metrics for each such group are summed up in accordance with their mathematical nature (for instance, packet counts are added and response times are averaged per packet), so only one line of data is retained for each conversations group. This line still contains a relevant summary of your network and application performance, but it's storage takes up a lot less disk space.

Example: A user checks out a Web page once at 16:38...

Info

Query begin: 2010-07-30 16:38:00-02:00
Query end: 2010-07-30 16:40:00-02:00
Aggregate Level: 120s
Number of collected results: 1

Sync	Begin Time	End Time	Client Zone	Client IP	Server Zone	Server IP	Application	Traffic	Packets	Handshake	Transactions	EURT
	2010-07-30 16:38:47	2010-07-30 16:39:11	Private	192.168.10.2	Internet	88.191.122.7	http	139.0KB	240	12	32	100ms

Figure 2.11: Flow example at 16:38 to 16:40

... and once at 16:41.

Info

Query begin: 2010-07-30 16:40:00-02:00

Query end: 2010-07-30 16:42:00-02:00

Aggregate Level: 120s

Number of collected results: 1

Sync	Begin Time	End Time	Client Zone	Client IP	Server Zone	Server IP	Application	Traffic	Packets	Handshake	Transactions	EURT
	2010-07-30 16:41:34	2010-07-30 16:41:56	Private	192.168.10.2	Internet	88.191.122.7	http	55.9KB	152	7	30	112ms

Figure 2.12: Flow example at 16:40 to 16:42

Then here is the aggregated line for both events if you query between 16:38 and 16:42:

Info

Query begin: 2010-07-30 16:38:00-02:00

Query end: 2010-07-30 16:42:00-02:00

Aggregate Level: 120s

Number of collected results: 1

Sync	Begin Time	End Time	Client Zone	Client IP	Server Zone	Server IP	Application	Traffic	Packets	Handshake	Transactions	EURT
	2010-07-30 16:38:47	2010-07-30 16:41:56	Private	192.168.10.2	Internet	88.191.122.7	http	194.9KB	392	19	62	106ms

Figure 2.13: Flow aggregation from 16:38 to 16:42

Observe that the traffic, and the packet, handshake and transaction counts have been added, and the EURT averaged. For example handshake is now 19 (12 + 7).

Note: Performance Vision requires a complete set of data for an aggregate level to compute its summary. This is the reason why captured network events don't appear right away on your probe: the probe first waits until the end of the minimal aggregate time of 2 minutes, computes its summary, and only then is the aggregated data for these last 2 minutes made available in the interface.

DEPLOYMENT

3.1 How to integrate Performance Vision in your network?

3.1.1 Preliminary steps

Performance Vision is dedicated to analyzing the performance of business critical applications in a corporate network. Hence the very first step before considering integrating Performance Vision in your network, is:

- identifying an up-to-date list of business critical applications (including applications directly supporting business processes, but also applications on which these may rely – e.g. DNS, Microsoft-DS etc...).
- locating the servers hosting these applications.
- defining which network devices clients are using to access these applications.

3.1.2 Positioning the probe

Performance Vision appliance will be installed as close as possible to the servers to provide the best analysis. Measurements are more accurate if the probe is located in a central location next to the server and you will get a wider view on the performance experienced by all the users connecting to this server.




Figure 3.1: SPV network positioning synoptic

3.1.3 Choosing a traffic capture method

Two main methods may be used to establish a permanent point of traffic capture: TAP or SPAN. A TAP is a network device which will be installed in-line on the network and will send a copy of the traffic on one or two listening ports of the probe. A SPAN (also commonly called port mirroring) is a feature of network switches that enables a network administrator to send a copy of a given traffic (on one or several interfaces / VLANs to a mirroring port).

The most commonly used method is the `SPAN` port (port mirroring) mainly because it enables administrators to monitor potentially any traffic going through the switch, with an existing network device. Collecting traffic through a `SPAN` port will likely not generate any additional point of failure on the network and will be regarded as a minor modification of its existing configuration. Network `TAPs` are also an option (if no `SPAN` is doable for example) but the traffic captured will be limited to the network link(s) going through the `TAP`. A connection via `TAP` induces additional costs.

If you choose to capture network traffic through a `SPAN`, you should pay a specific attention not to copy twice the same traffic to the listening interface of the probe (which would degrade the statistics provided by the probe).

3.2 How to capture traffic?

Performance Vision can rely on two mechanisms to capture network traffic: Port Mirroring (commonly called `SPAN`) & `TAP` (Terminal Access Point).

3.2.1 Port mirroring

Port mirroring, also known as `SPAN` or *roving analysis*, is a method of monitoring network traffic which forwards a copy of each incoming and/or outgoing packet from one (or several) port(s) (or `VLAN`) of a switch to another port where the analysis device is connected. Port mirroring can be managed locally or remotely. To configure the port mirroring, an administrator selects one or several ports from which all packets will be copied (source ports) and another port or ports where the copy of the packets will be sent (destination port). The administrator can include either all packets in the port mirroring or only the transmitted/received packets. In case both transmitted and received packets are included, a packet going from a 1st monitored port to another monitored port will be copied twice to the destination port. This will have an impact on the measures and performance provided by the analysis device (e.g. retransmission rates, response times, ...). Performance Vision captures and evaluates the data without any impact on the original traffic.

The port mirroring is the most commonly used solution to capture traffic, because it is inexpensive, flexible in terms of how much traffic can be captured at once and remotely configurable.

Please note that a port mirroring may have some drawbacks, such as:

- It can consume significant CPU resources while active
- There is a risk of not receiving some packets (like media errors)
- In the case of traffic congestion at the switch level, the port mirroring is likely to drop some traffic (because the `SPAN` process does not have priority).

In some cases, a better solution for long-term monitoring may be a passive `TAP` or an Ethernet repeater ("hub").

Advantages

- Low cost (this feature is embedded in most switches)
- Can be configured remotely through IP or Console port
- The only way to capture intra-switch traffic
- A good way to capture traffic on several ports at once

Drawbacks

- Not adequate for fully utilized full-duplex links (packets may be dropped)
- Filters out physical errors
- Impact on the switch's CPU
- Can alter the timing of the frame (with an impact on response time analysis)

- SPAN has a lesser priority than port to port data transfer

3.2.2 Network TAP

A network TAP (Terminal Access Point) is a hardware device which can passively capture traffic on a network. It is commonly used to monitor the network traffic between two points in the network. If the network between these two points consists of a physical cable, a network TAP may be the best way to capture traffic. The network TAP has at least three ports: a port A, a port B, and a monitor port. To place a tap between points A and B, the network cable between point A and point B is replaced with a pair of cables, one going to the TAP's A port, one going to the TAP's B port. The TAP passes all traffic between the two network points, so they are still connected to each other. The TAP also copies the traffic to its monitor port, thus enabling an analysis device to listen. Network TAPs are commonly used by monitoring and collection devices. TAPs can also be used in security applications because they are non-obtrusive, are not detectable on the network, can deal with full-duplex and non-shared networks, and will usually pass-through traffic even if the tap stops working or loses power.

Advantages

- No risk of dropped packets
- Monitoring of all packets (including hardware errors -MAC & media)
- Provides full visibility including congestion situations

Drawbacks

- The device may require two listening interfaces on the analysis device
- Costly
- No visibility on intra-switch traffic
- Not appropriate for the observation of a narrow traffic range.

3.3 Supported Protocols

The SPV sniffer can detect all Ethernet packets even if those packets have a VLAN tag in their Ethernet header. SPV also accepts both IPv4 and IPv6 protocols.

Note: Non Ethernet flows are invisible for the SPV solution.

3.3.1 Non IP Protocols

If the Ethernet protocol is not an IP protocol, it will appear in *Non IP* submenu. All those data will not appear elsewhere.




Figure 3.2: Non IP protocols menu

3.3.2 IP Protocols

Ipv4 and Ipv6 are both captured and splitted in four *Level 3/4* protocols: TCP, UDP, ICMP and OtherIP.




Figure 3.3: Level 3/4 protocol filter

Some of those data are duplicated in other specialised categories: *Web*, *VoIP*, *DNS* to display more specific metrics.

Begin Time	End Time	Request Name	Packets	Traffic	DNS rt
2011-09-07 10:38:31	2011-09-07 10:43:32	pypi.rd.securactive.lan	1 712	198.5 KiB	< 1ms
2011-09-07 10:38:31	2011-09-07 10:43:32	pypi.mimors.rd.securactive.lan	596	73.8 KiB	< 1ms
2011-09-07 10:00:50	2011-09-07 10:59:47	No request name	570	100.9 KiB	< 1ms
2011-09-07 10:00:21	2011-09-07 10:59:23	mail.google.com	120	18.3 KiB	27 ms
2011-09-07 10:00:23	2011-09-07 10:58:37	reviewsboard.rd.securactive.lan	112	14.0 KiB	3 ms
2011-09-07 10:01:25	2011-09-07 10:56:59	WORKGROUP	97	8.7 KiB	-
2011-09-07 10:01:16	2011-09-07 10:59:14	safebrowsing.clients.google.com	62	13.8 KiB	54 ms
2011-09-07 10:05:34	2011-09-07 10:56:10	proxy.securactive.lan	61	7.3 KiB	< 1ms
2011-09-07 10:04:57	2011-09-07 10:51:37	git.rd.securactive.lan	56	6.4 KiB	< 1ms
2011-09-07 10:02:26	2011-09-07 10:56:03	sdouche@babbarge_presence_tcp.local	55	5.5 KiB	-

Figure 3.4: DNS specialised view

3.3.3 Limitations

If the rate of incoming packets exceeds the rate at which the sniffer can parse the traffic for too long then some packets may be dropped by the Linux kernel. These packets won't get accounted for in the GUI.

As a realtime protocol analyzer, the sniffer is also limited in what protocols it supports and how deep it inspects packets. Here is a quick overview of the most blatant limitations:

- Ethernet parser supports Linux cooked capture extension (used when capturing on “any” interfaces) and 802.1q vlan tags. All other Ethernet extensions are ignored.
- Http parser does not support multi-line headers.
- ARP parser knows only Ethernet and IP addresses.
- DNS parser support MDNS, NBNS and LLMNR in the extend where these protocols mimic legacy DNS (with the exception that it can unscramble NetBios encoded names).
- FTP connection tracking merely look for PASSV or PORT commands in the TCP stream without much care for the actual protocol.
- TCP options are ignored.
- Postgresql parser supports only protocol version 3.0 and Mysql parser supports only protocol version 10. This should cover most of the installed base, though.
- TNS parser (for Oracle databases) was roughly reverse engineered from various sources, especially the wireshark source code. It should thus not be expected to understand all messages in all situations.

- SIP parser implements no proprietary extensions, however prevalent.
- As there are no concept of connections for UDP, UDP conversations are ended after a timeout period of 2 minutes without any packet in any direction. This might not match the underlying protocol.
- VoIP dialogs are identified by their call-id only, which imply that if the sniffer listens to various independent SIP proxys or servers then call-id collisions can not be ruled out (this choice was made because it proven useful in practice).

3.4 Port-mirroring and duplicated packets

3.4.1 Introduction

The configuration of a port-mirroring session has to respect some specific rules and standards. The main goals of a port-mirroring session are to:

- Gain insight into the highest number of flows, which are seen as strategic by the IT manager
- And ensure that all collected flows are appropriately analysed.

It is crucial to ensure that a minimum number of flows are not duplicated to the interfaces.

3.4.2 Detail

SPV solutions takes into account duplicated packets (packets may be dropped). However, this will involve a significant loss of performance. There are two main rules:

- Basic port-mirroring sessions, also called **1 to 1** port-mirroring session. This configuration does not generate duplicated packets. However, increasing the number of **1 to 1** port-mirroring sessions could produce this phenomenon.

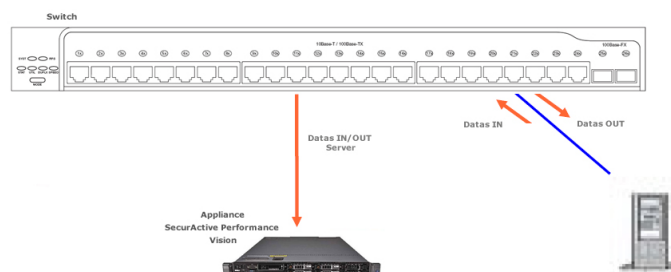


Figure 3.5: “1-to-1” port mirroring session

- Multiple port-mirroring sessions, also called **N to 1** port-mirroring session. In this specific event, the duplicated packets phenomenon can occur.




Figure 3.6: “N-to-1” port mirroring

Warning:

- According to the number of listening points, in a multi-switch mode this phenomenon can occur despite the use of a **1 to 1** port-mirroring session.
- A VLAN is a definition of a set of ports; this means that the port-mirroring session is a **N to 1** port-mirroring session.

3.4.3 Some examples of duplicated packets / non-duplicated packets

In a standard port-mirroring configuration (**N to 1**), it is highly likely that some transmitted packets to the appliance are duplicated. In the following example, configuring a port-mirroring session on both the IN traffic and the OUT traffic of the switch means that the appliance will receive twice the same traffic:




Figure 3.7: Example with duplicated packets

By only listening to the IN traffic (or only the OUT traffic) on the Ethernet ports concerned, we will ensure the flow transmission to be in a unique way for the sessions between the client and server, thus avoiding the duplication of packets:




Figure 3.8: Example without duplicated packets

Note: In the event of a **N to 1** port-mirroring session, the total bandwidth of the “source” Ethernet ports of the mirror should not exceed the maximum bandwidth of the “destination” Ethernet ports of the mirror.

3.4.4 Removal of duplicated packets

The SecurActive system checks and controls the duplicated packets phenomenon on all listening ports. It also ensures all duplicated packets are removed. However, in some cases, some duplicated packets could be mixed up with retransmitted packets.

It is therefore crucial to minimize the duplicated packet rate. In order to reach a low rate of duplicated packets, the appliance provides information on the duplicated packet rate through the Pulsar command:

This means that 5.12% of the listening traffic is duplicated.

```

Welcome to Pulsar - the SPV shell - v1.13.0

Type 'help' to display Pulsar commands.
Type 'help COMMAND' to display the command help details.

vsonde73# analyzer mirror
5.12%
vsonde73#

```

Figure 3.9: Information on the duplicated packets rate in Pulsar

3.4.5 Deduplication algorithm

The sniffer usually receive frames from multiple locations on a network, and so it can be cumbersome (if not impossible) to avoid the situation where the same frames are mirrored toward the probe. Deduplication is the process of ignoring selectively packets that are artificial duplicates due to the network infrastructure. The following chapters covers the deduplication system in order to help minimizing duplication issues.

The packet sniffer detects and drop duplicate frames before parsing their content according to the following algorithm.

First of all, frames smaller than the *Ethernet* header size are not checked against duplication.

Then, only a selected set of frame bytes are compared:

- For small frames (which size is below the size of an IP header) all bytes are taken into account,
- for bigger frames bytes after the Ethernet header (including the Vlan tag if collapsing VLans) and up to the 64th byte of the frame (or less if the frame is smaller), excepting the TOS, TTL and IP checksum fields are taken into account.

The rational behind skipping *Ethernet* header is that we want to pair two packets if only their Ethernet addresses or Vlan tag differ (one is a copy of the other, merely one switch away from it). The rational behind excluding TOS, TTL and checksum fields of the IP header is to be able to pair two packets when one is a copy of the other, only one hop away from the first one (after traversing one or several routers).

Then a packet signature is build from the remaining bytes and compared to those of previously received packets. If a packet with the same signature was previously received then the new packet is merely dropped.

The deduplication algorithm makes use of a few parameters for limiting the number of packets to check:

- The maximum delay between two potential duplicates (default: 100ms). Any duplicate frame that is received after this delay will be processed as a legitimate frame.
- The current average delay between echoed packets ; this one is recomputed by a short but frequent comprehensive search through all packets received up to the maximum allowed delay.
- How often the current average should be computed (default: every 10s).
- How far back in time we should look for duplicates relative to the current average (default: average + 1 x sigma)


These default values should fit most settings.

3.5 Distributed Architecture

3.5.1 How does the distributed infrastructure work?

Appliances hosting only the sniffer component of SPV are called “pollers”. The appliance hosting the components in charge of collecting, merging and integrating the data from the pollers into a single database is called “collector”. The collector appliance may also host one sniffer component.

The pollers listen and analyze the network traffic. The collector receives data from the pollers, integrate them in the database, and then provides an access to the data through the Web UI.



You can add a new *poller* via *Pulsar* (page 26) by using the command `poller add <IP>`. The specified IP of the *poller* must be reachable with SSH port 22.

The *Pollers Status* page in the *Configuration* menu display some status information about pollers.

3.5.2 Where is data being merged / segregated?

The data is merged (i.e. the data is integrated in the reports with no consideration for the poller, which has captured it) in:

- Business Critical Application Dashboard.
- Business Critical Network Dashboard.
- Application dashboards.
- Graphs (performance, bandwidth, matrix).
- Comparison tables (Client / Server, Network performance, Application performance).

Please note that in these reports, you can enter a filter to view the data captured by one poller only. The data is segregated (i.e. the data is kept separated depending on the poller which captured the data) in all other tables.¹ Please note that in these reports for a single conversation viewed by two pollers, you will get two lines.

3.5.3 What happens if a poller does not answer?

If a connection to a poller is broken, the collector wait for it during 10 minutes. After this time interval, the collector will flag the poller as 'missing'. After these 10 minutes, the collector stops waiting for the missing poller and restarts its activity. Data integration will be 10 minutes shifted upon missing poller response again. See example bellow:

```
min00
^^^^
poller1 ok
poller2 ok
=> data integration

min02
^^^^
poller1 ok
poller2 fail
=> wait for poller2 [min02]

min04
^^^^
poller1 ok
```

¹ This may never be developed.


```
poller2 fail
=> wait for poller2 [min02, min04]

... same, wait more and more poller2 data ...

min12
^^^^
poller1 ok
poller2 fail
=> integrate data of poller1 for "min02"
=> wait for poller2 [min04, min06, min 08, min10, min12]

min14
^^^^
poller1 ok
poller2 ok
=> integrate all data poller1 and poller2

Conclusion
^^^^^^^^^^
Data lost: poller2 [min02]
```

3.5.4 How configure a poller?

All pollers are available via SSH using the Pulsar shell, just like you access to the collector (please refer to *Pulsar* (page 26)). A poller shell allows you to configure the poller IP, hostname, etc. But some commands like `reset` or `poller` are not available.

The collector's shell allows you to show and to create or delete pollers. To do this, please use the `poller` command (`help poller` for details).

3.5.5 Limits

The distributed architecture provided by version 2.5 has some intrinsic limits:

- There is no feature for deduplication between pollers (i.e. a network flow captured by two pollers will be counted twice in reports that merge data from several pollers).²
- If there is some load balancing at the packet level (and not at the session level) and two pollers view two different parts of the traffic, the collector will not be able to rebuild this flow and no performance metric will be available in this case.³
- The positioning of each poller with regards to client and server will have some impact on some metrics (SRT, RTT Server, RTT Client, RR Server, RR Client, ...)
- The zone and application objects do not integrate the concept of poller (i.e. you cannot distinguish between two applications based on the fact that they are viewed by two different pollers).
- The maximum number of sessions handled by the collector remains unchanged (approx.. 100k concurring sessions).

3.5.6 Prerequisites

- All pollers have to be synchronized to a single *NTP*.
- All pollers and collector require an administration port connected to the network and a fixed IP address.

² This corresponds to a rare case ; this case is not handled by the non distributed implementation of Performacne Vision, nor by most competitors. The bypass option would be to use TAPs to re-aggregate both flows before it reaches the interface of the poller.

³ This is already the case in a non distributed implementation. The only new element is the fact that data will be more readable if all pollers have the same capture points.

- Connectivity between pollers and collector on port TCP/22 is required.
- Some network capacity is required to transfer the data from the pollers to the collector (current evaluation is 0.2% of the bandwidth analyzed).

3.5.7 Adequate / non-adequate implementations

Situation	Fit for version 2.5	Comments
Two data-centers (Active / passive)	Distributed may or may not be required.	Most applications will be deployed in normal conditions on DCa; if in normal conditions DCb, receives no production traffic, hence a second probe may not be required; if applications are, in normal conditions, distributed between DCa and DCb, then a distributed implementation is required.
Two data-centers (Active / Active)	Distributed is adequate.	If the traffic between servers is captured, it may be double counted ; traffic from clients to servers should not be double counted.
N data-centers through WAN.	Distributed is adequate.	Traffic between servers will be captured twice and double counted.
N Datacenters and M remote sites	Distributed may not be adequate.	The traffic going from the remote sites to the datacenters will be double counted. The cost of deploying physical units may be superior to the benefit.

3.6 Virtual Performance Vision

Note: For more details about step-by-step virtual appliance installation cf [Virtual Appliance Step-by-Step](#) (page 87).

If you are installing the virtual image of Performance Vision then you have to take into account a few additional facts.

3.6.1 How to get the image

This section is based on version 2.5.13, the filename will evolve depending on the version number.

The ZIP archive will contain the following files:

- SPV-2.5.13-r2.mf
- SPV-2.5.13-r2.ovf
- SPV-2.5.13-r2.disk1.vmdk

3.6.2 Virtual hosts settings

Performance Vision virtual appliance is designed to run in a VMWare ESX v4 or v5 environment. It can be launched with a minimum of 512MB of RAM although a larger quantity is recommended to ensure satisfactory performance rates.

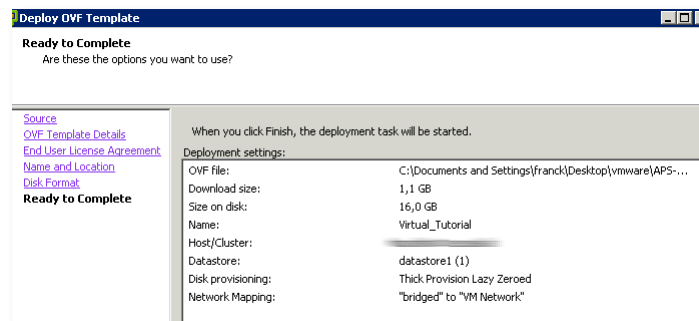
However all settings cannot be tested; in case of doubt it is recommended to fall back on these tested settings:

- RAM: 512MB, 4GB, 6GB, 8GB, 12GB or 16GB;
- CPU: 1, 4 or 8;

3.6.3 Installation

1. Connect to your Vsphere Client and then in the Virtual Machines tab, in the “File” menu, select “Deploy a new OVF template”.
2. Find and open the Performance Vision OVF file.
3. Click on “Next” twice and then accept the license agreement
4. Name the Virtual Machine appropriately (SPV appliance for example).
5. The system detects the space available on the disk for the new Virtual Machine, we recommend to allocate the following spaces:
 - Trial Virtual Appliance: 4GB RAM, 2 vCPU > 2,0 GHz
 - Virtual Poller: 8 GB, 2 vCPU > 2,0 GHz,
 - Virtual Appliance: > 16 GB, 4 vCPU > 2,4 GHz

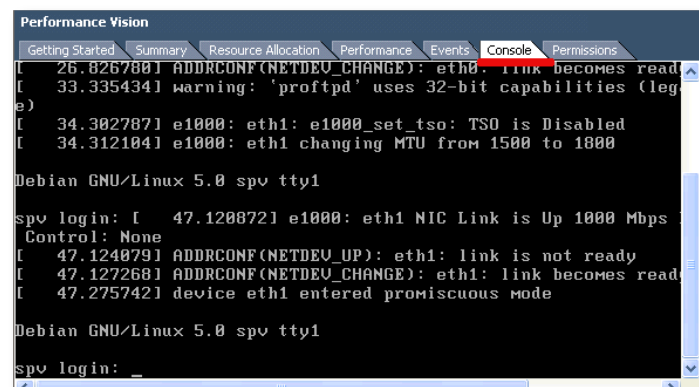
You get:



6. Click on “Finish”, the Virtual Appliance gets installed. You will get notified when the installation is complete.
7. Once the Virtual Appliance is installed, you have to start it by clicking on “Power on the Virtual Machine” or on the green triangle.

3.6.4 Access the virtual console

Display the Console tab and access the CLI interface named “Pulsar”.



The probe is launched. When the network interfaces turn into promiscuous mode, click on the Console view and then “Enter” to display the login prompt. Please note: clicking on the black screen deactivates your mouse. To reactivate it, you can use the key combination “Ctrl + Alt”. To configure the probe, please refer to the *Pulsar* (page 26) chapter. After configuration you have to reboot the virtual appliance.

3.6.5 License

Except the experimental virtual appliances for testing provided from our Web site, the virtual appliances are delivered without license key. You normally receive this key by e-mail at the product's delivery. If it is not the case, please contact our sales department: sales@securactive.net (sales@securactive.net).

To install a license package (as well as an upgrade package), proceed as usual (see *License and upgrade installation* (page 25)).

3.6.6 Capturing traffic

Virtual appliances are configured with only two network interfaces:

- eth0 for administration
- eth1 for sniffing traffic

Any additional virtual adapters you may add will be listened for traffic by the packet sniffer.

Actual packet capture depends on the virtual switch you are using.

In the realm of VMWare's bundled Virtual Switch the `promiscuous` mode (beware that name is misleading) is actually a port mirroring. Also, depending on the virtual switch configuration, if the packet sniffer sets the promiscuous bit of the eth1 virtual adapter, the mirroring mode will be activated automatically. Refer to the *Virtual Infrastructure Client manual* (<http://www.vmware.com>) for further details.

Under VMware Player you need to configure eth1 as a bridged device, and give permission to the virtual appliance to turn it into promiscuous mode.

Other virtual switches may have different/more features.

3.6.7 Data storage

Virtual appliances come with no data disk, thus everything (traffic data as well as pcaps and reports) will be written to the system disk only.

If you plan to keep a long history of data then a dedicated data disk is mandatory. To create one, attach a new drive to your VM and then run the `format_data_disk` command from `pulsar`.

Notice that:

- you will not be able to resize this data disk hereafter (the required size depends on the traffic you plan to monitor but anything below 500GB seems dubious);
- the data previously acquired will be lost;
- you are required to reboot the appliance once done.

CONFIGURATION

4.1 Hardware

The first thing to do is to plug a screen and a keyboard to the probe (for first set-up only) and then to provide electrical power. Once done, just turn power on.

For the screen, the connectivity is a standard VGA port. Two are available, one is located on the front side of the probe, the other is located on the rear side of the probe.

For the keyboard, you can plug it to any of the four USB ports. Two of them are located on the front side of the probe, the two others are located on the rear side of the probe.

By default the probes are equipped with four *Gigabit Ethernet* interfaces labeled 1 to 4. The first one is the administration port used to connect to the probe. Plug the Gb1 network interface to your network to be able to connect to the probe. The three others interfaces, 2 to 4 are dedicated to network traffic sniffing. Connect one or more of these interfaces to your network according to the network traffic you want to analyze and monitor.

4.2 License and upgrade installation

All SPV entities: virtual, poller and collector (see *Distributed Architecture* (page 19)) **needs** a specific license. The licenses are specific to a given hardware serial number (the *device id*), so that each device must be sent its own license package.

The same procedure must be performed for all the entities either for license or upgrades, please follow the steps below:

1. Connect to the *FTP* server of the probe (*user: ftp, password: S3c7r!*).
2. Upload (put) your license or upgrade file.

Wait a few minutes and it's done! Check your license or new version with the `status` or `poller` commands. For upgrades, please redo the same procedure on all the entities.

Warning: It is **STRONGLY** recommended to reboot all the probes after upgrading (use the `reboot` command in *Pulsar* (page 26)).

Note: Security

The FTP access is writable only (no read). It allows only to put a Securactive signed and encrypted file. This file will be automatically moved, checked and executed by an internal process.

ServicePack

In rare cases, it's needed to upgrade some third-party internal softwares. The information is available in the release note of the new version. These packages are called Service Packs. To apply them, put the file (SPV-ServicePackX-rY.bin) using the same method.

4.3 System

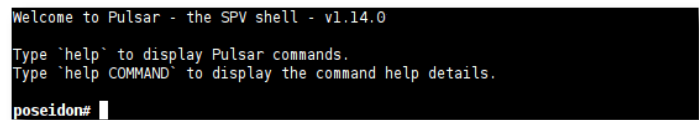
The probes come with a Command Line Interface named *Pulsar*. This allows the user to check the probe state and configure it when needed.

4.3.1 Connect to the probe

If this is your first encounter with the probe, you will have, for the first time only, to access to the probe physically (just use a screen and keyboard plugged to the probe). Log in with user `admin` and default password `admin`. Once the network address of the the probe will have been set-up you will be able to access to it directly through SSH on port 22 also with the same user `admin`.

4.3.2 Pulsar

When logged in you should see the following prompt (version number can vary).



```
Welcome to Pulsar - the SPV shell - v1.14.0
Type 'help' to display Pulsar commands.
Type 'help COMMAND' to display the command help details.
poseidon#
```

Figure 4.1: Pulsar prompt on the *poseidon* probe

Note: Pulsar uses 3 colors while displaying informations.

- **Green** outputs are *informations*.
- **Yellow** outputs are *warnings*.
- **Red** outputs are *errors*.

If needed you can set the keyboard mapping with the `kb <mapping>` command. Typing `kb` displays the list of available mappings.

Pulsar allows you to change the administration password through `passwd` command. This should be your first command. Typing `passwd` in the pulsar shell launches the standard UNIX password-change process.

Warning: At this point, there is no way to retrieve the password. If you totally lost the password, the Securactive support team can generate a new one. See [Support access through VPN](#) (page 28). You can also restore the probe, see [Restore probe state](#) (page 27).

4.3.3 Configure the probe

Use the `config` command to setup up the probe.

```
pulsar# config
service:
1. dns
2. hostname
3. network
4. ntp
5. smtp
6. support
7. **all [default]**
Your choice?
```

Typing `enter` will launch the whole interactive configuration process.

Warning: This command is **mandatory** as it will configure key elements needed for proper operations (DNS servers, hostname, IP address, NTP, SMTP...).

Some changes in configuration require to reboot the probe (command: `reboot`).

Restore probe state

You may need to restore some probe original configuration. There are three way of achieving this. As these are destructive commands a strong confirmation will be requested.

You want to erase any single data from your previous network captures. This preserves configuration settings and IHM user accounts.

```
pulsar# reset data
...
Stopping services...
Deleting data...
Done.
```

The command `reset all` will destroy both your configuration and capture database. You will have a fresh new database. Configuration settings, users and pollers will be reset to default values.

```
pulsar# reset all
...
Stopping all services...
Resetting...
Creating default settings...
Done.
```

Restoring data hard drive disks

This is to be used when you are delivered new data disk(s). If you want to use it anyway, any existent data (capture and configuration) will be lost. Default values will be restored.

```
pulsar# format_data_disk

These processes should not be interrupted. Do NOT use Ctrl-C.
Preparing disk ...
Formatting disk ...
Installing disk ...
Generating database ...
This may be quite long (5 min) ...
Done.
```

More about pulsar

`help` provides both global and command help. Tab-completion is enabled for commands and subcommands such as `help`, `config` and `show`.

Configuration example

```
pulsar# config network
[NETWORK]
Connection Type:
1. Static network
2. DHCP
```

```
Pulsar shell: configure your probe
help (?)          quit (exit)

===== SPV =====
analyzer         csv_status   format_data_disk poller
process          reset        status          supervisor
=====
config           kb           show
=====
===== system commands =====
bmon             date         dig            ethtool
halt             host         ifconfig       ip
nslookup         ntpdate     partitions (df) passwd
ping             reboot      set_date       set_time
tcpdump          top          traceroute     uname
uptime

===== system info =====
dns              hostname     log            network
ntp              smtp         snmp           support
```

Figure 4.2: Available commands

```
Your choice? 1
IP address: 192.168.1.1
netmask: 255.255.255.0
gateway: 192.168.1.254
```

Support access through VPN

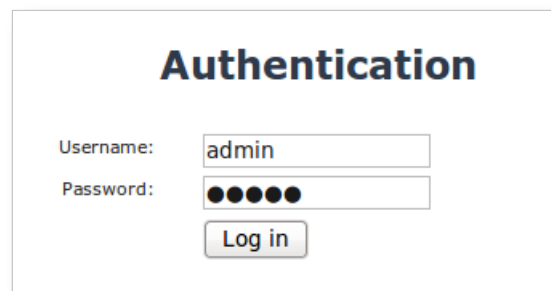
The probes come with an already configured VPN connection to allow access for support operations, if needed. The VPN address is set by default and should normally not be changed. If it needs to be changed, this can be done by the command `config` and option 7. The VPN service is stopped by default. It can be started or stopped at any moment by the corresponding commands `support start` or `support stop`.

Note: In order to have the VPN connection of the probe working fine, you will probably have to configure your network and/or security equipment like your firewalls. Default Host IP address is 88.191.121.167 and default port is 443.

4.3.4 Access Performance Vision

Through a Web browser

We assume here that the probe has been previously configured through the command line interface and the user knows the probe IP address. The probe can be accessed either with SSH or with a Web browser. To connect with a Web browser, the port to use are the 80, 8080 or 443.



```
Authentication

Username:  admin
Password:  ●●●●●
          [Log in]
```

Figure 4.3: Login parameters in SPV

Thus, if the IP address of the probe has been configured as 10.0.0.1, then just open the URL `http://10.0.0.1` with your Web browser (or `https://10.0.0.1` for using the HTTPS protocol).

Please note that you can verify that you are actually connected to a Performance Vision appliance, by checking that the certificate serial number is 00:90:26:d5:46:2a:5e:66:ec.

To log in, please use: `admin` as user and `admin` as password.

You are now logged in and ready to use the Graphical User Interface. In order to offer the best performances, the use of Mozilla Firefox is recommended.

4.3.5 How to configure User Interface language ?

User interface is available in English and French languages. The language is detected automatically based on the default language of the browser used to access the probe. So, to get the User interface to use the desired language, the administrator should check and configure the default language of its browser.

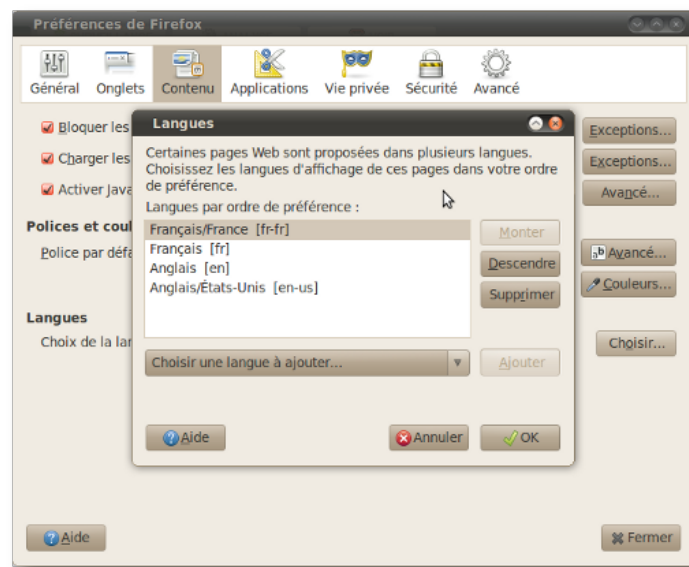


Figure 4.4: Configuration of the French language in Firefox

4.4 SPV Functional Configuration

4.4.1 User Management

There are two groups of users in the "Users Configuration" interface:

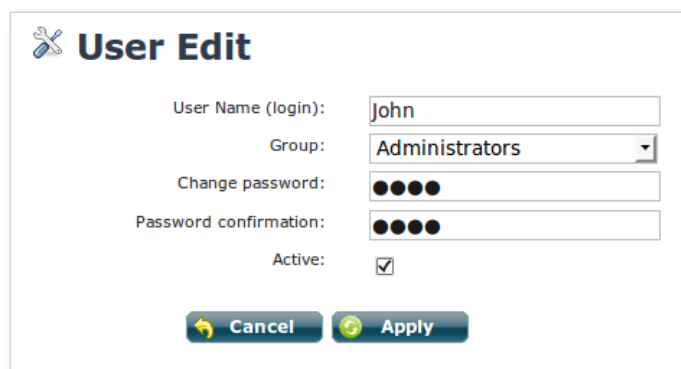
- The *Administrators* group
- The *Users* group

These two groups have different access permissions to the application pages: the administrators group provides its members a full access to the "Configuration" pages. Users group members will be able to read reports but will not have access to the configuration page.

In order to create a new user account you must be logged into the appliance as a member of the *Administrator* group. As mentioned in the above paragraph, the default admin group has the right to create, modify and access the configuration. You can add a new user account by clicking on the *Users* tab found on the configuration menu on the left hand side. Then click on the *Add* button and fill in the "User information" (username, password, and

group). Make sure the `Active` button is checked, otherwise the user won't be able to login. Thanks to this option you will be able to disable or enable an account without deleting it.

Example: Adding a new member to Administrators group. In the example below, we have created a user account in the `Administrators` group with the user name `John` and `foo2` as the password:



The 'User Edit' form contains the following fields and controls:

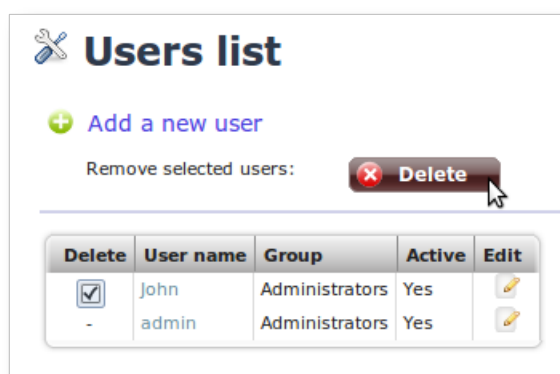
- User Name (login):** Text input field containing 'John'.
- Group:** Dropdown menu showing 'Administrators'.
- Change password:** Password input field with five dots.
- Password confirmation:** Password input field with five dots.
- Active:** Check box that is checked.
- Buttons:** 'Cancel' and 'Apply' buttons at the bottom.

Figure 4.5: Edit User

The user name is case sensitive, and it is required to be non-empty and to contain only letters, numbers, or _ (underscores).

You can modify a user account by clicking on the `Users` tab found on the configuration menu on the left hand side, and then clicking on the user name of the desired user account in the user list. You will be able to modify any field on a created user. Please note that the password field will appear empty on edition to avoid giving out information and will not be modified upon edition if it is left empty. In order to save any modifications, click on the `Apply` button.

You can delete a user account by clicking on the `Users` tab found in the configuration menu on the left hand side, and then clicking on the check box next to the user name of the account you wish to delete. Then, clicking on `Delete` button will delete all selected Users.



The 'Users list' interface shows a table of users with the following data:

Delete	User name	Group	Active	Edit
<input checked="" type="checkbox"/>	John	Administrators	Yes	
<input type="checkbox"/>	admin	Administrators	Yes	

At the top of the interface, there is a '+ Add a new user' link and a 'Remove selected users: Delete' button. A mouse cursor is pointing at the 'Delete' button.

Figure 4.6: User account 'John' is about to be deleted.

4.4.2 Zone configuration

The aim of this chapter is to help the administrator of the platform to configure zones. When you change or create a *zone*, the modifications will be immediately applied for future integrated data but not to the already captured data which keep their old zone attribute.

How to access the configuration menu?

After clicking on the top right `configuration` button, you will observe a tree configuration menu with different items.

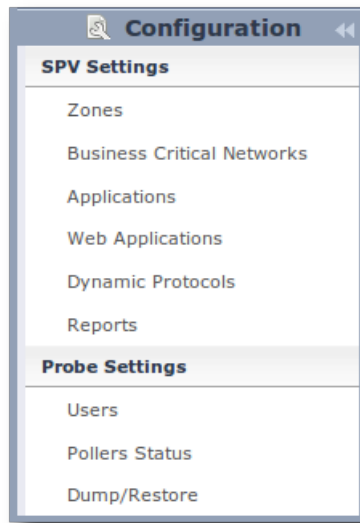


Figure 4.7: Configuration menu

Zones management

Please refer to *Zones & Fallbacks* (page 5) for *Zone* tree and *Fallback* explanations.

You can reach the zone configuration page by clicking on the `Zones` label of the menu. The illustration below displays a list of zones and their subnets. This list of zones enable you to add a zone, edit a zone or delete a zone as needed.

In order to create a zone, you need to click on the green button below:




Figure 4.8: Button to create a zone

You can select the zone's name by filling the *Name* field and the subnet by filling the new subnet field. In order to add several subnets you can use the `more` button.

Here are some examples of valid subnets

- 192.168.100.0/24
- 192.168.100.12/32
- ::ffff:192.168.0.0/128

The administrator can:

- rename the zone



Add a zone

Name:

Parent zone:

Zone inner subnets

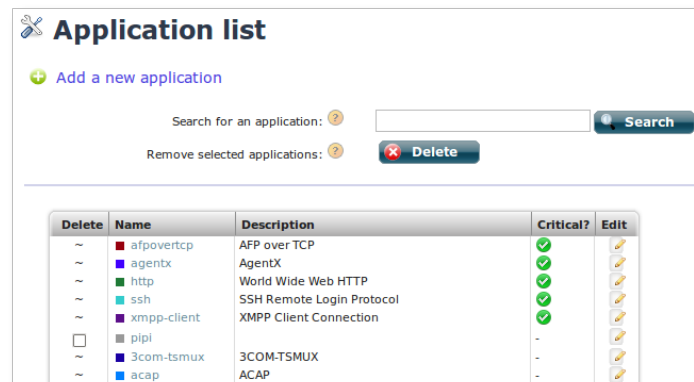
New subnet:

- add a new subnet or IP Address
- change the previous subnets
- delete a zone.

4.4.3 Application configuration

You can configure *Applications* in the configuration page. Applications represent the business applications running on your network and make the reports provided easily understandable to everyone in your organization.

To access the configuration of *Applications*, click on the *Configuration* button, on the top right of the user interface.



Application list

[+ Add a new application](#)

Search for an application:

Remove selected applications:

Delete	Name	Description	Critical?	Edit
<input type="checkbox"/>	afpovertcp	AFP over TCP	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
<input type="checkbox"/>	agentx	AgentX	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
<input type="checkbox"/>	http	World Wide Web HTTP	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
<input type="checkbox"/>	ssh	SSH Remote Login Protocol	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
<input type="checkbox"/>	xmpp-client	XMPP Client Connection	<input checked="" type="checkbox"/>	<input type="button" value="Edit"/>
<input type="checkbox"/>	pipi		-	<input type="button" value="Edit"/>
<input type="checkbox"/>	3com-tsmux	3COM-TSMUX	-	<input type="button" value="Edit"/>
<input type="checkbox"/>	acap	ACAP	-	<input type="button" value="Edit"/>

Figure 4.9: Application list screen

To create an application, go the to *Application* submenu, in the left menu.

This panel displays the existing Applications (by default or user defined). To create an Application, click on *Add*; you will see the configuration screen.

An Application can be defined using the following elements:

- **Name:** it corresponds to the designation of each Application, which will be used in displays. This is a mandatory field.
- **Color:** it is the color which will be used to display this specific Application in graphs. This is a mandatory field.
- **Comment:** it is a description field, which should be used to track information related to this Application.

Application Add

Application information

Application Name:

Application Color:

Application Description:

Protocol

Set of protocols and ports ranges (any of them will flag the application).

New protocol: Ports:

New protocol: Ports:

Signature

Choose dynamic ports and web applications (any of them will flag the application).

New signature:

Server zone

Smaller zones have priority over largest ones.

New server zone:

Client zone

Smaller zones have priority over largest ones. Server zones have priority over this client zones.

New client zone:

Business critical application thresholds

Application is business critical ☐

Figure 4.10: Application configuration screen

The following elements are combined to define patterns which the application flows must match:

- **Destination ports:** each line can be used to define a range of ports on a specific protocol (UDP or TCP), which the flows for this Application must match.
- **Signature:** a signature is an layer 7 Application pattern which makes it possible to define an Application based on pattern matching in the payload of packets. Signatures may be for dynamic ports Applications or Web patterns (more details in the section hereunder).

Warning: In case of a web application, these two parameters may be filled and but not necessarily (we consider a **OR operand** between this two parameters **Destination ports** and **Signature**).

Client or *Server* zone indicates the zones in which the servers of the Application are located.

Warning: In case of a application associated to one server (or a server farm), the server zone edit box must be filled with the right definition of the zone defined in the zone configuration.

In the example here above, we created this Application: application `SAP-Sales` will be display in red, corresponds to flows on server port TCP 8080 **OR** range of UDP ports from 8080 to 8090 only if the flow is sent to a server in Zone 'VLAN_Sales'.

Application signature

application signature is defined as a pattern recognized in the payload of a packet that makes the identification of the application possible. There are two types of **application** signatures:

- *Signature Dynamic port:* these signatures are used to identify applications using dynamically negotiated ports (e.g. Passive FTP, Bittorrent, ...) through connection tracking. Additionally, Performance Vision supports protocol recognition for the following applications: FTP, SIP, MGCP.

The ports on which to search for these applications can be configured on this panel:

Protocol	Port Range
ftp:	21
dns:	53
http:	80,8000,8080,3128
mgcp:	2427,2727
sip:	5060
bittorrent:	6881-6999

- *Signature Web application:* these signatures are used to identify applications using URLs in HTTP flows. They are defined as patterns matched against the URLs contained in HTTP requests. The patterns should contain at least a domain name, optionally including wild card characters like '*' or if you check 'regex mode', you can set POSIX regular expressions:

4.4.4 Business Critical applications

An application can be tagged as **Business Critical**. Those applications are used to display the '*Business Critical Application Dashboard*' for now and will be also used in future dashboard. Business Critical is an additional




Figure 4.11: Web application signature configuration

attribute of the application.

From here you can:

- Add a new Critical Application (Add button).
- Edit the parameters of a Critical Application (Edit button).
- Remove a Critical Application (Remove button).

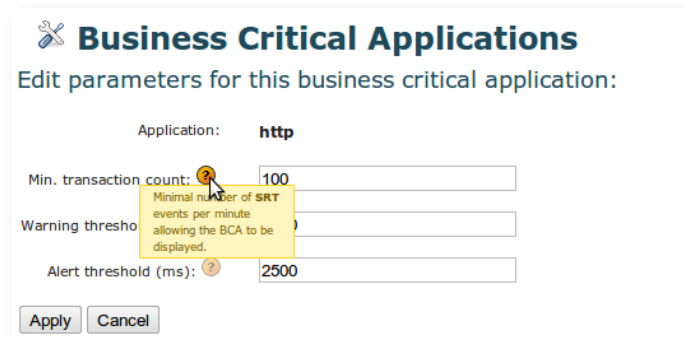


Figure 4.12: Business critical application edition

When you add a new *Critical Application*, three parameters are required for defining a critical thresholds:

- **The minimum transaction count.** It indicates, for one minute, the minimum of SRT (Server Response Time) events to be seen on the network for being considered as a pertinent measurement. If no transaction at all is seen during the period of time analyzed, the color displayed on the *BCA dashboard* will be “white”. If the number of events seen during the period of time analyzed is above zero but under this value, the color displayed on the *BCA dashboard* will be “grey”. It means that some events have been seen, but not enough to be considered as a pertinent measurement. If the number of events seen during the period of time analyzed is above or equal to this value, the color displayed on the *BCA dashboard* will be either “green”, “orange” or “red” depending on the *EURT* values.
- **The warning threshold level of the EURT (End User Response Time) value in milliseconds.** When the value is above or equal to this level, the color displayed on the *BCA dashboard* will be “orange”. When the value is under this level, the color displayed on the *BCA dashboard* will be “green”.
- **The alert threshold level of the EURT value in milliseconds.** When the value is above or equal to this level, the color displayed on the *BCA dashboard* will be “red”.

Note: To be useful and pertinent, these parameters must be accurate values adjusted to your network configuration. These values can be easily changed for fine tuning or to cope with any change in the network or applications you are using.

A new critical application, will benefit of all the data history. So after having defined an application as critical, if the data has already been collected for this application, the thresholds levels will be automatically applied on the *BCA dashboard*, even for a period back in time.

4.4.5 Business Critical Networks

A BCN consists of a *virtual link* between two zones; its objective is to monitor normal volume and performance levels between two network segments, which represent a strategic network link for your organization (e.g. link from the data center to a remote site, from the server VLAN to a user VLAN). An administrator can configure thresholds for warning and alert on bandwidth consumption, Retransmission Rate (RR) and Round Trip Times (RTT).

A specific configuration screen allows configuring the specified BCN. To access it, just go to the *Configuration* menu and choose the entry labeled *Business Critical Networks*.

Figure 4.13: Editing an existing Business Critical Network

From here you can add a new BCN or edit the parameters of an existing BCN. Modifications will also be applied on already captured traffic.

For each Critical Network, you have to configure the following parameters:

- The source/destination network zones.
- One or several thresholds for both **Warning** and **Alert** levels, all these thresholds are computed from source to destination and not from client to server. We call this an “oriented” metric:
 - Oriented latency (RTT in ms)
 - Oriented retransmission rate (%)
 - Utilization rate (%) according to bandwidth available (Mib/s)
- A minimum volume for triggering (Mib/s). This value represents the minimum bandwidth observed from which you will consider the performance and volume thresholds as relevant.
- The thresholds values can be configured as symmetric by ticking the **Symmetric Link** check-box or be configured as distinct values for both directions. This is particularly useful when the critical network:
 - refers to **asymmetric** connections like ADSL,
 - has one of its zones **closer** to the poller than the other zone and latency (RTT) computation is impacted (see *Distributed Architecture* (page 19)).

You can define thresholds from either one criterion or more (any of the following: latency, retransmission rate and consumption level). But you cannot define a BCN from one zone to itself, as their intended purpose is to check the performance of most important links or routes between two network segments.

By applying your changes, the *BCN Dashboard* will be updated in accordance with the new threshold values (including already captured data). To be useful and pertinent, these parameters must be accurate values adjusted to your network configuration. These values can be easily changed for fine tuning or to cope with any change in the network or applications you are using.

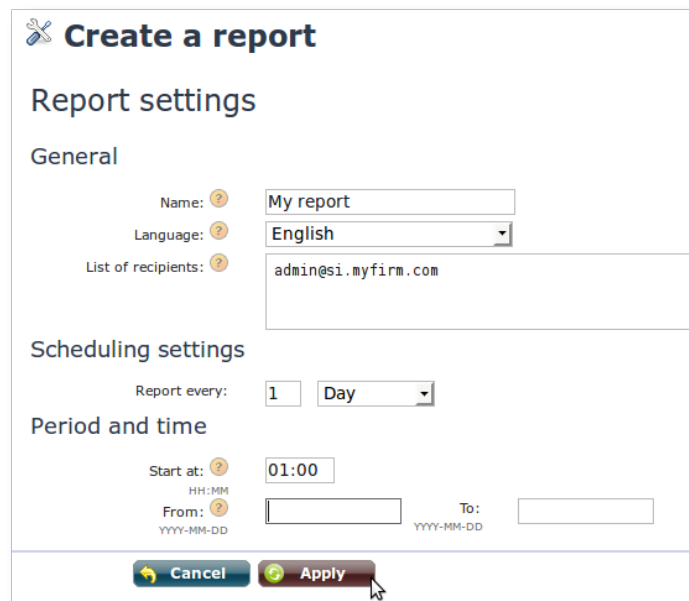
4.4.6 Reports

Creating *Reports* is just a matter of a few clicks. You can easily create and define exactly the level of information you want to get. You will receive it directly in your mailbox or via *FTP* at the frequency you prefer.

Configuration

In the first step, you start by creating a template that will mainly define the name of the report, the list of recipients and the scheduling settings. In the second step you just have to add the different views you want to see to the appropriate template. Then you're done, just check your mailbox.

To create a report template, in the *Configuration* area, select *Reports* in the menu list on the left. This will display the list of existing report templates. Use the button *Create* to create a new report template. Please note that this feature is only available for users with administration rights.



The screenshot shows a web-based form titled "Create a report". It is divided into two main sections: "Report settings" and "Scheduling settings".

Report settings:

- General:**
 - Name:** A text input field containing "My report".
 - Language:** A dropdown menu set to "English".
 - List of recipients:** A text input field containing "admin@si.myfirm.com".

Scheduling settings:

- Report every:** A numeric input field set to "1" and a dropdown menu set to "Day".
- Period and time:**
 - Start at:** A time input field set to "01:00".
 - From:** A date input field with a placeholder "YYYY-MM-DD".
 - To:** A date input field with a placeholder "YYYY-MM-DD".

At the bottom of the form are two buttons: "Cancel" and "Apply". A mouse cursor is pointing at the "Apply" button.

Figure 4.14: Create a new report

To create a report template you must fill some information:

- The name of the report for easy identification purpose,
- The language option defines the language that will be used for this reports (thus the language for the report can be different than the language of the web screen),
- The list of recipients defines the email addresses to which the reports will be sent (the recipients email addresses can be separated by a comma, a semi-colon or a new line),
- Scheduling settings define the frequency at which the reports will be sent. Available options are:
 - **Day:** Generates the report every x day(s); example: every two days.
 - **Week:** Generates the report every x week(s) the selected days; example: every two weeks on Friday (several days in the week can be chosen).
 - **Month:** Generates the report every x month(s) on y day; example: every month, the first of the month (be careful, if you choose the day 29, 30 or 31, you will only receive your reports if there is such day in the corresponding month).
- **Start at** defines the hour (format HH:MM) at which the generation of the report will start. Once the report will have been generated it will then be sent to the recipients email addresses.

- From and To fields are optional. This allows you to define a validity period for the report. In such case, the report will only be sent in the period ranging from the first date up to the second date.

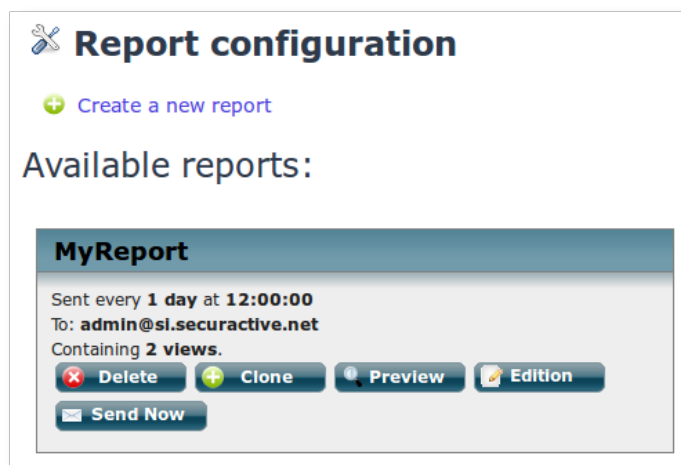


Figure 4.15: Report: A template just created

The new report template just created will appear in the list of available report templates. A summary is displayed (scheduling frequency, generation time, first recipient emails). At this stage it is empty and does not contain any view, this is why you have Containing 0 views indicated. After having added some views to the report, here will be indicated the number of views contained in the report.

Add views to report

To add a view to a report template, just go to the screen with the desired view. Select a time period and run the search. Once search is completed, the link `Add this page to a report` becomes active. When you click on it, a drop box with the list of available template reports is displayed. You can chose the template report to which you want to add the current view and click on the button `Add`. If you need, you can click on `Show report list`, it will open the configuration area with the list of available report templates.

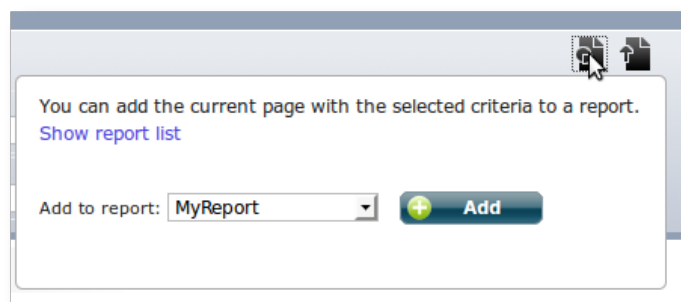


Figure 4.16: Add a view to a report template

Please note that while the time is fixed, the date will remain relative to the moment the report is sent. If the view you're adding starts yesterday at 20:00 and ends today at 8:00, and the report is scheduled to be sent next Friday, then the effective capture time bracket will be from Thursday at 20:00 to Friday at 8:00.

(Note: Before release 2.9, an additional time delta was added under certain circumstances. As of 2.9 it's not longer the case; all dates are relative to the day the report is being sent.)

Actions on reports

A report template can be deleted with the button `Delete`. You can clone a report template: all its parameters and included views will be duplicated. A new report template is created with `(copy)` added to the report name.

`Preview` will start the generation of the report right now and you will be able to see the PDF file with your favorite PDF viewer once it has been generated.

`Edit` allows you to change the parameters of the report template (name of the report, the list of recipients and the scheduling settings...).

`Send now` will start the generation of the report right now and the report will be sent by mail once it has been generated.

Sending Email

So that your reports could be sent properly to the recipients email addresses, you need to configure the *SMTP* server within *Pulsar* (page 26). You can do that with the `config smtp` command. Then just add a valid *SMTP* host, and in option a *login* and *password* if you use an authenticated *SMTP* server. You also can modify (with the same command) the *From* header of the emails generated by the probe.

After that you can either reboot the probe or use `smtp stop` followed by `smtp start` commands to activate the new configuration.

4.4.7 SNMP

Optionally, *SNMP* requests are answered on default `SNMP port` (see *Pulsar* (page 26) documentation). The *SNMP* objects that are thus made available are twofold. First there are the standard *SNMP* objects then SPV specific objects.

System MIB

The probe uses the UNIX `Net-SNMP`¹ daemon, which serves standard *MIB*. So you can monitor your probe from your *SNMP* console as you would normally monitor any UNIX server. For instance the usual statistics about network interface usage, file system available spaces, I/O operations, etc, are available.

Monitoring specific MIB

In addition to these default information the probe provides various statistics under `iso.org.dod.internet.private.enterprises.securactive`.

The comprehensive *MIB* files are available from our web site² so this section only sketches what kind of information is available. You are encouraged to download the actual *MIB* for use with your common purpose *SNMP* console. This will give you access to:

- *Interface statistics* for each network interface, such as the count of received packets, dropped packets and duplicated packets.
- *Protocol statistics* for each recognized protocol, which can give a good impression on the realtime composition of the whole network stream.
- *Various CPU/RAM information* that are destined to troubleshoot an SPV more than to reveal anything about the network.
- *License related information* such as date of expiry and so on.
- *Averaged metrics* such as *RTT* or *DNS* response time.

¹ <http://www.net-snmp.org>

² <http://www.securactive.net/en/documents/250-securactive-mibs/download>

BCN and BCA MIBS

Since the 2.9 version, two new modules are available: BCA and BCN. Please update your MIB file if you use a SPV MIB before 2.9. Here is a tree description of the BCA and BCN MIB:

BCA module

```
+--sactSPVBCAModule(1)
|
+--spvBCAStateTable(1)
| |
| +--spvBCAStateEntry(1)
| |   Index: spvBCAName
| |   |
| |   +-- -R-- String      spvBCAName(1)
| |   +-- -R-- EnumVal     spvBCAStatus(2)
| |   |   Values: Ok(1), Warning(2), Alert(3), NA(4), Nothing(5), NotEnough(6)
| |   +-- -R-- Gauge      spvBCAEURT(3)
| |   +-- -R-- Gauge      spvBCASRT(4)
| |   +-- -R-- Gauge      spvBCASRTCount(5)
| |   +-- -R-- Counter     spvBCASRTCountSum(6)
| |   +-- -R-- Gauge      spvBCARTTClient(7)
| |   +-- -R-- Gauge      spvBCARTTServer(8)
| |   +-- -R-- Gauge      spvBCADTTClient(9)
| |   +-- -R-- Gauge      spvBCADTTServer(10)
| |   +-- -R-- Gauge      spvBCATrafficClient(11)
| |   +-- -R-- Gauge      spvBCATrafficServer(12)
| |   +-- -R-- Counter     spvBCATrafficClientSum(13)
| |   +-- -R-- Counter     spvBCATrafficServerSum(14)
| |   +-- -R-- Gauge      spvBCAThresholdMinSRTcount(15)
| |   +-- -R-- Gauge      spvBCAThresholdWarning(16)
| |   +-- -R-- Gauge      spvBCAThresholdAlert(17)
| |
+--spvNevraxBCATime(2)
```

BCN module

```
+--sactSPVBCNModule(2)
|
+--spvBCNStateTable(1)
| |
| +--spvBCNStateEntry(1)
| |   Index: spvBCNName
| |   |
| |   +-- -R-- String      spvBCNName(1)
| |   +-- -R-- String      spvBCNZoneA(2)
| |   +-- -R-- String      spvBCNZoneB(3)
| |   +-- -R-- EnumVal     spvBCNGlobalStatus(4)
| |   |   Values: Ok(1), Warning(2), Alert(3), NA(4), Nothing(5), NotEnough(6)
| |   +-- -R-- EnumVal     spvBCNStatusAtoB(5)
| |   |   Values: Ok(1), Warning(2), Alert(3), NA(4), Nothing(5), NotEnough(6)
| |   +-- -R-- EnumVal     spvBCNStatusBtoA(6)
| |   |   Values: Ok(1), Warning(2), Alert(3), NA(4), Nothing(5), NotEnough(6)
| |   +-- -R-- Gauge      spvBCNRttAtoB(7)
| |   +-- -R-- Gauge      spvBCNRttBtoA(8)
| |   +-- -R-- Gauge      spvBCNRrAtoB(9)
| |   +-- -R-- Gauge      spvBCNRrBtoA(10)
| |   +-- -R-- Counter     spvBCNRetransCountSumAtoB(11)
| |   +-- -R-- Counter     spvBCNRetransCountSumBtoA(12)
| |   +-- -R-- Gauge      spvBCNBandwidthAtoB(13)
```

```
|      +--- -R--- Gauge      spvBCNBandwidthBtoA(14)
|      +--- -R--- Counter    spvBCNTrafficSumAtoB(15)
|      +--- -R--- Counter    spvBCNTrafficSumBtoA(16)
|      +--- -R--- Counter    spvBCNPacketsCountSumAtoB(17)
|      +--- -R--- Counter    spvBCNPacketsCountSumBtoA(18)
|      +--- -R--- EnumVal    spvBCNThresholdSymetricLink(19)
|      |      Values: True(1), False(2)
|      +--- -R--- Gauge      spvBCNThresholdBandwAvailableAtoB(20)
|      +--- -R--- Gauge      spvBCNThresholdBandwAvailableBtoA(21)
|      +--- -R--- Gauge      spvBCNThresholdBandwMinAtoB(22)
|      +--- -R--- Gauge      spvBCNThresholdBandwMinBtoA(23)
|      +--- -R--- Gauge      spvBCNThresholdBandwrateWarningAtoB(24)
|      +--- -R--- Gauge      spvBCNThresholdBandwrateWarningBtoA(25)
|      +--- -R--- Gauge      spvBCNThresholdBandwrateAlertAtoB(26)
|      +--- -R--- Gauge      spvBCNThresholdBandwrateAlertBtoA(27)
|      +--- -R--- Gauge      spvBCNThresholdRttWarningAtoB(28)
|      +--- -R--- Gauge      spvBCNThresholdRttWarningBtoA(29)
|      +--- -R--- Gauge      spvBCNThresholdRttAlertAtoB(30)
|      +--- -R--- Gauge      spvBCNThresholdRttAlertBtoA(31)
|      +--- -R--- Gauge      spvBCNThresholdRrWarningAtoB(32)
|      +--- -R--- Gauge      spvBCNThresholdRrWarningBtoA(33)
|      +--- -R--- Gauge      spvBCNThresholdRrAlertAtoB(34)
|      +--- -R--- Gauge      spvBCNThresholdRrAlertBtoA(35)
|
+---spvNevraxBCNTime(2)
```

Note: Notice that none of these *MIB* objects is currently settable.

INTERPRETING THE RESULTS

Note: Note about terms used: starting from version 2.8, The in/out notion has been fully replaced by Server/Client. So in our Graphs, any RTT and RR (in/out) should be considered as RTT,RR (Server/Client) as in the following rules.

- RTT in stands for RTT Server.
- RTT out stands for RTT Client.
- RR in stands for RTT Server.
- RR out stands for RTT Client.

5.1 Business Critical Application Dashboard

To customize this view for your own needs, just go to the *Configuration* menu and choose the application you want to be a 'business' one. (see the *Business Critical applications* (page 34)).

The purpose of the Business Critical Application Dashboard (BCA) is to have, regrouped into one single view, the most important elements that are critical for your business. In one single screen vital information is presented to people in charge in order to radically improve early diagnostics and impact analysis. The right information is directly available through a completely configurable and dynamic dashboard view. What is monitored is the EURT (End User Response Time) metric. Thus, this dashboard reflects the quality of experience of the users for the selected critical applications.

- In red: poor quality
- In orange: medium quality
- In green: good quality
- In grey: not enough data gathered




Figure 5.1: Business Critical Application Dashboard view

5.1.1 Business Critical Application Dashboard Capabilities

- You can customize the business critical dashboard to view specific applications and metrics corresponding to your specific business.

- From the BCA dashboard, you can drill-down from the general view to detailed analysis and problem resolution views.



Figure 5.2: Quick links in the Business Critical Application Dashboard view.

Thus, from each Business Critical Application, with a single click on the appropriate icon, you can:

- Directly access to the corresponding Application Dashboard,
- Add a filter on this specific Critical Application (in case you have defined a lot of Critical Applications and you want to see only one for a moment),
- Edit Application characteristics.
- Directly access to the details of Conversations for this Application.

Note: If you click on the icons that are next to the name of the application at the beginning of each line, the quick links will take into account the complete period of time currently displayed. If you click on the icons associated to a specific period of time, the quick links will use this specific period time when redirecting you to a detailed screen.

- You will always see up-to-date information with the auto-refresh feature of the BCA dashboard. The information will be automatically refreshed based on the data aggregation level (see [aggregation period](#)). For example if the “Aggregate level” is “2 minutes”, the BCA will be updated every two minutes; if the “Aggregate level” is “15 minutes”, the BCA will be updated every fifteen minutes.

5.2 Business Critical Networks Dashboard

To customize this view for your own needs, just go to the *Configuration* menu and choose the entry labeled Business Critical Network (see the [Business Critical applications](#) (page 34)).

The Business Critical Network Dashboard (BCN) is aimed at presenting in a single screen the status of your organization’s most critical network “links”. You can customize the business critical network dashboard to view the status of the most strategic links corresponding to your business.




Figure 5.3: Business Critical Network Dashboard

From the Business Critical Network Dashboard, you can drill down from the general view to more detailed information for analysis and problem resolution:

By pointing with the mouse, you can view the threshold values for each direction at each point of time (indicating status OK, Warning or Alert as well as the value for each direction). You can also access to the bandwidth graphs and the conversations table for each link. If you click on the icons that are next to the name of the link at the beginning of each line, the quick links will take into account the complete period of time currently displayed. If you click on the icons associated to a specific period of time, the quick links will use this specific period time when redirecting you to a detailed screen. You will always see up-to-date information with the auto-refresh feature of the *BCN dashboard*. The information will be automatically refreshed based on the data aggregation level (see [aggregation period](#)). For example if the “Aggregate level” is “2 minutes”, the BCN will be updated every two minutes; if the “Aggregate level” is “15 minutes”, the BCN will be updated every fifteen minutes.




Figure 5.4: Detailed values for a point of time

5.3 VoIP Module

A specific reporting for *Voice over IP* traffic is provided. The aim of this module is to show the volume and quality of service associated with VoIP flows.

5.3.1 Supported protocols

These two VoIP set of protocols are supported:

- SIP + RTCP + RTP
- MGCP + RTCP + RTP

For more information, please consult the corresponding *RFCs*:

- SIP as defined in **RFC 3261** (<http://tools.ietf.org/html/rfc3261.html>)
- MGCP as defined in **RFC 3435** (<http://tools.ietf.org/html/rfc3435.html>)
- RTP as defined in to **RFC 3550** (<http://tools.ietf.org/html/rfc3550.html>) and **RFC 3551** (<http://tools.ietf.org/html/rfc3551.html>)
- RTCP as defined in **RFC 3605** (<http://tools.ietf.org/html/rfc3605.html>)

5.3.2 Basics of VoIP

Voice Over IP relies on three protocols to operate over IP networks:

- **Signalization protocol:** the role of this protocol is to establish and control the voice communications. It usually consists of communications between the IP phone and a call manager / IPBX. The 2 signalization protocols supported are SIP (Session Initiation Protocol) and MGCP (Media Gateway Control Protocol). Please note that SIP may follow the same route as the RTP traffic or not, while MGCP follows the same route as RTP.
- **Media protocol:** the role of this protocol is to carry the voice signal from one IP phone to the other IP phone (it can eventually go through the call manager / IPBX). RTP is the only media protocol supported by *Performance Vision*. It stands for *Real Time Protocol*; it usually runs over UDP.
- **Control protocol:** the role of this protocol is to carry quality and control information from one phone to the other phone. RTCP is the only control protocol supported. It stands for *Real Time Control Protocol*.

5.3.3 Quality of service & MOS

MOS stands for Mean Opinion Score. It is a numeric indication of the perceived quality of service of VoIP. It is expressed by a number ranging from 1 to 5, 1 corresponding to the lowest quality and 5 to the highest (close human voice).

MOS Rating	Meaning
5	Excellent
4	Good
3	Fair
2	Poor
1	Bad

Please note that in real network a MOS note of over 4.4 is unachievable. A low MOS will translate into echo and degraded signal. MOS is in principle the result of a series of subjective tests; in the context of network analysis, MOS will be estimated using a formula that integrates 3 factors:

- Network latency (RTT recommended value: <100ms)
- *Jitter* (recommended value: <30ms)
- Packet loss rate (recommended value: <5%)

5.3.4 Prerequisites

To provide MOS values for VoIP traffic, it is necessary to capture the three flows: signaling (SIP or MGCP), media (RTP) and control protocol (RTCP). If one of these flows is not present in the traffic capture brought to the listening interface(s), the MOS value will not be calculated. Other quality of service metrics will remain available.

Protocol	Metrics obtained by analysis of the protocol
SIP/MGCP	<ul style="list-style-type: none"> • Sign. RTT (network latency between each phone – value in & out interval between a request and the first response (definitive or temporary) from the signaling server) • Sign. SRT (signaling server response time) • Sign. RD (retransmission delay for the signaling traffic) • Sign. RR (retransmission rate for the signaling traffic) • Code (indicates how the VoIP call ended – e.g. error or not; please note that the code depends on the protocol used)
RTP	<ul style="list-style-type: none"> • Jitter (standard deviation of latency for the media traffic going from one IP phone to the other) • Packet loss (percentage of packets lost in the conversation at the point of capture of the probe-based on RTP sequence numbers)
RTCP	<ul style="list-style-type: none"> • RTT (network latency between the two IP phones – based on the timestamps provided by both IP phones)

Note: RTT and MOS values depend to some extent on the quality of the measurement provided by RTCP. Please note that MOS is not very sensitive to “normal” latency values. When referring to voice or media, we refer to the RTP traffic, which may correspond to different things (human voice, prerecorded message, ring back tone, busy

line tone, ...) The *VoIP* module discards the *jitter* and packet loss data present in the RTCP flow and replace them with equivalent values computed internally. This is so for several reasons:

- It was observed that many softphones do not place accurate (or even credible) values in these fields,
- RTCP stream is more often missing than present, probably because it is firewalled and of little use to the *VoIP* client software.

For the *VoIP* module to remain passive, there is no other option than compute these values for every RTP stream to generate jitter and packet loss values which will be a good estimate of the real jitter and loss experienced by both users. This is how, even, in the absence of RTCP stream, we can display a jitter and packet loss count (and no RTT, and thus no MOS).

5.3.5 VoIP Views


VoIP Overview

VoIP Overview is a view of all VoIP traffic in the network, zone per zone:

- Number of calls
- MOS value
- Packet loss (global or caller / callee)
- Jitter (global or client / server)
- RTT (global or client / server)

Note: The value “caller” / “in” corresponds to the metric for the RTP/RTCP traffic from the caller to the callee and the value “callee” / “out” corresponds to the metric for the RTP/RTCP traffic from the callee to the caller. From each line, you drill down:

- to the MOS chart,
- to the *VoIP* conversations.




MOS over time

This view shows the evolution of the *Mean Opinion Score* through time. A second graph shows the evolution of the number of calls, to help you evaluate how many were impacted by a MOS degradation.

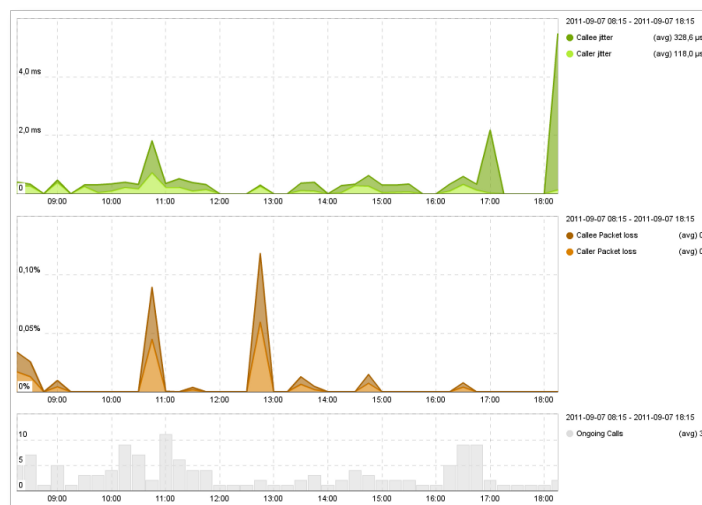
- By pointing a specific point of time on the graph, you can display the exact value for each metric on the right side of the graph.
- By clicking on a specific point of time, you are directly to the *VoIP* conversations for this point of time.

Jitter / Packet Loss

This view shows the evolution through time of the jitter and the packet loss. This view can help you understand MOS variations and see which metric is impacting MOS.



- By pointing a specific point of time on the graph, you can display the exact value for each metric on the right side of the graph.
- By clicking on a specific point of time, you are directly to the *VoIP* conversations for this point of time.



VoIP Bandwidth & Call Volume

This view shows a chart of:

- bandwidth used for voice and signaling for the first one.




Figure 5.5: VoIP Bandwidth Chart

- the evolution of the volume of calls through time. Calls are distributed between successful and unsuccessful calls. Successful calls are conversations where some voice was exchanged; unsuccessful calls are conversations without any voice exchanged.

VoIP Conversations & Details

The two last views show each call individually with some *usage* metrics for *VoIP Conversations*. The *VoIP Details* view is the same table but with *performance* metrics.




Figure 5.6: VoIP Calls Volume

Symc	Begin Time	Call Duration	Caller	Called State	Called	Called State	Application	MOS	Packet loss	Tranfer	Code	Last Call State
1	2011-04-07 15:44:51	00:01:00	101 P.BLANCARD	Private (callback)	83208	ra	sp	-	0.00 %	2.24KB	200	closed
2	2011-04-07 15:44:50	00:01:00	101 system	ra	000000000000	ra	sp	-	0.00 %	1550bytes	200	closed
3	2011-04-07 15:44:50	00:01:00	101 system	ra	02700	ra	sp	-	0.00 %	3.24KB	200	closed
4	2011-04-07 15:44:49	00:01:00	101 system	Private (callback)	83052	Private (callback)	sp	-	0.00 %	7.34KB	200	closed
5	2011-04-07 15:44:49	00:01:00	101 system	Private (callback)	84700	Private (callback)	sp	-	0.00 %	7.34KB	200	closed
6	2011-04-07 15:44:41	00:01:00	101 P.MAGNI	ra	83182	Private (callback)	sp	-	0.00 %	5.60KB	200	closed
7	2011-04-07 15:44:31	00:01:00	101 LYON	Private (callback)	83784	Private (callback)	sp	-	0.00 %	1.04KB	200	voice
8	2011-04-07 15:44:23	00:01:00	101 LYON	Private (callback)	83784	Private (callback)	sp	-	0.00 %	3.16KB	200	voice
9	2011-04-07 15:44:23	00:01:00	101 system	Private (callback)	000000000000	Private (callback)	sp	-	0.00 %	3.16KB	200	closed
10	2011-04-07 15:44:23	00:01:00	101 system	Private (callback)	000000000000	Private (callback)	sp	-	0.00 %	3.16KB	200	closed
11	2011-04-07 15:44:11	00:01:00	101 P.MEURIS	Private (callback)	000000000000	Private (callback)	sp	-	0.00 %	4.36KB	200	closed
12	2011-04-07 15:44:08	00:01:00	101 system	ra	000000000000	ra	sp	-	0.00 %	5770bytes	200	closed
13	2011-04-07 15:44:08	00:01:00	101 LYON	ra	83961	ra	sp	-	0.00 %	3.24KB	404	not_req
14	2011-04-07 15:44:00	00:01:00	101 A.CRUZ	Private (callback)	83404	ra	sp	-	0.00 %	81.14KB	180	not_req
15	2011-04-07 15:44:00	00:01:00	101 A.CRUZ	Private (callback)	83404	Private (callback)	sp	-	0.00 %	5.16KB	200	voice
16	2011-04-07 15:43:48	00:01:00	101 LYON	Private (callback)	83010	ra	sp	-	0.00 %	7099bytes	200	closed
17	2011-04-07 15:43:48	00:01:00	101 system	Private (callback)	83115	Private (callback)	sp	-	0.00 %	4.64KB	200	closed
18	2011-04-07 15:43:45	00:01:00	2nd 9e P.VIGOURDUX	Private (callback)	83127	Private (callback)	sp	4.39	0.00 %	605.34KB	200	voice
19	2011-04-07 15:43:45	00:01:00	101 P.VIGOURDUX	Private (callback)	83127	Private (callback)	sp	-	0.00 %	4.94KB	200	voice
20	2011-04-07 15:43:43	00:01:00	101 H.METZ	ra	83012	Private (callback)	sp	-	0.00 %	3.66KB	200	closed
21	2011-04-07 15:43:39	00:01:00	101 P.ADOCH	ra	83078	Private (callback)	sp	-	0.00 %	3.66KB	200	closed

Figure 5.7: VoIP Calls

5.4 Application dashboards

Dashboards are a report fitting on a single screen that put together all relevant information to understand how the application is doing. They are present in APS from version 1.7.

Note: Those dashboards are not available in *Securactive NPS*.

It is extremely useful:

- as a starting point for troubleshooting,
- as a tool to communicate to management and business users on how the application is actually performing.

It is a set of three elements that display key information on the performance of a business application.

5.4.1 How can it help?

For reporting

In a single report you have enough to explain a business user or a manager how the application performance went through time, which servers were doing worse and which zones were impacted. On top of the EURT, all this is based on three synthetic metrics that are easy to explain, so that you can address non-technically aware people with an understandable speech about “*what is going on*”:

- RTT – network performance
- SRT – Server Performance
- DTT – Delivery of application response through the network.

For troubleshooting

For network administrators this report brings together all the information about a business application required to:

- validate whether there is a slowdown or not
- identify the origin of a slowdown (network, application, response delivery)
- which users or servers were impacted

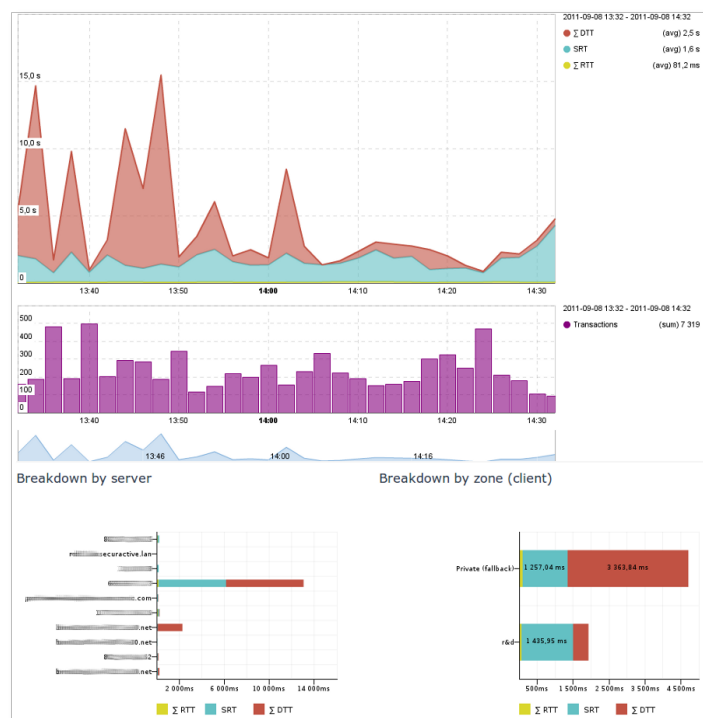


Figure 5.8: Overall view of the application dashboard

In no more than one click, you can conclude on whether there was a slowdown or not, what was the origin of the degradation, which client zones were impacted. With a single additional click (i.e. two clicks in total!), you can view whether all clients in a *zone* were impacted or if the server response time degradation was due to another *application* hosted on the same server machine.

5.4.2 Components

1s element: the evolution of End User Response Time through time




Figure 5.9: End User Response Time (EURT) graph

This EURT graph shows:

- the evolution of the quality of experience for users of this application over the period of time,
- the number of transactions help you consider the evolution of EURT with rigor and common sense (you would not consider a degradation of *EU Response Time* for 10 applicative transactions in the same way as for 10 000).

The breakdown of EURT in three intelligible components (RTT for network latency, SRT for *Server Response Time* and DTT for *Data Transfer Time*) let you know at first glance what is the origin of the possible performance degradation. For example in the screenshot here-above, we can observe an increase in the SRT; the network and the time required to send the response to the client have not increased. Either the server overall responded slower or some specific queries required a much larger treatment time (you can determine this by drilling down to that specific point of time).

2nd element: EURT by Server





Figure 5.10: EURT by server

What we can see here, is a comparison between the EURT for that application on each server that provides this *application*. In this case, it is obvious that *Atlantis* tend to respond much slower than *Brax*. By clicking on it having a looking at a second dashboard called *Server/Application Dashboard*, we shall be able to determine if this permanent or punctual and whether this due to the load on this application or on another one hosted on the same server.

3rd element: EURT by Client zone



Client zone / application dashboard

You can access this dashboard either through the menu or by clicking on a specific client zone in the *Application Dashboard*. This dashboard contains three bits of information:

- EURT graph through time for this client *zone* and this *application*.
- *EURT breakdown by server* (so that you can compare the performance offered by different servers to that client zone).
- *EURT per client* (so that you can identify whether all clients are impacted by a slowdown, or which individual client generates more volume or has worse application performance).

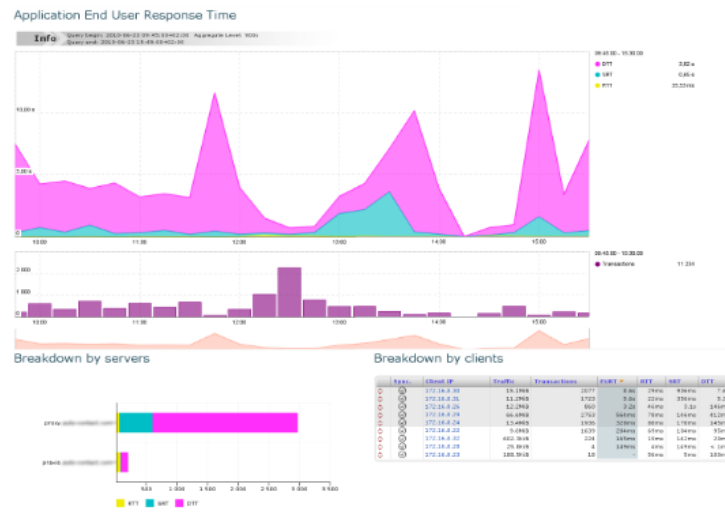


Figure 5.12: Client zone / application dashboard

The breakdown by client is interesting to know whether all the zone was impacted or just some individual users and on which component of the EURT (network latency, server response time or data transfer time and for which number of transaction and amount of traffic).

	Sync.	Client IP	Traffic	Transactions	EURT	RTT	SRT	DTT
⬇	⬇	172.16.8.30	15.1MiB	2677	6.6s	29ms	906ms	7.6s
⬇	⬇	172.16.8.31	11.2MiB	1723	5.6s	22ms	356ms	5.3s
⬇	⬇	172.16.8.26	12.2MiB	860	3.2s	46ms	3.1s	146ms
⬇	⬇	172.16.8.29	66.8MiB	2753	564ms	78ms	146ms	432ms
⬇	⬇	172.16.8.24	13.4MiB	1936	328ms	88ms	178ms	145ms
⬇	⬇	172.16.8.22	6.3MiB	1639	204ms	69ms	104ms	95ms
⬇	⬇	172.16.8.22	6.07.3KiB	274	1.65ms	19ms	145ms	20ms
⬇	⬇	172.16.8.28	25.8KiB	4	1.46ms	4ms	145ms	< 1ms
⬇	⬇	172.16.8.23	188.5KiB	18	-	56ms	5ms	100ms

Figure 5.13: Breakdown by client

Server / application dashboard

You can access this dashboard either through the menu or by clicking on a specific server in the *Application Dashboard*. This dashboard contains three bits of information:

- EURT graph through time for this server and this application
- *EURT breakdown by client zone* (so that you can compare the performance offered to different client zone from that server)
- Comparison with other applications provided by that server (so that you can identify whether a peak of transactions on another application is impacting the performance of that application, and see the volume of data, transactions and performance metrics for all applications provided by this server).

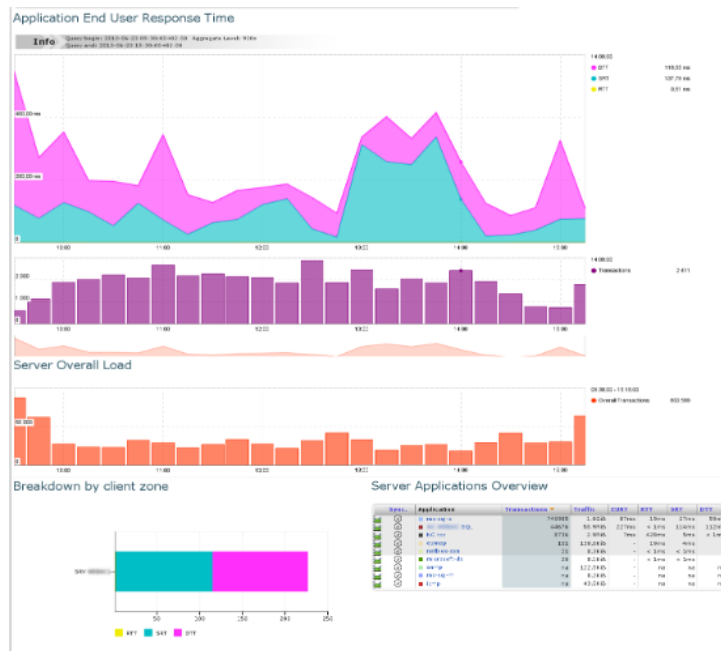


Figure 5.14: Server / Application Dashboard

5.4.4 Interactions

Dashboards have been developed so that a single click drives on more detailed information on the object you are most interested in:

- If you click on the EURT graph in any of these three dashboards, you make a focus on a shorter period of time (for example a SRT peak – depending on the aggregation level you either reach a lower aggregation level for a shorter period or the corresponding performance conversations, see [Data Aggregation](#) (page 11)). At the same time you will get the server and zone breakdown for that more specific period of time.
- If you click on a server, you reach the Server / application dashboard.
- If you click on a client zone, you reach the Client zone / application dashboard.

5.5 SPV Comparison tables

5.5.1 Objectives

SecurActive SPV presents performance metrics in the form of comparison tables to make it easier for network managers to compare the performance depending on where users are located. This feature provides an easier way to:

- compare application performance between client zones
- locate the furthest zones from a network latency stand point
- isolate communications where retransmission are impacting service delivery
- ...

5.5.2 Network Performance Table

This table provide an easy way to compare network performance between zones:

- RTT Server

- RTT Client
- RR Server
- RR Client
- RD Server
- RD Client

As many reports in SecurActive SPV, the navigation through the data is based on a drill-down mechanism, which makes it possible to go from a very wide view of network performance throughout the network to a focus on the communication between 2 zones, and then down to the detailed conversations.

Network Performance Client / Server table (page 54) and *Network Performance Client / Server table (unfolded)* (page 54) are two illustrations of this benefit:

- In this table, we have identified the two zones, between which the network response time was the highest.

Action	Client Zone	Server Zone	Traffic	CT	RTT	RR	RD
All	All	Internet	577.8 MB	68ms	67ms	0.21 %	326ms
Paris LAN	Paris LAN	Internet	577.8 MB	68ms	67ms	0.21 %	326ms
VLAN R&D	VLAN R&D	Internet	387.7 MB	63ms	63ms	0.08 %	892ms
VLAN Sales	VLAN Sales	Internet	190.1 MB	77ms	74ms	0.41 %	147ms
VLAN Top	VLAN Top	Internet	3.5 KB	na	na	< 0.01 %	na
VLAN VWR	VLAN VWR	Internet	1.3 KB	na	na	< 0.01 %	na
VLAN Labo	VLAN Labo	Internet	0 Bytes	na	na	na	na
MPLS Danton	MPLS Danton	Internet	0 Bytes	na	na	na	na
VPN Danton	VPN Danton	Internet	0 Bytes	na	na	na	na
Zone IPv6	Zone IPv6	Internet	0 Bytes	na	na	na	na
Paris LAN (failback)	Paris LAN (failback)	Internet	0 Bytes	na	na	na	na
Internet	Internet	Internet	684 Bytes	na	na	< 0.01 %	na
Extranet-SP	Extranet-SP	Internet	0 Bytes	na	na	na	na
Trip-ESG	Trip-ESG	Internet	0 Bytes	na	na	na	na

Figure 5.15: Network Performance Client / Server table

- In this table, we have identified the two zones, between which we could observe the highest retransmission rate.

Action	Client Zone	Server Zone	Traffic	CT	RTT	RR	RD
All	All	Internet	577.8 MB	68ms	67ms	0.21 %	326ms
Paris LAN	Paris LAN	Internet	577.8 MB	68ms	67ms	0.21 %	326ms
VLAN R&D	VLAN R&D	Internet	387.7 MB	63ms	63ms	0.08 %	892ms
VLAN Sales	VLAN Sales	Internet	190.1 MB	77ms	74ms	0.41 %	147ms
Policy	Policy	Internet	190.1 MB	77ms	74ms	0.38 %	146ms
VLAN Sales (failback)	VLAN Sales (failback)	Internet	0 Bytes	na	na	0.89 %	59ms
Printer HP4100	Printer HP4100	Internet	0 Bytes	na	na	na	na
Printer HP2040	Printer HP2040	Internet	0 Bytes	na	na	na	na
Printer HP4250	Printer HP4250	Internet	0 Bytes	na	na	na	na
Server Production	Server Production	Internet	0 Bytes	na	na	na	na
VLAN Top	VLAN Top	Internet	3.5 KB	na	na	< 0.01 %	na
VLAN VWR	VLAN VWR	Internet	1.3 KB	na	na	< 0.01 %	na
VLAN Labo	VLAN Labo	Internet	0 Bytes	na	na	na	na
MPLS Danton	MPLS Danton	Internet	0 Bytes	na	na	na	na
VPN Danton	VPN Danton	Internet	0 Bytes	na	na	na	na
Zone IPv6	Zone IPv6	Internet	0 Bytes	na	na	na	na
Paris LAN (failback)	Paris LAN (failback)	Internet	0 Bytes	na	na	na	na
Internet	Internet	Internet	684 Bytes	na	na	< 0.01 %	na
Extranet-SP	Extranet-SP	Internet	0 Bytes	na	na	na	na
Trip-ESG	Trip-ESG	Internet	0 Bytes	na	na	na	na

Figure 5.16: Network Performance Client / Server table (unfolded)

5.5.3 Application Performance Table

Securactive APS This table provide an easy way to compare application performance between zones:

- EURT
- RTT Server
- RTT Client
- SRT
- DTT

As many reports in SecurActive SPV, the navigation through the data is based on a drill-down mechanism, which makes it possible to go from a very wide view of application performance to a focus on a specific client or server zone, and then down to the detailed conversations. In the illustration *Application Performance Client / Server table* (page 55) you can see the Application performance between client zones for the application 'Salesforce':

Application Performance Client / Server Table - Traffic volume and application performance by application, client zone and server zone

Begin: 2011-04-14 13:22 End: 2011-04-14 14:22 More...

Client Zone: All Application: CRM-SalesForce Server Zone: Extranet-SF Search

Info: Begin: 2011-04-14 13:22:09+02:00 End: 2011-04-14 14:22:00+02:00 Aggregate Level: 2 minutes

Action	Client Zone	Server Zone	Traffic	EURT	RTT	SRT	DTT
+	All	Extranet-SF	1.4 MB	672ms	161ms	370ms	140ms
-	Paris LAM	Extranet-SF	1.4 MB	672ms	161ms	370ms	140ms
-	VLAN Sales	Extranet-SF	1006.3 KB	581ms	177ms	350ms	53ms
-	VLAN R&D	Extranet-SF	449.2 KB	903ms	115ms	425ms	362ms
-	VLAN Labo	Extranet-SF	0 Bytes	na	na	na	na
-	VLAN Web	Extranet-SF	0 Bytes	na	na	na	na
-	VLAN Tool	Extranet-SF	0 Bytes	na	na	na	na
-	WPLD Berlin	Extranet-SF	0 Bytes	na	na	na	na
-	VPN Darwin	Extranet-SF	0 Bytes	na	na	na	na
-	Zone IP4	Extranet-SF	0 Bytes	na	na	na	na
-	Paris LAM (fallback)	Extranet-SF	0 Bytes	na	na	na	na
-	Extranet-SF	Extranet-SF	0 Bytes	na	na	na	na
-	Tout - BIG	Extranet-SF	0 Bytes	na	na	na	na
-	Internet	Extranet-SF	0 Bytes	na	na	na	na

Figure 5.17: Application Performance Client / Server table

5.6 TCP Errors / Events

5.6.1 Objectives

These two tables expose to the user many TCP statistics in order to reveal dysfunctions or unusual events.

5.6.2 TCP Errors

For each TCP conversation the following fields are displayed:

- RD Server/Client
- Duplicate acks
- number of SYNs
- number of handshakes
- number of session ends
- number of FINs from client
- number of FINs from server
- number of RSTs from client
- number of RSTs from server
- number of timeouts

By sorting on the RD or duplicate ack fields one can quickly check the worst conversations in term of TCP performance. Also, number of reset packets are usually noteworthy. One can then jump to the IP summary page of either the client or the server (depending on who is to blame) to gather further data on this event.

5.6.3 TCP Events

This page does not focus explicitly on TCP errors but aims at giving various overall statistics about each TCP conversation, in order first to give an accurate view of the actual traffic in term of payload and number of connections, and second to notice unexpected patterns.

This page can also serve as a way to find which conversations are important/relevant and thus which zone / application could be split to help distinguish more closely between significant flows.

For each TCP conversation the following fields are displayed :

- payload
- number of packets
- number of handshakes

- number of timeouts
- number of RSTs from client
- number of RSTs from server
- number of FINs from client
- number of FINs from server

5.7 Packet level analysis

5.7.1 Objectives

Once you have identified the origin of an issue, you may want to analyze it further by looking at the packets themselves. You have two ways to realize this:

- Manual capture through Pulsar's `tcpdump` command
- Automatic packet capture

5.7.2 Manual packet capture

By connecting through Pulsar, you can start a manual capture of any traffic viewed on the interface of your device. To do so, you need to go through 3 steps:

1. Connect to Pulsar (see *Pulsar* (page 26))
2. Enter the command to launch the trace: for example, `tcpdump -i <interface> host <host_ip>`.
3. Enter `Control+C` to stop the trace.

Note:

- you can access a help by entering `help tcpdump`.
 - you can refer to `tcpdump` command (<http://www.tcpdump.org>). Please, have a look at the online [manual](http://www.tcpdump.org/tcpdump_man.html) (http://www.tcpdump.org/tcpdump_man.html).
 - all parameters are available except the `-w`.
-

Accessing the tracefile

You will not be able to view the trace through Pulsar; to access the PCAP file, you should connect to the probe via FTP, using a FTP client and the Pulsar admin user (see *Pulsar* (page 26)).

5.7.3 Automated packet capture (“AutoPCAP”)

Principles

Performance Vision can capture packets automatically, in case abnormal values are observed on critical servers. These packets are presented for later analysis as *PCAP* files, which can be downloaded through the web graphical interface at the conversation level.

Applications

These files are presented in the following views:

- Conversations
- DNS messages
- VOIP details

In each of these views, a column at the right end of the table indicates *PCAP*; a small icon indicates that packets have been captured for a given conversation or not. If the *PCAP* file is available, you can download it by clicking on the icon. Once the file has been downloaded, you can view the packets with any protocol decoder (capable of reading *PCAP* files).

Conn. established	Transactions	EURT	RTT in	RTT out	SRT	DTT cil	DTT srv.	CT	CRT	RR in	RR out	TTFB	Pcap
0	48				< 1ms	1 m 15 s	< 1ms	< 1ms	< 1ms	-	-	-	
1	66	55.3 s	210 ms	< 1ms	55.1 s	< 1ms	< 1ms	187 ms	95 ms	-	-	55.4 s	
1	20	23.2 s	81 ms	< 1ms	21.1 s	794 ms	1.3 s	74 ms	76 ms	-	-	151 ms	
0	40		97 ms	-	17.5 s	38.4 s	1.5 s	-	11.4 s	-	-	-	
2	1	3.6 s	179 ms	< 1ms	3.4 s	< 1ms	2 ms	188 ms	-	-	10.00 %	3.6 s	
4	3	913 ms	126 ms	< 1ms	743 ms	< 1ms	43 ms	123 ms	-	-	-	1.4 s	
1	2	891 ms	63 ms	< 1ms	628 ms	< 1ms	198 ms	51 ms	2.1 s	-	-	1.2 s	
10	31	778 ms	71 ms	< 1ms	621 ms	64 ms	20 ms	57 ms	16.3 s	25.43 %	0.91 %	436 ms	
2	7	549 ms	50 ms	< 1ms	472 ms	< 1ms	26 ms	41 ms	10.7 s	-	-	336 ms	
6	9	961 ms	222 ms	< 1ms	450 ms	< 1ms	288 ms	220 ms	570 ms	-	-	790 ms	

Figure 5.18: PCAP column in Performance conversations

192.168.25.254	Private (fallback)	45	7.1 KB	mail.google.com	6 ms	A	NoError	
192.168.10.254	R&D	44	5.7 KB	pyip.mirrors.rd.securactive.lan	< 1ms	A	NoError	
192.168.10.254	R&D	44	5.9 KB	pyip.mirrors.rd.securactive.lan.secur...	< 1ms	AAAA	NoDomain	
192.168.10.254	R&D	44	5.2 KB	pyip.mirrors.rd.securactive.lan	< 1ms	AAAA	NoError	
192.168.10.254	R&D	44	5.8 KB	pyip.mirrors.rd.securactive.lan.aps.secu...	< 1ms	AAAA	NoDomain	
192.168.10.254	R&D	44	6.1 KB	pyip.mirrors.rd.securactive.lan.aps.secu...	< 1ms	AAAA	NoDomain	
192.168.10.254	R&D	44	6.0 KB	pyip.mirrors.rd.securactive.lan.rd.secur...	< 1ms	AAAA	NoDomain	
192.168.10.254	R&D	44	8.1 KB	mail.google.com	17 ms	A	NoError	
192.168.10.5	R&D	42	7.2 KB	No request name	-	Unknown DNS type 0	NoError	
192.168.60.255	Private (fallback)	36	3.2 KB	WORKGROUP	-	NB	NoError	

Figure 5.19: PCAP column in DNS messages

Callee	Callee Zone	Application	MOS	Packet loss	Server sign. traffic	Voice traffic	Code	Last Call State	Pcap
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	1.05 %	3.1 KB	1.2 MB	200	closed	
aainr1@172.25.51.150	na	■ NC udp	-	0.00 %	989 Bytes	0 Bytes	200	closed	
aainr1@172.25.51.150	Internet	■ NC udp	-	0.19 %	6.9 KB	2.3 MB	200	closed	
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	0.00 %	3.1 KB	1005.8 KB	200	closed	
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	5.11 %	2.6 KB	241.6 KB	200	closed	
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	0.00 %	3.0 KB	267.0 KB	200	closed	
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	0.00 %	3.0 KB	891.4 KB	200	closed	
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	0.00 %	2.6 KB	195.0 KB	200	closed	
aainr1@172.25.51.139	Private (fallback)	■ NC udp	-	0.00 %	6.3 KB	0 Bytes	250	closed	
aainr1@172.25.51.150	Private (fallback)	■ NC udp	-	0.00 %	3.4 KB	515.9 KB	200	closed	

Figure 5.20: PCAP in VOIP details

For instance, if you are using *Wireshark* to decrypt the packets, you can directly view the packets.

To view the query and the beginning of the response, you can use the feature *Follow TCP stream* (in the Analysis menu).

Conditions

Packets are saved by Performance Vision, as soon as the conversation they belong to matches a certain number of conditions:

- One of the hosts (either client or server, for whichever protocol) is a server for one of the Business Critical Applications
- One of the following metrics is considered as out of the norm:
 - Server Response Time (SRT) for TCP flows
 - Retransmission Rate
 - DNS Response Time

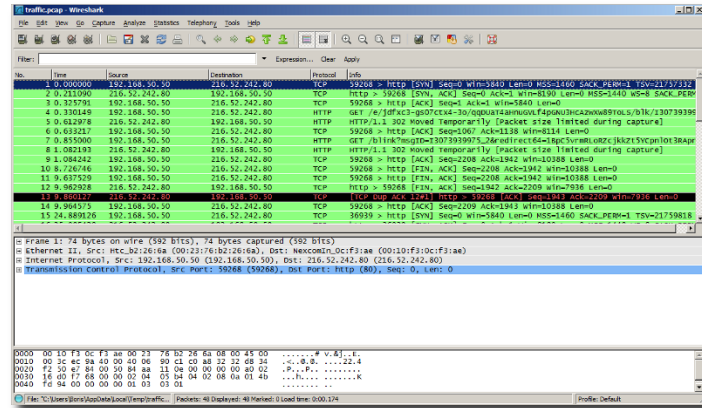


Figure 5.21: Viewing packets in Wireshark

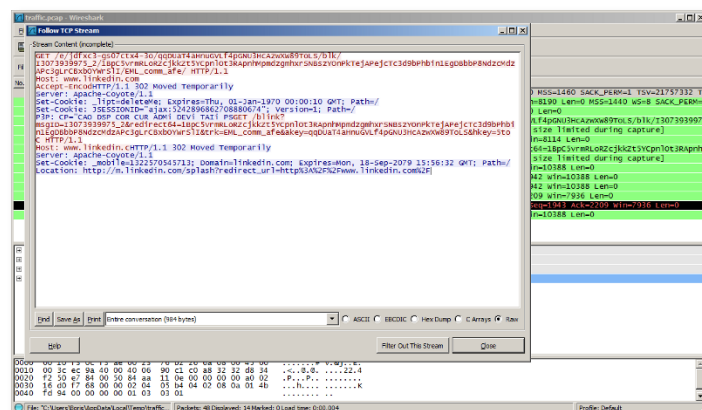


Figure 5.22: Viewing query and response

Note: *PCAP* files are a **sample** of the conversation. If you request on a one hour interval and get a *PCAP* file, the *PCAP* will not contain one hour of data but only the data which match the above conditions.

Limitations

The Automatic Packet Capture feature works under a certain number of conditions to ensure the proper execution of other services provided by Performance Vision. Among these necessary limitations, you need to observe the following:

- The retention of *PCAP* files is limited by the disk space allocated for captures; in the current version, this space is limited to 10GB (for both manual and automatic captures). When all 10GB are used, no new *PCAP* file is saved.
- The maximum retention time for Automatic captures is set to 48 hours; after this delay, Automatic *PCAP* will be deleted.
- The sniffer component of Performance Vision cannot forge more than 5 000 *PCAP* files simultaneously; if more than 5 000 conversations are hold on Business Critical Applications servers at once, some conversations will not be recorded at packet level.

Please note that the threshold values and voluntary limitations will be reviewed in newer versions in the light of our experience and the customer feedback we will receive. Please note that if you need an exhaustive trace of a given set of conversations, you can also use the manual capture feature available through Pulsar.

5.8 Interpretation Guidelines

The objective of this section is to help our customers to make the best use of the performance reports provided by their appliance. You will find enclosed a brief overview of how application performance issues can be solved with SPV. This first section focuses on synthetic metrics to produce a measure of the quality of experience of users (QoS - End User Response Time) and give you a simple explanatory framework to understand the cause of application slowdowns (Round Trip Time, Server Response Time and Data Transfer Time).

Note: Some metrics and views described below are only available in *Securactive APS*.

5.8.1 Objectives

Before you start analyzing performance reports, there is a certain number of elements which you must bear in mind: **Performance metrics should not be considered as absolute values, but in comparison with different time intervals, servers and user groups.** Performance metrics represent time interval. Although most of them correspond to the measurement of a concrete phenomenon, it is almost impossible to provide a scale of what is a good or a bad response time, with no experience of the impact it has on users. For example, indicating that the Network Round Trip Time from a site A to a site B is 200ms does not mean you have a measure which is acceptable or not. In the same way, a Server Response Time (SRT) of an application A of 100ms may be very “bad” when the same value would be excellent for an application B. As a consequence, it is important to consider performance metrics as relative values; one of the key to a good interpretation of performance metrics is to compare systematically performance metric value:

- to another time period,
- to another users group.

Mixing up performance metrics for several applications does not make sense. When looking at application performance metrics, you should be very careful of isolating applications for analysis. As a consequence the metrics which very much depend on the application’s specific behaviour should not be considered altogether: this

is true for metrics such as EURT (End User Response Time), SRT (Server Response Time) and DTT (Data Transfer Time).

RTT measurements can marginally be impacted by the behaviour of the operating system. Network Round Trip Times for TCP are based the TCP acknowledgment mechanism. This means that, although RTT is generally a good measurement of round trip latency, if the operating system of one of the parties is so overloaded that the acknowledgment process becomes slower, RTT values will be impacted. RTT Server would be impacted on the server side and RTT Client on the client side. RTT should then be analyzed in parallel to CT (Connection Time) - because the treatment of new session by the IP stack has a higher priority).

Some values are averaged measures. For each conversation, two kinds of values are reported:

- counters, for instance packets or byte counters, which are the sum over all connections aggregated for this conversation;
- performance metrics, for instance RTT, SRT, DTT and the likes, which are average values over all samples aggregated for this conversation.

EURT

EURT stands for End User Response Time.

This metric is an aggregate of various other measures meant to give an idea of the perceived overall end user experience. It is taken as the sum of RTT, SRT and DTT.

EURT has no meaningful physical counterpart. Only its evolution makes sense, and allow the system administrator to check at a glance whether a network *zone* is behaving as usual or not. Notice that expected correct values for both SRT and DTT depend on the protocol at hand. As a consequence you should not try to compare two EURT of different applications.

RTT

RTT stands for Round Trip Time.

RTT gives an approximation of the time required for a packet to reach its destination, and can be further decomposed into a RTT Server (delay between a data packet send by the client and its ACK from the server) and a RTT Client (in the other way around). As a typical IP implementation will delay acknowledging of incoming data, additional tricks are exploited in order to rule out these software biases :

- make use of SYN/FIN acknowledgment and some exceptional conditions such as TCP resets, that suffer no such delays, to estimate a realistic upper bound.
- exclude unusually high RTT values.
- bound RTT Server/Client by SRT/CRT if RTT sample set looks suspicious.

RTT is meaningful of the bare speed of the physical layer. It is unaffected by packet retransmissions, packet loss or similar occurrences. RTT may be affected by (from most common to the rarest):

- Slow network equipment between client and server (such as a router or a switch);
- Link layer overloaded (ethernet collisions for instance);
- Malfunction of one of the involved network adapter.

These troubles should be further investigated by comparison with other client and/or server zones in order to locate the misbehaving equipment. Notice that a degradation of RTT will almost invariably impact other metrics as well.

SRT

SRT stands for Server Response Time.

SRT gives an estimation of the elapsed time between the last packet of an applicative request and the first packet of the server's response.

SRT represents the processing time of the server, at the application layer, for a given request. SRT may be affected by (from the most common to the rarest):

- Time greedy application request (a complex SQL command can let the server processes during many seconds);
- Application layer overloaded (too many requests, such that the server can't handle all of them in a small period of time);
- Marginally SRT can be affected by the increase of network latency between the point of capture and the server (parallel increase of the RTT `Server` value);

To pinpoint the root cause of the slowdown, we firstly want to compare the SRT for a given couple server/application to other applications on the very same server. If there is a blatant difference, the application is guilty. Otherwise, we want to compare it to other servers in the same zone, then different zones.

DTT

DTT stands for Data Transfer Time.

DTT `server` is defined as the time between the first data packet of the response (with ACK flag and a non null payload) from the server and the last packet considered as part of the same response (if the packet has the same acknowledgement number); FIN, RST packets from server or client will also be considered as closing the sequence. A Timeout will cancel a DTT. Note that if the answer is small enough to be contained in only one packet, the DTT will be of '0'.

DTT `client` is the same metric in the other direction.

DTT (sum of both server and client DTT) is meaningful of the time the user is going to have to wait for the response to circulate on the network from the server to the client. It is not dependent on the Server Response Time (e.g. a DTT might be short for a long SRT):

- the request might require a large calculation, but the result represents a small volume of data; or a DTT might be very large, but SRT very short because the request is easy to handle but the response is very large). DTT depends on (from the largest impact to the smaller):
- the size of the response (the more data is contains the longer it takes to transfer it),
- the level of retransmission (the more packets are retransmitted, the longer it will take to transfer the whole response),
- the network latency (the longer it take to transfer packets through the network, the longer it will be to transfer the response - minor impact),
- the actual throughput which can be reached to transfer the response from the server to the client.

DTT may vary (for most common to a the rarest):

- globally or not on a per transaction basis (if only for some transaction, it may be linked to the size of some specific application response),
- for all client zones or for some only (if for some client zones only, it may be linked to specific network conditions — retransmissions),
- for all servers or for some server (if for a specific server, it may be due to a specific server issue in broadcasting the response).

5.8.2 Scenario guidelines

Slow site connection

Hypothesis:

One or several end users complain about a slow access to all applications (both in and out the LAN).

Diagnosis:

You will find in this section the classical informations to grab in order to diagnose the issue:

- is the application really slower for this site? You can get this information from several places:




Figure 5.23: Zone comparison in the Application Performance Dashboard

Action	Client Zone	Server Zone	Traffic	EURT	Σ RTT	SRT	DTT
✕	All	All	2.9 GiB	1.2s	34ms	907ms	297ms
✕	Private	All	2.9 GiB	1.2s	34ms	907ms	297ms
✕	R&D	All	2.8 GiB	891ms	20ms	566ms	304ms
✕	Sales	All	86.0 MiB	2.6s	85ms	2.2s	272ms
✕	Private (fallback)	All	10.4 KiB	na	na	na	na
✕	Labo	All	0 Bytes	na	na	na	na
✕	Voip	All	0 Bytes	na	na	na	na
✕	Internet	All	1.1 KiB	na	213ms	na	na

Figure 5.24: EURT comparison in the Application Performance Comparison Table

- Does the slowdown occur for a specific application? If so, check [Slow application](#) (page 65); You can achieve this through the client / server table.

Begin

2011-04-04 13:00

End

2011-04-04 14:00

More...

Previous report

Client Zone

R&D

Server Zone

Labo

Application

Add this page to a report

Search

Info

Begin: 2011-04-04 13:00:00+02:00

End: 2011-04-04 14:00:00+02:00

Aggregate Level: 2 minutes

Action	Client Zone	Server Zone	Applications	Traffic	EURT	SRT	CRT	Σ RTT	DTT	CT
✕	R&D	Labo	11 applications	162.2 MiB	na	18ms	1.2s	< 1ms	15ms	< 1ms
✕	R&D	Labo	http	74.9 MiB	6ms	5ms	139ms	< 1ms	< 1ms	< 1ms
✕	R&D	Labo	http-alt	29.6 MiB	79ms	72ms	na	< 1ms	7ms	< 1ms
✕	R&D	Labo	ssh	28.2 MiB	538ms	4ms	< 1ms	< 1ms	538ms	< 1ms
✕	R&D	Labo	ftp-passive	27.9 MiB	na	na	< 1ms	< 1ms	na	< 1ms
✕	R&D	Labo	git	1.1 MiB	4ms	3ms	5ms	< 1ms	< 1ms	< 1ms
✕	R&D	Labo	microsoft-ds	300.5 KiB	55ms	< 1ms	5.0s	< 1ms	36ms	na
✕	R&D	Labo	snmp	202.1 KiB	na	na	na	< 1ms	na	na
✕	R&D	Labo	NC	29.4 KiB	9ms	1ms	16.5s	< 1ms	7ms	na
✕	R&D	Labo	redisearch	16.6 KiB	na	na	na	< 1ms	na	na
✕	R&D	Labo	ftp	2.6 KiB	318ms	2ms	< 1ms	< 1ms	315ms	< 1ms
✕	R&D	Labo	netbios-ssn	1.6 KiB	na	na	na	na	na	na

Figure 5.25: EURT comparison between applications for a given zone (Client / server table)

- Does the slowdown occur for a specific server? If so, check [Slow server](#) (page 69);
- Did you upgrade the clients workstations recently? If so, it's a specific system issue, you may ask the System Administrator for more details;
- Did you upgrade your network equipment? If so, the router/switch configuration is probably involved;
- Now we might inspect deeply in the SPV dashboards. Check the **Monitoring -> Network Performance Chart**
- Do the **Retransmission Rate** and **Retransmission Delay** vary? If so, we might face a congestion issue. Take a look at the router's load, etc;




Figure 5.26: EURT comparison between servers in the Application Performance Dashboard

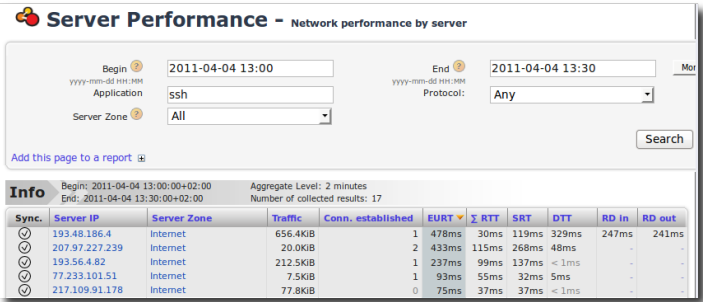


Figure 5.27: Server Response Time comparison through Server Performance.




Figure 5.28: Network Round Trip Time analysis




Figure 5.29: Retransmission analysis




Figure 5.30: Retransmission analysis

- The general slowdown for a client zone may also be the consequence of a crucial service: the DNS. Check out *DNS Response Time* (page 75);
- Look at the **Monitoring -> Bandwidth Chart**, to inspect the bandwidth variation, and the number of TCP /UDP flows as well.




Figure 5.31: Bandwidth charts

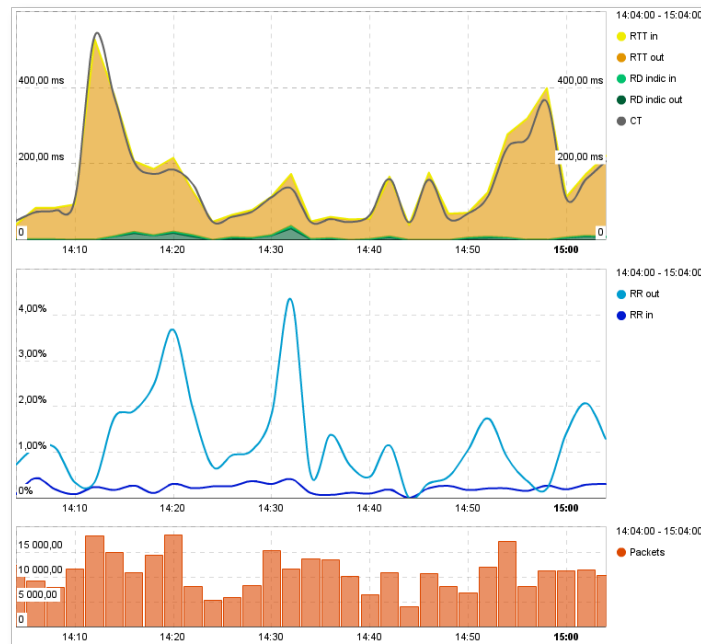


Figure 5.32: Impact of congestion on retransmissions and network latency or connection time

They might have overcome a QoS threshold, such that all the new application requests are blocked. A hint would be the increasing number of TCP RST packets. To be sure, you may take dive into the **Analysis -> TCP Errors** menu.

Slow application

Hypothesis

One or several end users complain about a slow access to a specific application : a fileserver.

Server IP	Application	Traffic	Packets	Conn. established	Num. timeout	Client RST	Server RST
172.16.1.10	PROXY	6.2MiB	10939	234	47	120	0
172.16.1.12	http-private	1.0MiB	1722	25	9	74	1
proxyauto-contact.c...	PROXY	19.3MiB	22795	220	32	73	0
proxyauto-contact.c...	PROXY	9.0MiB	15181	184	44	67	0
172.16.1.10	PROXY	9.3MiB	14615	244	42	61	0
172.16.1.10	PROXY	4.7MiB	8376	208	36	55	0
172.16.1.10	PROXY	11.6MiB	15668	138	26	44	0
172.16.1.10	PROXY	6.6MiB	13134	348	58	43	0
172.16.1.10	PROXY	2.6MiB	4361	99	21	43	0
172.16.1.10	PROXY	3.0MiB	4368	81	16	38	0

Figure 5.33: Number of RST packets sent from the TCP servers

Prerequisites

Zones have been configured to reflect the customer's network topology. The application `Samba_CIFS` has been identified. The traffic to the fileserver is mirrored to one of the listening interfaces of the probe. Where to start: a global view of the application performance!

1st example

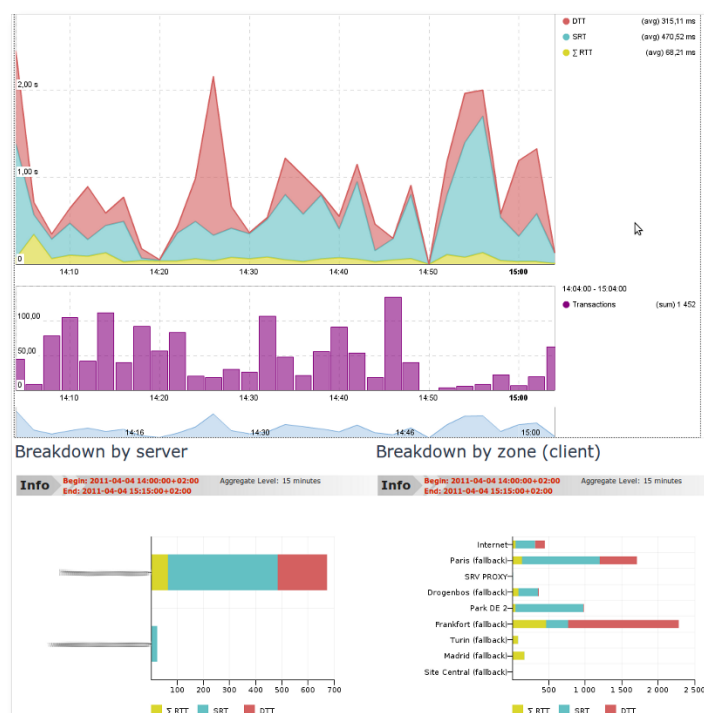


Figure 5.34: Peak in Server Response Time: application performance

Display the Application Dashboard for a relevant period of time. We can easily observe a peak in SRT from 6 to 18:15. From the breakdown by zone, we can easily conclude that only one zone has been impacted.

By clicking on that zone, we can see this client zone application dashboard:

From this, you can conclude that only one client (= user) was impacted. This issue was definitely due to a slow response of the server; it may be due to an application issue or a request which is specifically hard to respond to.

2nd example

Application Dashboard for a relevant period in the past (48 hours for example).

This dashboard shows in the upper part the evolution of the End User Response Time (EURT) through time for this fileserver.

- We can easily observe that the quality of experience of users accessing to this application got much worse yesterday afternoon.

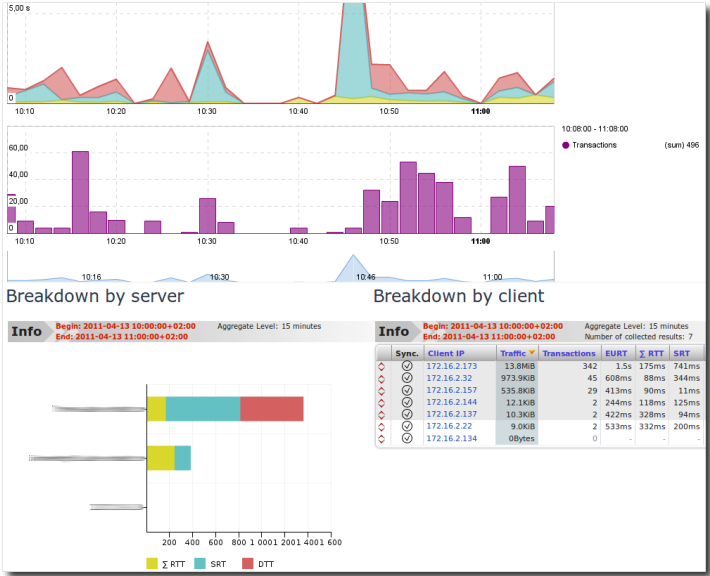


Figure 5.35: Peak in server response time: Application EURT

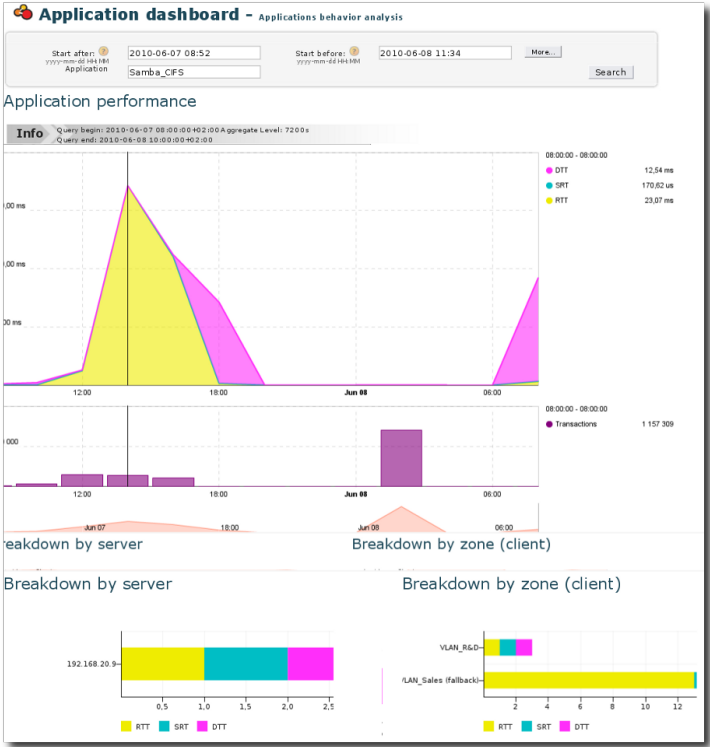


Figure 5.36: Peak in server response time: Application dashboard

- We can easily identify that this was due to a degradation of RTT (Round Trip Time - indicator of network latency) and not to the Server Response Time (SRT) or the Data Transfer Time (DTT).

From this graph, we can conclude that the server and the application are likely not to have any relationship with the slowdown. By looking at the two bar charts which show respectively the breakdown by server and by client zone, we can draw the following conclusions:

- This application is distributed by one server only (192.168.20.9)
- The EURT vary in large proportion between client zones, mainly because of RTT.
- *VLAN_Sales* has a much worse access to the application than *VLAN_R&D*, mainly because of the network latency.

Getting confirmation of our first conclusions. By clicking on the peak of EURT in the upper graph, we can narrow our observation period to understand better what happened at that point of time.



Figure 5.37: Peak of RTT in Application Dashboard

This confirms the following conclusions: RTT went up for the *VLAN_Sales* (only).

Understanding what is the perimeter of the slowdown

We now know that only *VLAN_Sales* was impacted by this slowdown, due to a longer network RTT. We therefore need to understand whether this was general (i.e. impacted all clients in the zone) or isolated to certain clients.

To achieve this, we can simply display the Performance conversations for the application *Samba_CIFS* for the zone *VLAN_Sales*. Here is the result:

From this screen, we can draw the following conclusion:

Only the clients 192.168.20.205 and 192.168.20.212 seem to be impacted. The other clients have very short RTT values.

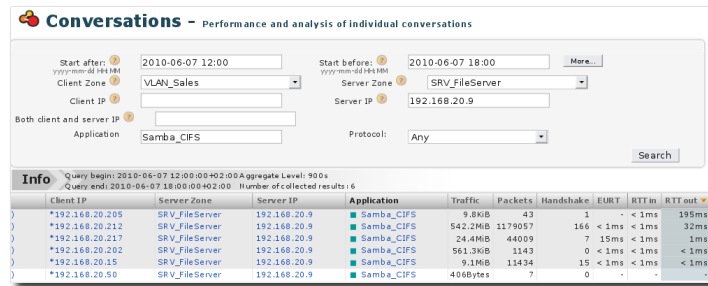


Figure 5.38: Peak in server response time: Conversations

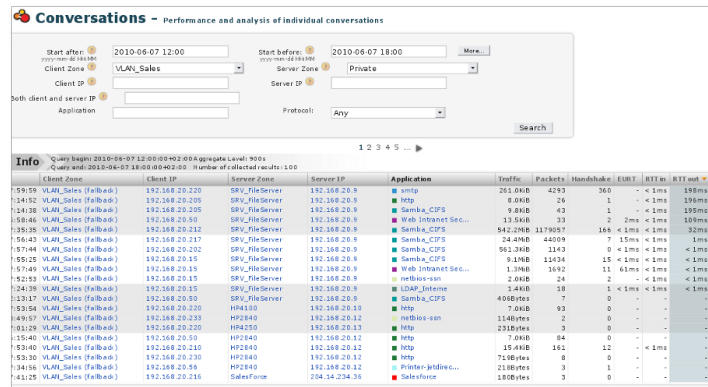


Figure 5.39: Peak in server response time: Conversations

To confirm this, we need to check that these two hosts are the only ones to be impacted and check whether they are impacted only when accessing to the Fileserver. To do that, we have a look at the Performance conversations between the *VLAN_Sales* and the *Private* zone. From this, we can draw the following conclusions:

- Not only 192.168.20.212 and 192.168.20.205, but also 192.168.20.220 and 192.168.20.50 are impacted.
- The *Samba_CIFS* (access to the fileserver) is not the only application impacted, but SMTP, HTTP and the *Web Intranet SecurActive*.

Actions to be taken after that analysis

- Check the windowing configuration on the operating system of these hosts (if high value, this is normal).
- Check the level of usage of the host (CPU, RAM usage).

Alternative scenarios:

- If we had seen some retransmission, check whether they are all on the same edge switch and check the interface configuration and media errors.

Slow server

Hypothesis:

Users complain about having to try several times to connect to a web-based application named “Salesforce”. The administrator suspects the application server hosting “Salesforce” is slow.

How to analyze the problem:

First, check to see if all applications on the application server hosting “Salesforce” are slow or if it is just the single web-based application “Salesforce” slow. If all applications are slow, then indeed, the application server may in fact be a slow server. If just the one web-based application “Salesforce” is slow, while the other applications (CRM) are responding quickly, the problem may be the application “Salesforce” and not a slow server.

To begin diagnosis, go to **“Monitoring”** -> **“Clt/Srv Table”**. Select the application server from the drop-down box labeled **“Server Zone”** and click **“Search”**.

- If we see that all applications on the server are responding slowly i.e. the SRT values are high for both **“Salesforce”** and **“CRM”**, the issue related to the server, not to applications.
- Second, check the Connection Time of the application server. If the connection times are high then this may also indicate a slow server.
- Third, check for retransmissions between the clients and the application server. If there are a lot of retransmissions then either the application server or a network device in between are dropping packets. Go to **“Monitoring”** -> **“Network performance chart”**. Select the application server **“Salesforce”** from the drop-down box labeled **“Server Zone”** and click **“Search”**.




Figure 5.40: Slow server: Network performance chart

Here we see that there is a high Retransmission Rate (RR Server) going from the clients to the application server. However, none of the packets from the server to the clients needed to be retransmitted (RR Client is around 0). This indicates that the application server is in fact dropping the packets and is therefore a slow server (Assuming that the route taken from the client to the server is the same route taken from the server to the client as is industry standard practice).

Lastly, check the *TCP errors* of the clients and the *Application server*. If the server reset count or number of timedout sessions is high, this is a further indication of a slow server. Go to **Analysis** -> **TCP errors**. Select the application server **“Salesforce”** from the drop-down box labeled **“Server Zone”** and click **Search**.

RD In	RD out	Dup ack	Conn. attempts	Conn. established	Sess. end	Client FIN	Server FIN	Client RST	Server RST	Num. timeout
334ms	333ms	48	330	165	452	165	0	122	330	2
263ms	346ms	58	330	165	463	165	0	133	330	1
215ms	273ms	37	328	164	446	164	0	118	328	0
338ms	365ms	37	18	9	25	9	0	7	18	0
-	-	5	15	15	20	16	0	5	15	3
-	301ms	2	10	5	12	5	0	2	10	0
-	1	1	10	5	10	5	0	0	10	0
-	310ms	1	10	5	10	5	0	0	10	0
< 1ms	-	0	10	5	10	5	0	0	10	0
-	300ms	1	10	5	10	5	0	2	8	1
-	-	4	7	7	7	7	0	0	7	4

Figure 5.41: Slow server: TCP Errors

Here we see that there are a lot of server resets and timeouts. Given all the above information, we can conclude that the application server is operating slowly. At this point, the server administrator should perform direct diagnosis on the application server to verify CPU, RAM and HD usage.

N-tier application performance issue

Hypothesis:

Users are complaining about slow response time from an in-house web application. This application being an N-tier architecture, its performance as seen by a client is tied to several parameters:

- DNS latency to resolve web server name from the client host (see **DNS Response Time**)
- Connection time to server
- Data Transfer Time between these hosts
- DNS latency to resolve other server names accessed from the web server (database servers for instance, cf. **DNS Response Time**)
- Connection and data transfer times between these hosts
- Server response time of these servers

Identification of the culprit:

First we need to find out if the experienced slowdown is due to the web front end itself. To this end, check every component of the EURT:

- If SRT is fast but RTT and/or DTT (see also Connection Time) then we are facing a network slowdown. Refer to previous sections of this guide to further track down the problem.
- If SRT is preponderant compared to DTT and RTT then the application itself is to blame. Proceed to find out what is affecting performance.
- Then check EURT between web server and each other involved servers (databases...)

If some of these EURT appear to be degraded then check recursively these other hosts. If not then check the web server load average.

5.8.3 Additional metrics

TCP anomalies

RST packets

A TCP connection is reset by a RST packet. There is no need to acknowledge such packet, the closure is immediate. A RST packet may have many meanings:

- If a TCP client tries to reach a server on a closed port, the server sends a RST packet. The connection attempt could be a malicious one (port scanning – nmap, etc), or the consequence of an unexpectedly down server, client/server misconfiguration, server restart, etc;
- A router might send a RST packet if the incoming TCP packet does not fit with the security policy (source range IP address is banned, the number of connection attempts is too high in a small period of time, etc);
- A QoS (Quality of Service) equipment limits the bandwidth (or the number of connections) by sending a RST packet to any new connection attempt;
- If a Intrusion Detection System (e.g. Snort) detects a malicious connection, he can send a RST packet to roughly close it;
- If a host between Client and Server wants to do a Denial of Service, it can reset the connection by sending RST to both peers. Basically it's the same mechanism than the previous one, but the motivation is quite different.

Retransmissions

One of the TCP metrics which is interesting to analyze is the retransmission. A TCP Retransmission is when a TCP packet is resent after having been either lost or damaged. Such retransmitted packet is identified thanks to its sequence number. In *SecurActive SPV* we do not consider packets with no payload, since duplicate ACKs are much more frequent, and not really characteristic of a network anomaly. There are several common sources of TCP retransmission:

- A network congestion. If a router can't cope with the whole traffic, its queue will grow bigger until it gets full and then start dropping the incoming packets. If you reach a predefined QoS limit, the exceeding packets will be dropped as well. Such drop will result in TCP retransmission. A common way to identify this kind of problem is by taking a glance to the traffic statistics. If you see a flat line at the max traffic allowed, then you get the root cause of retransmission. If the traffic graph looks OK, you can check over the load of the routers/switches you own (e.g. with the SNMP data). If the load is too high, you found the culprit.
- An overloaded server. Check the **Section Slow Server**.
- A hardware failure. Maybe a network equipment is simply down. It will obviously result in TCP retransmission until a new route is computed, or the issue fixed. This type of retransmission should occur with very short time effects and give some quite big peaks of retransmission, on very broad types of traffic on a specific subnet. If this happens often, it becomes important to find the faulty hardwares by tracking down which subnets are concerned.
- A packet header corruption. Network equipments are used to rewrite portions of packets (Ethernet source/destination, IP Checksum, maybe TOS field). A buggy firmware can result in corruption while rewriting protocol headers. In this case, the packet will probably be dropped within the network route. Even if it reaches the destination, the TCP/IP stack won't consider it as a valid packet for the current TCP sessions, and the stack will wait the correct packet. It will end in a TCP retransmission, anyway. This problem will likely occur on the same type of traffic and continuously.

ICMP

What is ICMP?

ICMP stands for *Internet Control Message Protocol* and is also a common IP transport protocol. It seems pretty explicit, although most people reduce ICMP to ping reply commands, a good way to test whether a host can be reached through a network and how much it takes for a packet to make a round trip through the network... Obviously ping and trace-route-like tools are very useful for network administrators... but there is much more to say about ICMP and the help it can provide for network administration & diagnosis. In total, ICMP can be used to send more than twenty types of control messages. Some are just messages, some others are a way for IP devices or routers to indicate the occurrence of an error.

Error messages

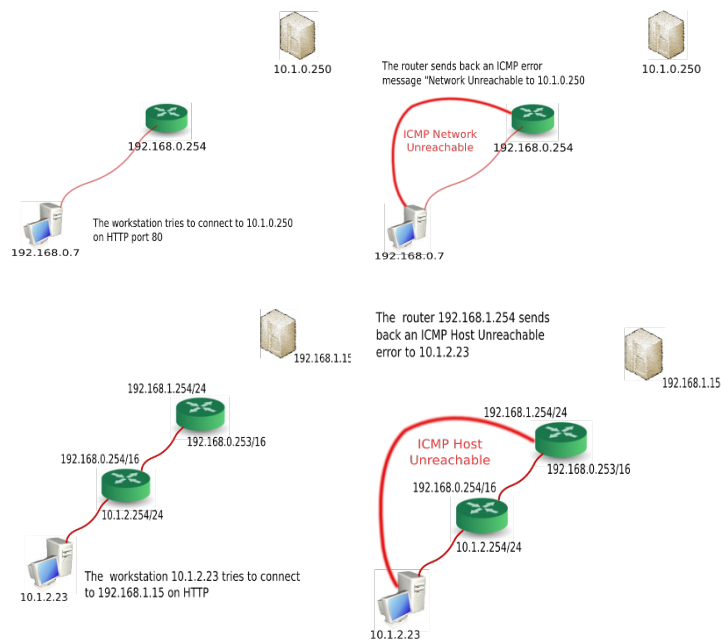
Let's describe the most typical ICMP error messages you can find on networks.

ICMP Network Unreachable

Let's take the simplest example: one machine sitting on a LAN (192.168.0.7), has one default gateway (192.168.0.254), which is the router. It is trying to reach a server, which does not sit on the LAN (10.1.0.250) and which cannot be reached, because 192.168.0.254 does not know how to route this traffic.

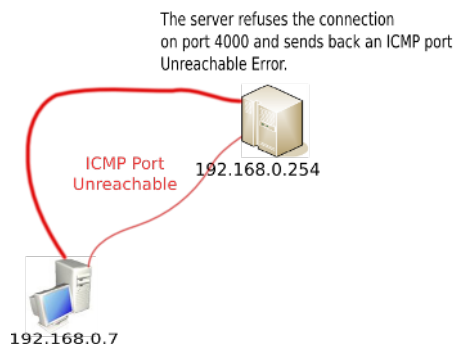
ICMP Host Unreachable

Let's take the simplest example: one machine sitting on a LAN (10.1.2.23), has one default gateway (10.1.2.254/24), which is the router. It is trying to reach a server, which does not sit on the LAN (192.168.1.15). The traffic flows and reaches the last router before the server (192.168.1.254/24); this router cannot reach 192.168.1.15 (because it is unplugged, down or it does not exist).



ICMP Port Unreachable

Let's take a second example: one machine sitting on a LAN (192.168.0.7). It is trying to reach a server 192.168.0.254, which sits on the LAN on port UDP 4000, on which the server does not respond.



Where is the challenge with ICMP?

You may be tempted to say: if it is that simple, why do we need SecurActive SPV on top of any sniffer? All the information sits in the payload. But in every network, you will find some ICMP errors... they may be due to a user trying to connect to a bad destination, or trying to reach a server on the wrong port. The key is in having a global view of how many errors you have normally and currently and from where to where. The key to leveraging ICMP information is in having a relevant view of it and understanding what it means.

How can ICMP help on network diagnostic and security monitoring?

From the explanation here above, we can keep in mind that by analysing ICMP errors we can identify machines that try to connect networks or machines, that are routable from the LAN's machine or ones that try to connect on actual servers but for services which ports are not open. Here are some examples of phenomena that can be identified that way:

Misconfigured workstation

A workstation repeats a large volume of missed attempts to connect to a limited number of servers: it may be that this machine does not belong to the company's workstations (external consultant on the network, whose laptop is trying to reach common resources on his home network -DNS, printers,...), or it may be the machine of someone coming from a remote site with its own configuration or a machine that has been simply wrongly configured.

How would we see it?

A large number of *ICMP Host Unreachable* errors coming from one or several routers to this machine or this group of machines. The *ICMP* information contained in the payload of each of these errors would probably show they are trying to reach a certain number of hosts for some services or applications.

Migration legacy

A certain number of machines keep requesting DNS resolution to a DNS server which has been migrated (this could be true for any application available on the network). Their users certainly feel worse performance when trying to use these services.

How would we see it?

A large number of *ICMP Host Unreachable* errors coming from one or several routers to a group of machines. The *ICMP* information contained in the payload of each of these errors would probably show they are all trying to reach the previous IP address of a given server.

Network device misconfiguration

A router does not have a route configured; some machines are trying to reach some resources, unsuccessfully.

How would we see it?

A large number of *ICMP Network Unreachable* errors coming from one router to many machines. The *ICMP* information contained in the payload of each of these errors would probably show they are all trying to reach the same network through the same router.

Port scanning

A machine is trying to complete a network discovery. It is trying to connect to all servers around to see on which ports they are open.

How would we see it?

A large number of *ICMP Port Unreachable* errors coming from one or several routers corresponding to a single machine (the one which is scanning).

Spyware / Worms

An infected machine is trying to propagate its spyware, virus or worm throughout the network; obviously it has no previous knowledge of the network architecture.

How would we see it?

A large number of *ICMP Host Unreachable* errors coming from one or several routers corresponding to a limited number of hosts, trying to reach a large volume of non existing machines on a limited set of ports.

Server disconnected/reboot

A service on UDP (DNS, Radius...) is interrupted because the server program is temporarily stopped or the host machine is temporarily shutdown. Many requests are then discarded.

How would we see it?

Many *ICMP Port Unreachable* messages (preceeded by some unreachable host if the host itself was shut down) are emmited during a short period of time for this service host/port.

DNS Response Time

Background:

The DNS (Domain Name System), which has been defined in detail in the [RFC 1034](http://tools.ietf.org/html/rfc1034) ([http://tools.ietf.org/html/rfc1034.html](http://tools.ietf.org/html/rfc1034)) and [RFC 1035](http://tools.ietf.org/html/rfc1035) (<http://tools.ietf.org/html/rfc1035.html>), is key to the good performance of TCP/IP networks. It works in a hierarchical way; This means that if one of the DNS servers is misconfigured or compromised, all the network, which relies on it, is also impacted. Although the DNS protocol is quite simple, it generates a significant number of issues: configuration issues, which affect the performance of the network as well as security issues, which jeopardize the network integrity. The purpose of this section is to cover the main configuration issues you may encounter with DNS when it comes to network performance.

Hypothesis:

You noticed a general slowdown for a specific host, zone, or the entire LAN. You didn't find out the issue with the previous methods. Maybe this problem has nothing to do with the business applications or you network equipment.

Diagnosis:

The DNS server(s) need to have a very high availability to resolve all the names into IP addresses that are necessary to good function of applications on the network. An overloaded DNS server will take some time to respond to a name request and will slow down all applications, that have no DNS data in their cache. An analysis of the DNS flows on the network will reveal some malfunctions like:

Latency issues

If we can observe that the mean time between the client request is significantly higher than the average (on a LAN it should remain close to 1 ms), we may face three kinds of issue:

- the client is not requesting the correct DNS server (DHCP misconfiguration, for example). You can check this out in the interface by looking at the **Server IP** fields;
- it means that the DNS server has an issue with regards to the caching of DNS names. The cache system makes it possible to resolve a name without requesting the DNS server, which has authority for the DNS zone, the IP address corresponding to the name. Hence, if the response time is high, first the application will be slow from the user's point of view, and secondly it will incude an unnecessary consumption of bandwidth. This bandwidth will be wasted both on the LAN and on the Internet link (if we make the hypothesis that the authority server sits on the Internet). If we consider the case of a fairly large organization, the bandwidth used by the DNS traffic will not be negligeable and will represent an additional charge;
- the DNS server may have system issues. If the server is overloaded, it cannot hold all the requests, and delay (or drop) some, which leads to a general slowdown of the network performances.

You can easily cast a glance at these issues: go in the **Analysis -> DNS Messages** menu, and fill the form with appropriate values (especially the **Requester Zone**), to verify if the requests are correctly answered, and in an acceptable timing.

Traffic issue

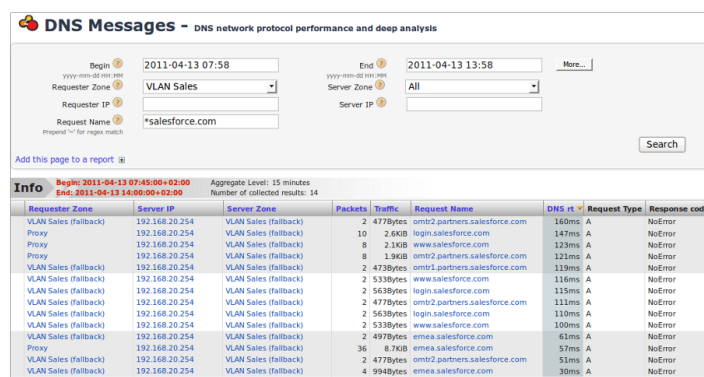


Figure 5.42: DNS Response Time for a specific requester zone (here, VLAN_Sales)

If we establish the top hosts making DNS requests, it will be possible to pinpoint misconfigured clients, not keeping in a local cache the DNS server responses; this approach makes it possible to distinguish between an issue coming from the user's workstation and one coming from the general function of the network. Please note that hosts making a very high volume of DNS requests may correspond to a malicious behaviour; for example, some malwares try to establish connections to Internet by resolving domain names and sometimes the DNS protocol is used in cover channels to escape information.

DNS errors issue

We can also ask for the top hosts receiving most DNS error messages (non existing hosts, etc.). This will also put the light on misconfigured stations, generating an unnecessary traffic and lowering the overall network performance.

DNS Internal misconfiguration

To do this, we need to identify the AXFR and IXFR transactions towards its authority server. If these updates occur too often (and therefore generate an unnecessary traffic), we can conclude that there is an issue. If the bandwidth used is too large, it means that our DNS server requests a full zone transfer (AXFR) when an iterative transfer (IXFR) would have been more adequate. If this is the case, then the network administrator can take some easy steps to improve his network's performance.

FREQUENTLY ASKED QUESTIONS

6.1 Firefox freezes randomly on some pages

This seems to be caused by the java plugin, and deactivating this plugin fixes the issue. This has no effect on SPV since it does not use java. To disable the Java plugin, enter the *Tools* → *Add-ons*. This will open a new window with a button bar on top, with a *Plugins* icon. Select it, and it will open the list of all currently installed plugins.

Locate your java plugin, that is the one that handles java applets (on the following screenshot it's titled *IcedTea NPR Web Browser Plugin*, but it may also appear under the name *OpenSDK*, or merely *Java*). Once located, select it and click on the *Disable* button. You should then restart firefox.




Figure 6.1: The Add-ons pop-up window of Firefox

6.2 Aggregate level changes when browsing from tables to charts

The aggregate level for tables is chosen to display a synthetic view on data, while the charts choose the aggregate level in order to have enough points to plot. So, this is not an error if the aggregate level changes from one page to another.

6.3 How can SRT be greater than DTT ?

Every DTT is preceded by a SRT but both are not computed simultaneously:

- DTTs are not stored until the data transfer is complete;
- SRTs are stored as soon as the first packet of the response is seen.

Thus it is frequent to have more SRTs than DTTs when browsing recent data.

6.4 How can we have 0 packets and no traffic at all on a conversation?

This is a common case when the observation period encompass the end of a timeout conversation. No packets have been sent during the observation period and the elapsed time since last packet have reached the timeout limit.

6.5 What is this timeout column (in Analysis/TCP Error)?

As there are no timeout in standard protocol (as TCP, UDP ...) this is an application level notion that the packet sniffer must guess. We consider the conversation as timeout after 2 minutes without packets exchanged.

6.6 Why are some DNS request names missing?

Although DNS protocol states that the question section must be present in the requests, not all DNS messages are name resolution requests. Some DNS server may use message types unknown of the traffic analyzer that do not embed anything meaningful in the question section of the message. For instance, the NBNS server statistic report is such a message that makes no use of the question section.

Note that you can search for empty DNS names using the regular expression `~^$` in the name search box.

6.7 Some TCP conversations are reported twice, what's wrong?

First make sure that the deduplication process is not configured too tightly. If the faulty TCP conversations keep being reported twice then maybe the duplicated packets are altered in some way that makes them too different from the originals. For instance, some firewall randomize the ISN (Initial Sequence Number) of TCP connections (for security reason). So if you mirror some traffic before and after passing through such a firewall this traffic will be reported twice since their sequence number will be different.

6.8 Pcap files generated by tcpdump are (mostly) empty

By far the most probable reason for this is that you are trying to use a filter on VLAN tagged packets. This won't work since *Tcpdump* filters look for fixed locations in the packet and the VLAN tag offsets the actual bytes that are being matched. Fortunately there is a workaround: by adding the filter `vlan` all following filters will be offset by the VLAN tag size. So for instance if you want to filter `ip proto \tcp` on an interface receiving only VLAN tagged packets then you must use the following filter instead:

```
vlan and (ip proto \tcp)
```

If the network interface receives both tagged and non-tagged packet then this somewhat cumbersome filter must be used:

```
(ip proto \tcp) or (vlan and (ip proto \tcp))
```

6.9 How to do complex searches on domain names?


On search boxes about domain names (Web and DNS reports), you can use a regular expression by prefixing the entry with a *tilde* character (`~`). For example, you can use this to filter **all but some names**. For instance, here is a valid input to filter all but *Google's* and *Amazon's*:

~^(?!(.*\.)?google\.[fr|com]\$)(?!.*amazon(\.\.{2,3}){1,2}\$)

6.10 What about Open Source?

SecurActive uses internationally proven and rock solid open source components such as *Linux*, *Python*, *Zope*, *Postgresql*, *Git*, *GCC* ... Our company has chosen to actively contribute to the open source community by regularly submitting patches to these projects and provide access to parts of its own code ¹.

6.11 Standard TCP Session



Securactive Performance Vision interpreting a standard TCP session - v1.2

Figure 6.2: Standart TCP Session

¹ <https://github.com/securactive>

KNOWN ISSUES

7.1 Configuration

If an application is defined with both a webpattern and a client or server zone then all conversations matching this webpattern but not the zone will belong to NC TCP, even if it should belong to another application according to the TCP ports.

7.2 Interface

There is no error message when a login attempt fails.

Sometimes when plotting data for the last hour, the chart ends with a value at zero.

In the charts, when a high value is immediately followed by a zero value then the smooth interpolation algorithm makes it go underneath the 0 line just after the 0 value.

When another language than English is requested some buttons are labelled in English nonetheless.

7.3 Various

SMTP delivery of reports lack retries.

There is no procedure to delete oldest data whenever the data disks become full.

Configuration dump/restore don't work across version boundary.

7.4 Sniffer

In case of IP fragmentation, the timestamps of involved packets are set to the last received one.

7.5 Upgrading

In some cases, the sniffer may fail to restart after an upgrade and leave some stalled processes if it is restarted on its own with *Pulsar* (page 26). One of the possible symptoms is that the `poller` command in Pulsar fails to display the poller and license status. Rebooting solves this issue.

7.6 Metrics

In versions prior to 2.9, the retransmission rate (RR) was computed as the number of retransmitted TCP segments divided by the total number of TCP segments. As of version 2.9, it is instead divided by the number of packets liable be retransmitted, such as the TCP segments carrying a payload.

In versions prior to 2.9, keep-alive packets occurring after the completion of a data transfer were taken into account in the computation of the Data Transfert Time (DTT) metric, resulting in abnormally large values. In order to avoid this issue, as of version 2.9, data transfers are considered complete after a 1 second timeout.

GLOSSARY

Aggregation period Time period over which all data are aggregated into flows (for each set of client, server and application). The Aggregation Period is defined for an aggregation level as time interval over which all flows are aggregated in the database on their IP src/dst, Zone src/dst and Application. Individual flows within the aggregated data cannot be viewed separately; The Aggregation Period defines the data resolution for an aggregation level.

Application Group logical or business related flow to emphasize valuable perspective an Application is identified with a name and a color, and defined by a set of Signature or a set of Port Range (at least one non-empty set of either), a set of client and server zones. A conversation is attributed to an application with the following rule: (PORT_RANGE1 OR OR PORT_RANGEn OR SIGNATURE1 OR ... OR SIGNATUREn) AND ((SERVER_ZONE1 OR ... OR SERVER_ZONEn) AND (CLIENT_ZONE1 OR ... OR CLIENT_ZONEn)); in case a conversation matches previous rule of several application, the priority will be given to the application whose definition is the most precise, i.e. the thinnest port range, signature or server/client zone.

Application NC NC stands for Non Classified. A NC Application is a special application that will match conversations that do not match any configured application.

Application Port Range Port or range of ports on a defined protocol (TCP, UDP). a Port Range is defined by a range of ports (described by a start port and an end port) and a protocol, either 'TCP' or 'UDP'.

Application Signature Mean of recognizing an Application based on a pattern in the payload. This pattern may be of two sorts: a dynamic port signature or web application signature. A Signature is either a Dynamic port Signature or a Web application Signature.

Connection Time (CT) Time taken by the exchange of the 3-way TCP handshake. CT stands for Connection Time. CT is defined as the duration of the three way handshake (SYN, SYN/ACK, ACK) of TCP session.

Conversation Regroups network exchanges between two network addresses for one application during the observation period. A conversation is defined as a group of flows between a client and a server over an observation period.

Data Transfer Time Time spent by the client or the server to send data. The DTT stands for Data Transfer Time. DTT server is defined as the time between the first data packet (with ACK flag and a non null payload) from the server and the last packet considered as part of the same answer. DTT client is the symmetric metric in opposite direction. Packets are considered part of the same answer if packet share the same acknowledgment number ; FIN, RST from server or client. A Timeout will cancel a DTT. Note that if the answer is small enough to be contained in only one packet, the DTT will be of '0'.

Delta sessions Number of session established minus those closed. Delta Session is a metric defined as the difference of the number of opened session to the number of closed session. Negative value means that more session were closed than opened.

Device Identifier Identifies the physical network adapter that received the network traffic associated to a conversation.

End User Response Time Total time the user waited to get an applicative answer. The EURT stands for End User Response Time. EURT is defined as the sum of the RTT (client + server), the SRT and the DTT (client + server). A timeout will cancel the computation of EURT.

Fallback Set of IP addresses which belongs directly to a zone (and not to any of its children zones). The Fallback of Zone A is an implicit Zone containing the set of the addresses which belongs to Zone A excluding all addresses belonging to A Child Zones.

Flow Regroups data exchanges between two network addresses for one application on the aggregation period. A flow is a group of communications between two network addresses for one application during the aggregation period. Notice that the VLAN tag, if present, as well as the device identifier, are considered components of the network address.

Initial Sequence Number The sequence number used in the SYN packet of a TCP connection.

Jitter Packet delay variation. The Jitter is defined as the variance of RTT (average difference between RTT measures and the average RTT). For more details, this equation is used: $\text{Sqrt}(\text{Average}(\text{RTT1}^2, \dots, \text{RTTn}^2) - \text{Average}(\text{RTT1}, \dots, \text{RTTn})^2)$.

Observation period In all reports, defines the observation time window. Observation Period is based on a starting time and an ending time provided by the user. These user-defined boundaries will automatically be moved to the closest previous aggregation boundary for the starting time and to the next aggregation boundary for the ending time: this modified time interval is the actual observation period.

Protocol The transport protocol relying on IP at the network level. Protocol is defined as one of the IP protocols that SPV can track. It can refer to 'OtherIP', 'TCP', 'UDP', or 'ICMP'. These protocols are detected by inspecting packet headers.

Retransmission Packets being resent, when they have either been lost or damaged. Packet Retransmission is identified thanks to their TCP sequence and acknowledgment numbers, and checksum values. Only packets with a non-null payload are checked.

Retransmission Delay Delay between a packet and it's the next retransmission. RD stands for Retransmission Delay. RD is defined as the time between a packet and its next retransmission.

Retransmission Duplicate ACK Duplicate acknowledgment Packet with null payload. Duplicate ACK are TCP ACK packets that are identified thanks to their same acknowledgment value and their empty payload.

Retransmission Rate Ratio of retransmitted packets to the total number of packets. RR stands for Retransmission Rate. RR is defined as the ratio of retransmitted packets to the total number of packet in a conversation.

Retransmission Total Delay between a packet and the last retransmission. TRD stands for Total Retransmission Delay. TRD is defined as the time between a packet and its last retransmission.

Round Trip Time Time between an applicative query and a response at the network level. RTT stands for Round Trip Time. RTT is defined as the time between a packet with a non null payload and the corresponding acknowledgment (a packet with a null payload and the TCP ACK flag).

Server Response Time Time between a query and an answer at the applicative level. Server Response Time is the elapsed time between a client packet with a non null payload and the corresponding server response (a packet with a non null payload which number of acknowledgment correspond to the first packet).

Session An established communication channel between two devices using TCP. a Session is defined as TCP communication between 2 devices beginning by a successful Handshake, and ending by a Timeout, or Packet with the RST flag from any of the devices, or a Packet with FIN from any of the device that is acknowledged by a FIN/ACK by the other device and followed by a FIN of this same last device. (no FIN/ACK is necessary to conclude that the connection is closed).

Signature Dynamic port Connection tracking of 'Application' based on dynamic TCP/UDP port negotiation detection a 'Dynamic Port Signature' is identified by a name in a set of internally predefined pattern name and an associated port. Each pattern name refers to an internal connection tracker that will start from the given associated port and will follow connexion on other port.

Signature Web application Allows to distinguish conversation that use HTTP by using simple pattern matching on the target URL. A Web application signature is defined by a single pattern. The pattern syntax allows hostname and optionally a path separated by '/' (ie: 'www.example.com/my/path', or 'www.example.com'). Notice that a wildcards character * is allowed in domain or path part of the pattern. Only Conversation which are detected to be based on HTTP will have URL of their GET/POST/CONNECT request matched against Web application signature's pattern. A match occurs when the pattern match the complete target URL.

Subnet Set of network addresses that have a common declared IP address routing prefix. A Subnet is defined by an IP address and a netmask.

TCP Handshake 3-Way negotiation that is part of TCP for establishing a TCP session. A TCP Handshake is defined between 2 devices as exchange of 3 TCP packets flagged SYN, SYN/ACK, ACK.

Time To First Byte Time for a user to connect to a server and receive a first response from the application. TTFB stands for 'Time To First Byte' and is defined as the interval between the SYN packet and the first packet with a non null payload from the server.

Timeout Session end by inactivity. Session Timeout will be reported after 120 seconds of complete inactivity (i.e. no packets seen).

Zone A logical group of subnets. A Zone is identified with a name, and defined as a set of subnets. More, it must be placed in the Zone container hierarchy following subnets natural subnet inclusion constraints. A Zone is attributed to an IP if the IP is included in one of the Zone subnets.

APPENDIX

9.1 Virtual Appliance Step-by-Step

9.1.1 How to get the image of the Virtual Appliance

This section is based on version 2.5.9, the filename will evolve depending on the version number.

The ZIP archive will contain the following files:

- SPV-2.9.4-r1.mf
- SPV-2.9.4-r1.ovf
- SPV-2.9.4-r1-disk1.vmdk

9.1.2 Virtual Appliance Specifications

The Performance Vision Virtual Appliance is designed to run in a VMWare ESX v4 or v5 environment.

It is designed to run with a minimum RAM of 500MB, although a larger quantity is recommended to ensure satisfactory performance rates. Here are the configurations which are validated:

RAM: 512MB, 4GB, 6GB, 8GB, 12GB or 16GB vCPU: 1, 4 or 8

9.1.3 Installation

The system detects the space available on the disk for the new Virtual Machine, we recommend to allocate the following spaces:

- Trial Virtual Appliance: 4GB RAM, 2 vCPU > 2,0 GHz
- Production:
 - Virtual Poller: 8 GB, 2 vCPU > 2,0 GHz
 - Virtual Appliance: > 16 GB, 4 vCPU > 2,4 GHz

You get:

Get it Started

Once the Virtual Appliance is installed, you have to start it.



Figure 9.1: Connect to your Vsphere Client.

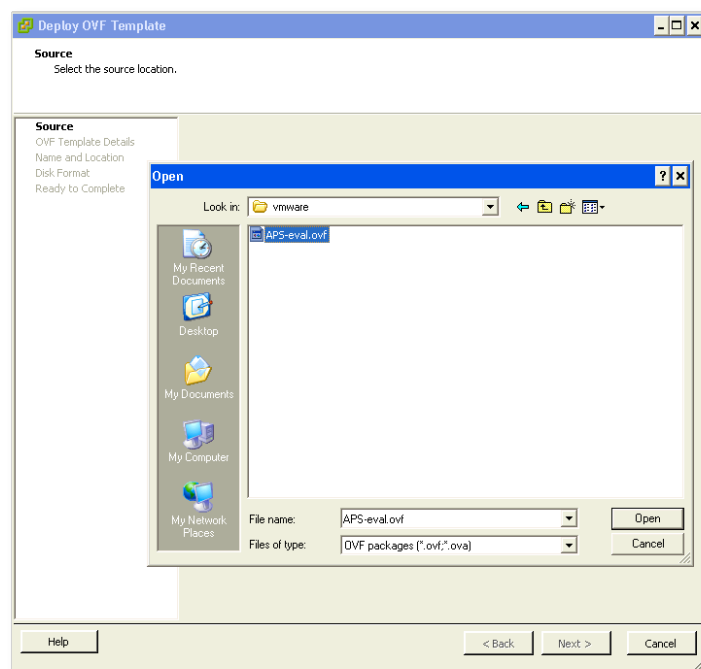


Figure 9.2: In the Virtual Machines tab, in the « File » menu, select « Deploy a new OVF template ». Find the Performance Vision OVF file. and Click on « Open ».

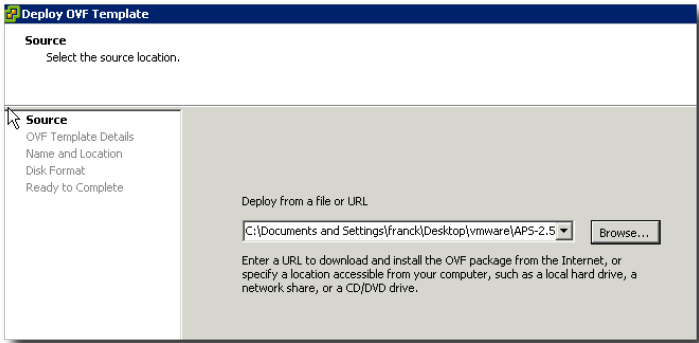


Figure 9.3: Click on « Next ».

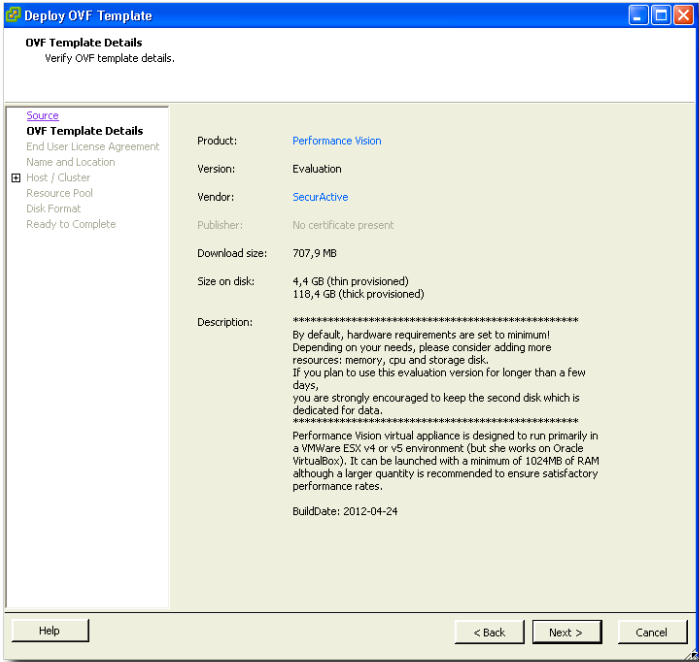


Figure 9.4: Click on « Next ».

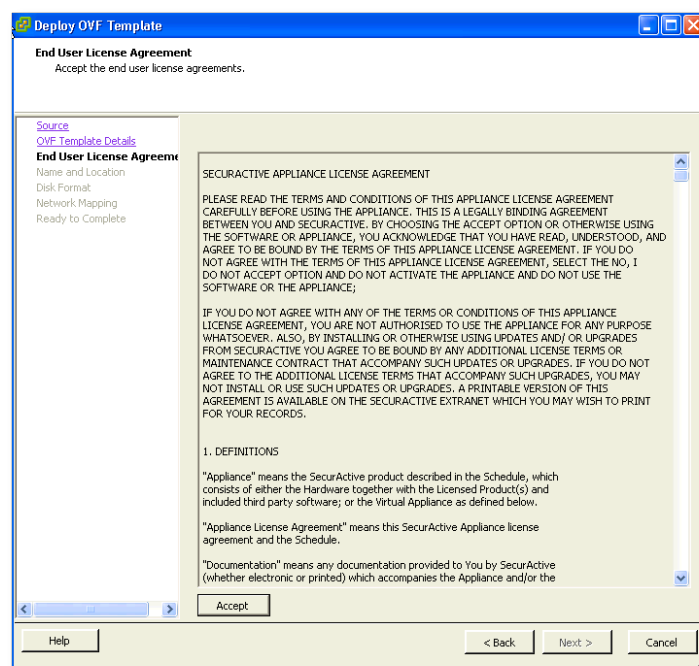


Figure 9.5: Read, then click on « Accept », then click on « Next ».

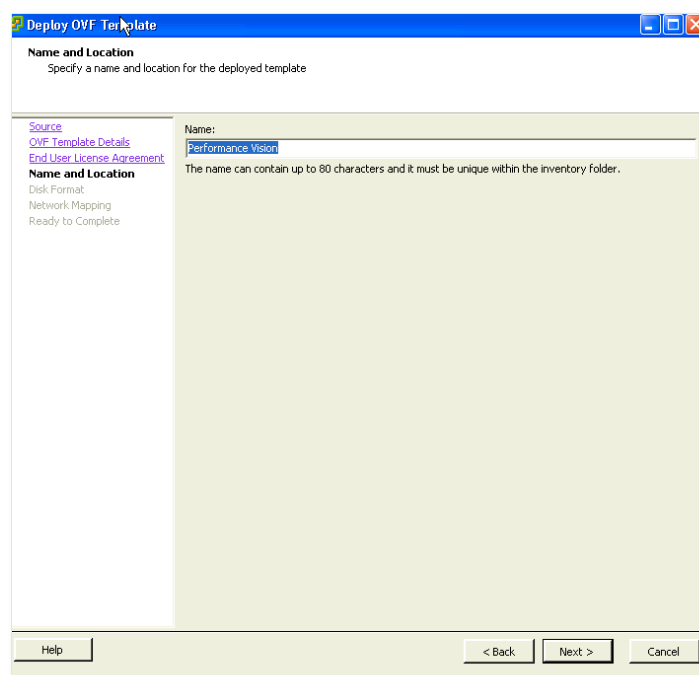


Figure 9.6: Name the Virtual Machine appropriately and click on « Next ».

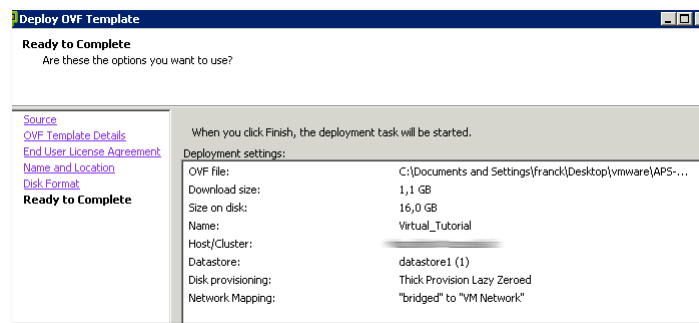
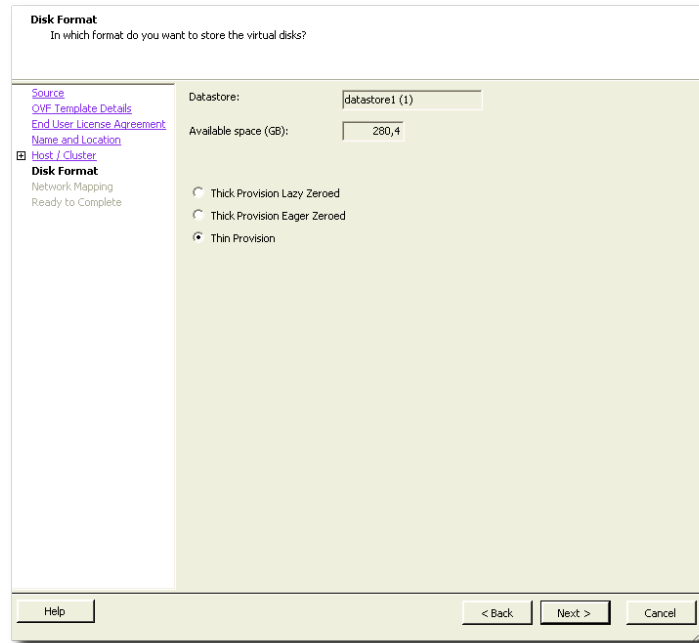


Figure 9.7: Click on « Finish », the Virtual Appliance gets installed.

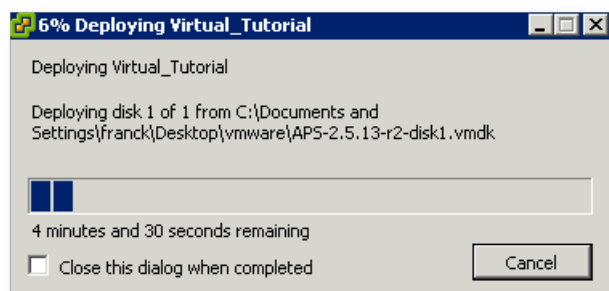



Figure 9.8: You get notified when the installation is complete.



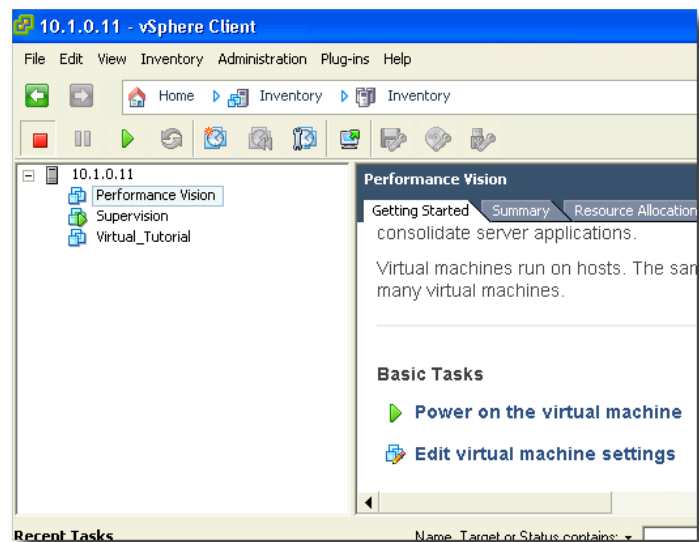
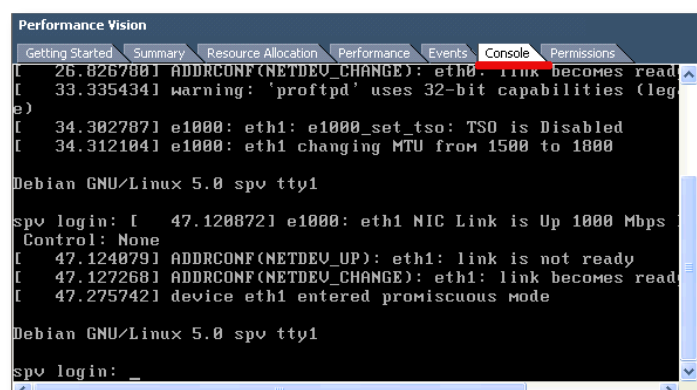


Figure 9.9: Click on « Power on the Virtual Machine » or on the green triangle.

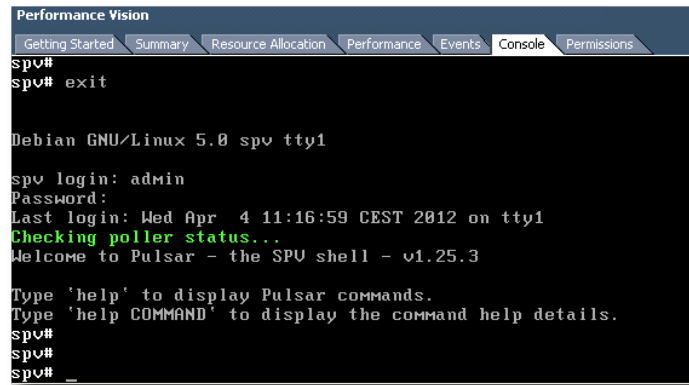


Access the virtual console

Display the Console tab and access the CLI interface named `Pulsar`.

The probe is launched. When the network interfaces turn into promiscuous mode, click on the Console view and then « Enter » to display the login prompt.

Note: Clicking on the black screen deactivates your mouse. To reactivate it, you can use the key combination `Ctrl + Alt`.



```

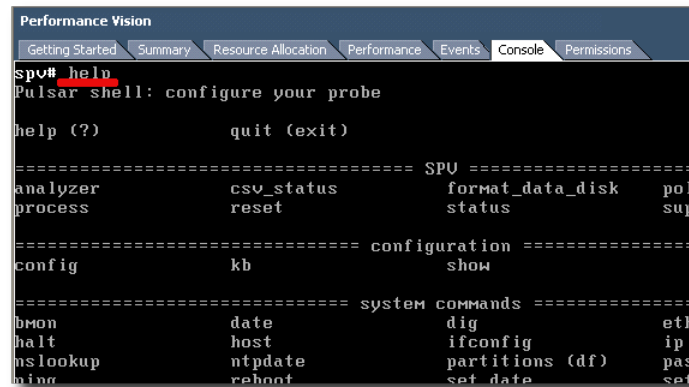
Performance Vision
Getting Started Summary Resource Allocation Performance Events Console Permissions
spv#
spv# exit

Debian GNU/Linux 5.0 spv tty1

spv login: admin
Password:
Last login: Wed Apr  4 11:16:59 CEST 2012 on tty1
Checking poller status...
Welcome to Pulsar - the SPU shell - v1.25.3

Type 'help' to display Pulsar commands.
Type 'help COMMAND' to display the command help details.
spv#
spv#
spv# _
  
```

Figure 9.10: The default credentials are : user = admin and password = admin.



```

Performance Vision
Getting Started Summary Resource Allocation Performance Events Console Permissions
spv# help
Pulsar shell: configure your probe

help (?)          quit (exit)

===== SPU =====
analyzer          csv_status       format_data_disk  pol
process           reset            status            sup

===== configuration =====
config           kb                show

===== system commands =====
bmon              date              dig                eth
halt              host              ifconfig           ip
nslookup          ntpdate           partitions (df)    pas
ming              reboot            set_date           set
  
```

Figure 9.11: The « help » command lists the possible actions.

Note: The virtual machine has a second 100 GB hard disk that you can resize depending on your needs, but then you'd have to format it (via pulsar's `format_data_disk` command).

To enter several DNS servers, the addresses must be separated by a space.

You then have to reboot the Virtual Appliance.

Insert a license key

Except the empirical virtual appliances of test provided from our Web site, the virtual appliances are delivered without license key. You normally receive this key by e-mail at the product's delivery. If it is not the case, please contact our sales department: sales@securactive.net. To install a license key, it is necessary to be connected to the virtual appliance by FTP in binary mode. Filezilla Client does it by default. Connect to the Virtual Appliance:

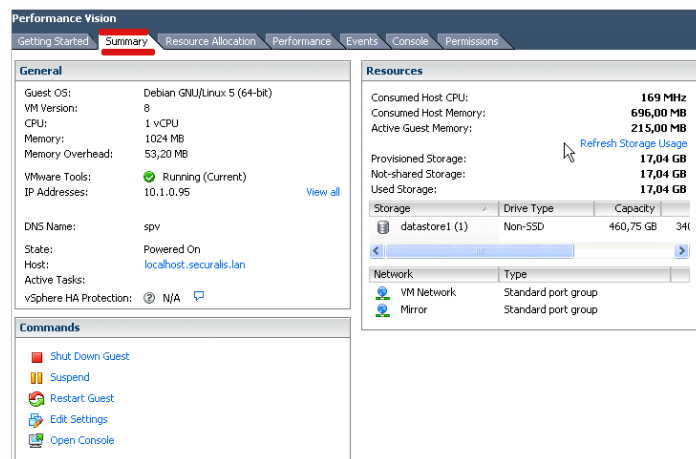


Figure 9.12: The summary view provided by Vsphere displays the parameters such as IP addresses:

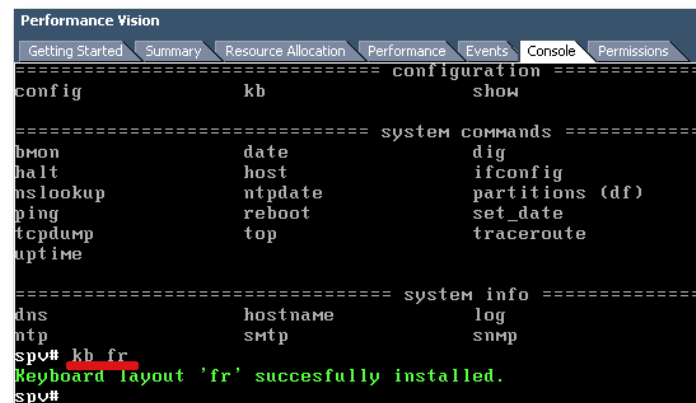


Figure 9.13: The command kb [parameter] enables to change the keyboard language configuration: for example, « kb fr » for French keyboard.

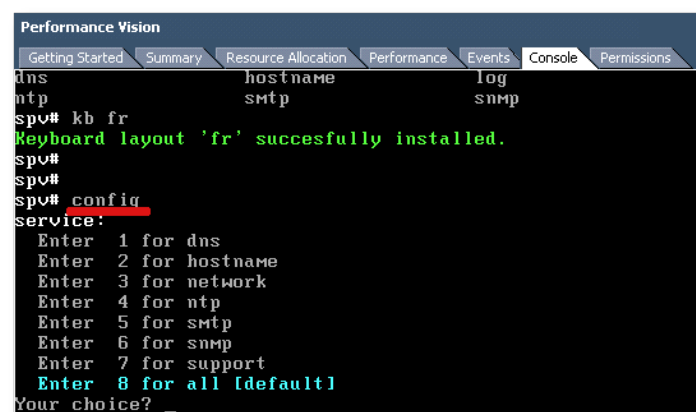


Figure 9.14: The « Config » command is used to setup the probe in function of your environment.

```

APS# config dns
[DNS] (type '?' for help)
DNS server [10.1.0.6]: 10.1.0.6 8.8.8.8
dns service has been updated
Configuration summary
[DNS]
servers      : 10.1.0.6 8.8.8.8

```

Figure 9.15: DNS configuration

```

APS# config network
[NETWORK] (type '?' for help)
Connection Type:
  Enter 1 for Static network [default]
  Enter 2 for DHCP
Your choice? 1
IP address [10.1.0.127]:
netmask [255.255.255.0]:
gateway (type 'empty' to remove) [10.1.0.1]:
network service has been updated
Configuration summary
[NETWORK]
interface      : eth0
connecttype    : static
address        : 10.1.0.127
netmask        : 255.255.255.0
gateway        : 10.1.0.1

Changes will not be effective before next probe reboot.
Reboot now?:
  Enter 1 for yes
  Enter 2 for no [default]
Your choice? _

```

Figure 9.16: Network configuration

Login: ftp
Password: S3c7r!

You can validate that the license is right from the Web interface of the virtual Appliance by clicking on « Configuration », then « Pollers Status ». The license key will be displayed (with its expiry date and its status).

See [License and upgrade installation](#) (page 25) for more details.

Access to the Graphical Interface

Use the IP address configured for the Virtual Appliance to access the GUI with a Firefox web browser.

Depending on the configuration, the probe can be accessed through the following ports : TCP/80, TCP/8080 or HTTPs. The default account to access the GUI is user = admin, password = admin. Beware that this account is distinct from the account used to access Pulsar.

Apart from the trial version, the Virtual appliances are provided with no license key. You have to get the license key, which will be provided by email by SecurActive.

Traffic capture

First of all:

- The port mirroring should be activated on yours switches (or TAP eventually)
- Connect the mirror destination port to the ESX server port dedicated to the traffic capture

We will now set the network in Promiscuous mode.

In The following example, we are using an ESX server with 8 physical ports. It is necessary to add a virtual network for traffic monitoring. How to do it?

- 1. Connect to Vsphere Client

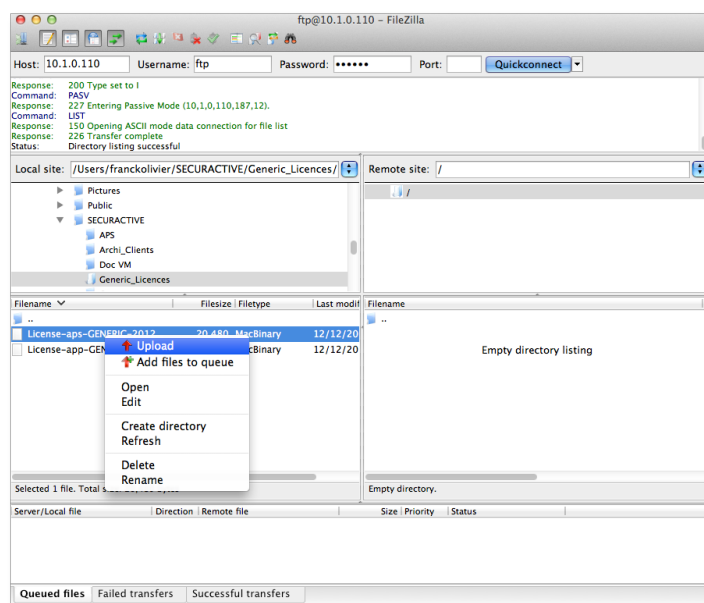


Figure 9.17: Find the Performance Vision license key file, right click on it and choose Upload




Figure 9.18: The installation is complete when the license key is not available anymore by refreshing the destination folder lists






Figure 9.19: Once logged in, you access this page.

```
Usage:
poller
poller add IP
poller modify IP
poller delete IP
APS# poller

Name:          APS
Address:       localhost
Created the:   2012-02-10 17:23
Device ID:    564D9B2D-C67F-A562-730B-6848A6850682
Device md5:   89fa72d1020baf6f6cc00da13b364d2a
Time:         2012-02-15 17:24
SPV Version:  2.5.13-r2
Sniffer status: ok, pid 2760
Sniffer version: 2.5.13
License:      invalid
Expiration:   no limit
```

Figure 9.20: The status of the license can be validated in Pulsar with the command « poller ».


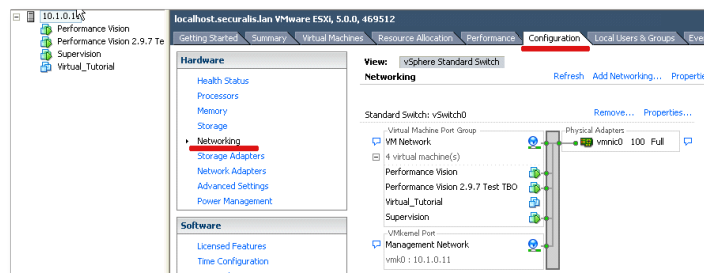



Figure 9.21: It can also be done through the web interface, in the page « Poller status » in the Configuration section.

- 2. Then on your ESX server icon, go to the « Configuration » tab:
- 3. Click on the « Networking » Menu on the left column



- 4. Click on « Add Networking »



Vlan ID (optional) for vlans tags:

0: Disables VLAN tagging on port group
 4095: Enables VLAN tagging on port group

- 5. Then click on « Next » and « Finish » to complete the operation.

Setup promiscuous parameters.

The Esx Server now manages 2 virtual networks.

The aim of the second vswitch (vSwitch1) is to show the flows in promiscuous mode.

To set up promiscuous mode on the Mirror Network:

Add a listening network card to virtual appliance.

Here we should add a listening network port in promiscuous mode:

- Power on the virtual appliance.
- Validate traffic Capture

There are 2 main methods to validate the traffic capture:

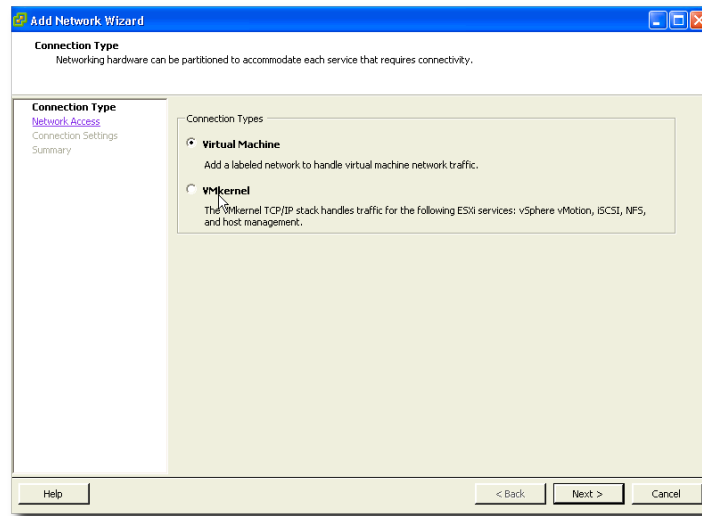


Figure 9.22: Select Virtual machine as Connection Types, then Click on « Next »

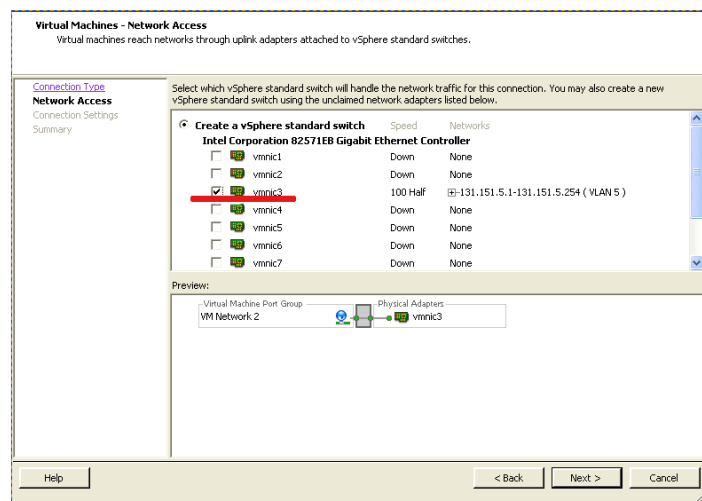


Figure 9.23: On « Network Access » Menu, select the Esx physical port dedicated to the traffic capture (here is vnic3) and unselect the others. The Esx physical will be binded to the new virtual network (here VM Network2) Click on « Next »

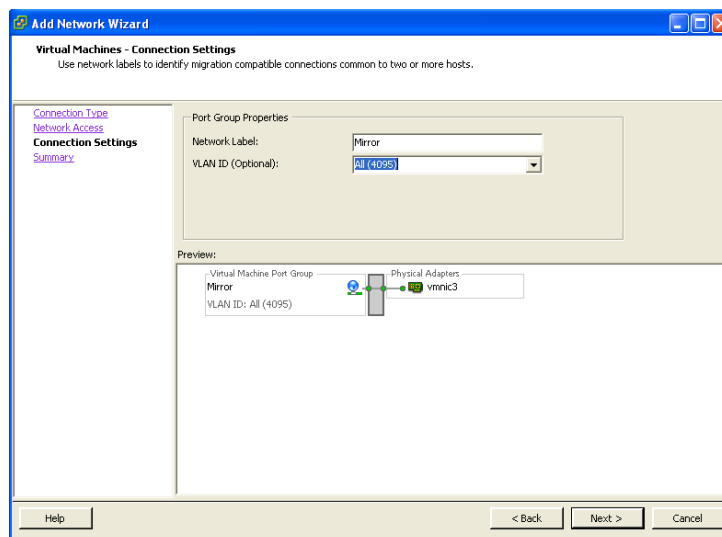
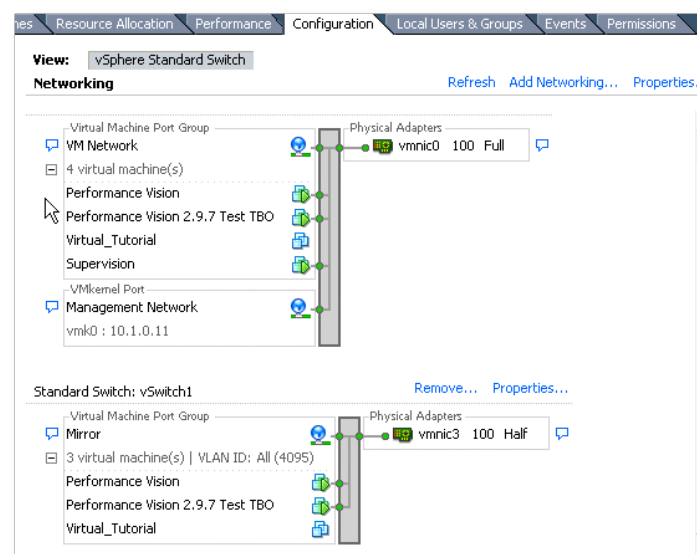
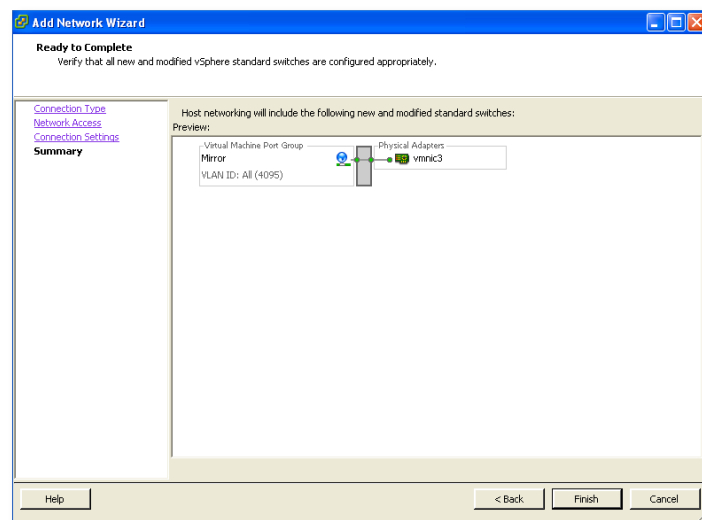


Figure 9.24: We can customize the new network label as “Mirror” here. The following option allows VLAN tags:






Figure 9.25: Click on «vSwitch1 Properties »

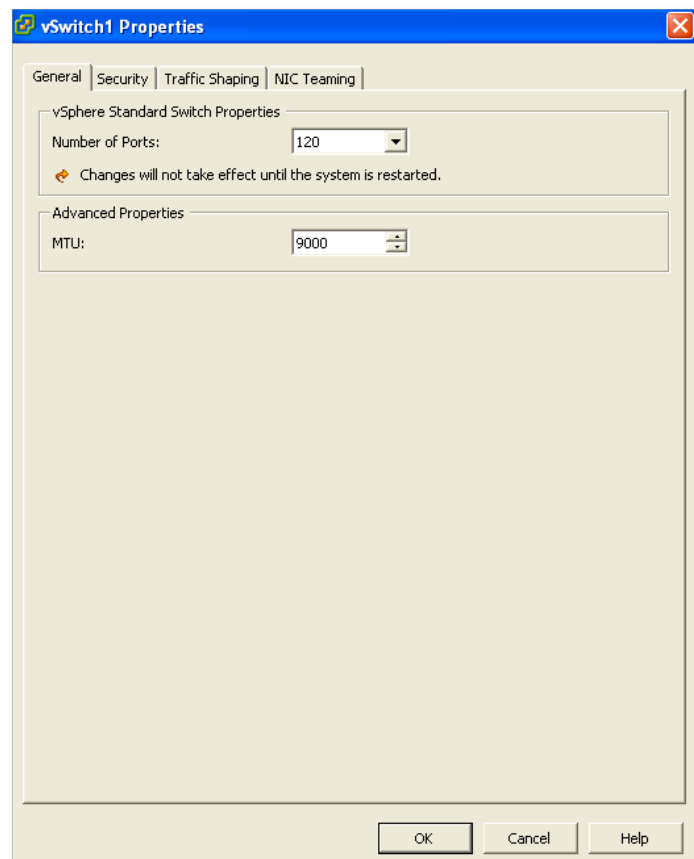


Figure 9.26: Double click on « vswitch »
In General Tab: Edit MTU settings to 9000

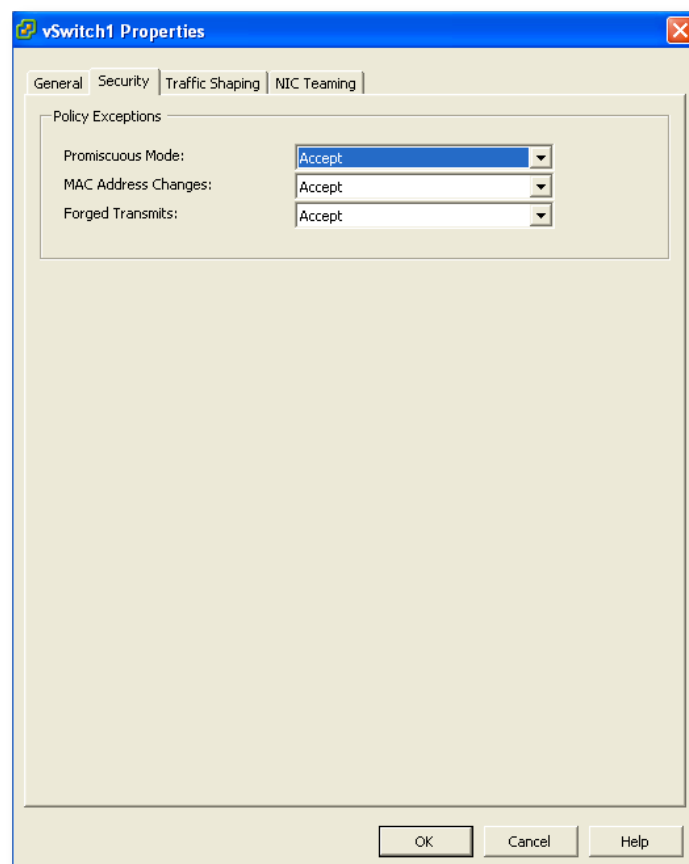


Figure 9.27: In Security tabs: Select « Accept » from the Promiscuous mode Listbox

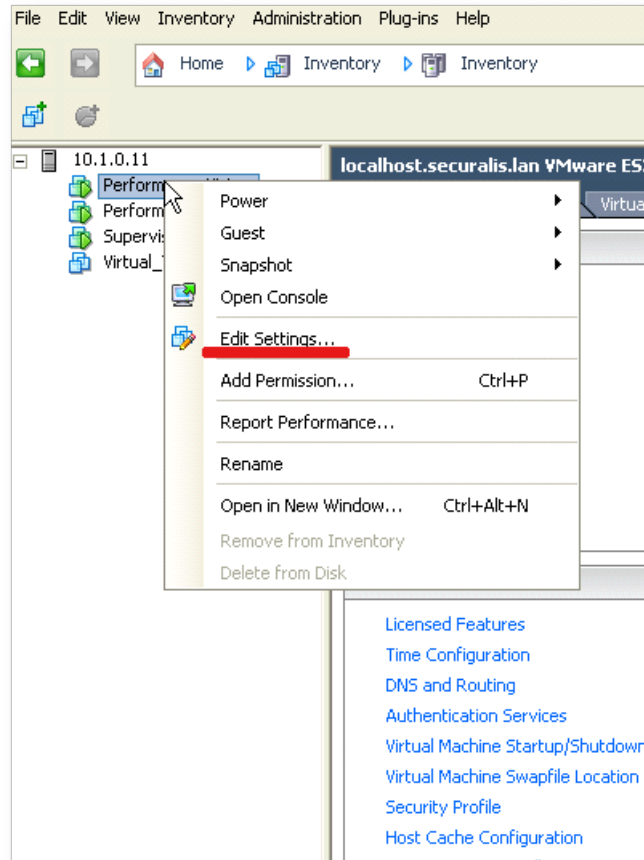


Figure 9.28: Right click of the virtual appliance then choose « Edit settings »

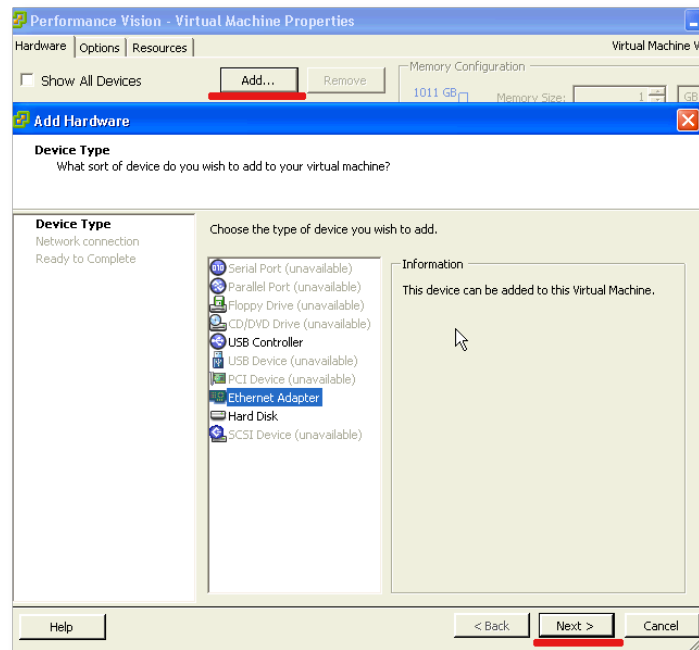


Figure 9.29: In the Hardware tab:
Click on « Add », then choose Ethernet adapter and Click on « Next »
Attach the New Ethernet adapter to the Network in promiscuous mode

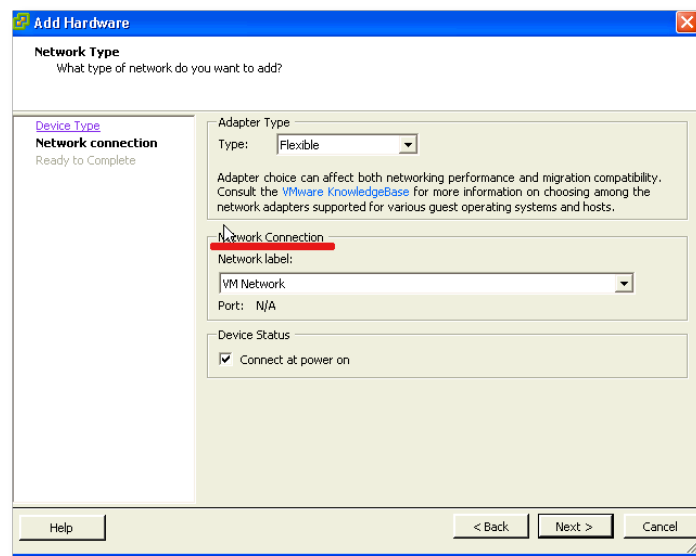


Figure 9.30: In the Network Connection Listbox, choose the accurate network configured above (Mirror here) Click on « Next »

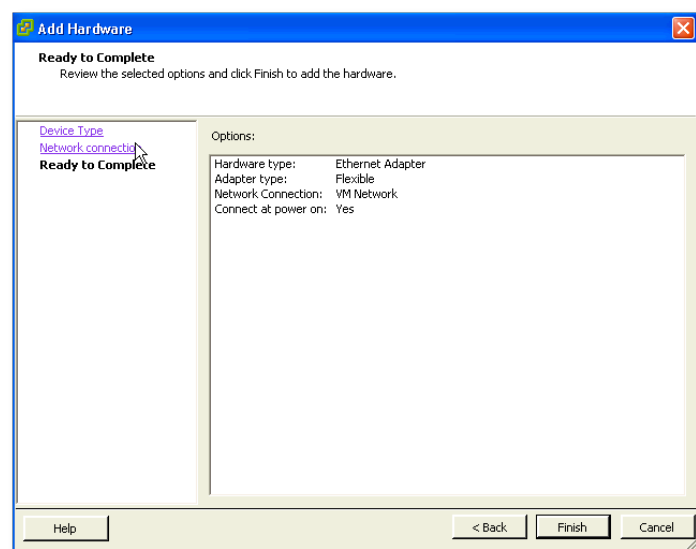
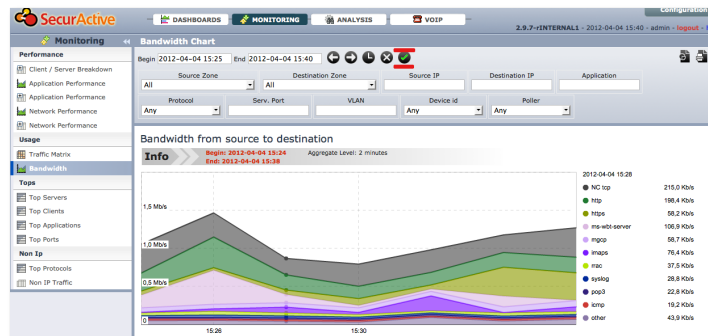


Figure 9.31: Click on « Finish » to complete the operation


Graphical Interface

As an example, you can monitor the bandwidth after 5 minutes of listening by clicking on the green validation button.



Pulsar

Connect to pulsar via ssh or from the virtual appliance console on the Esx.



9.1.4 Configuration

The Performance Vision Virtual Appliance is shipped with a default configuration that will likely not match your site very closely. For a better experience it is recommended that you spend some time configuring some additional zones and applications to suit your traffic.

Here are the sections you should consult, in order:

- *User Management* (page 29) for adding new users;
- *Zone configuration* (page 30) for adding new zones or modifying the preset configuration;
- *Application configuration* (page 32) for registering your specific applications;
- *Business Critical applications* (page 34) and/or *bcn_config* to define your business critical applications/links;
- *Reports* (page 37) to schedule periodic reports that will be sent via email.

9.1.5 Six pages you should not miss in Performance Vision

Network performance

Performance Vision provides a series of views on how your network is behaving. Here is a selection of views you should absolutely use.

Business Critical Networks

Provided you have configured some critical networks (setting thresholds on volume and quality indicators between 2 zones), you will get a summary screen of the performance of your most critical network links on this screen. This is an auto-refresh screen, whose data can be integrated in your SNMP based monitoring suite.

By pointing a specific time and link, you can view the origin of a degradation (latency, retransmission, excessive bandwidth consumption, and in which direction it occurred).




Figure 9.32: Business Critical Networks

You can access this view in the graphical interface in Dashboards / Critical Networks.

Network Performance table

This view will show you the main network performance metrics from zone to zone. By unfolding the different layers of zones, you will be able to identify between which parts of your network high latency, retransmissions can be observed. This view is excellent to assess the performance of your network depending on your topology.




Figure 9.33: Network Performance Table

You can access this view in the graphical interface in Monitoring / Network Performance Table.

Network performance chart

This view will show the main network performance metrics through time for a given selection (from one zone to another for example): round trip time, retransmission delay, connection time, retransmission rate, volume of packets. This shows the evolution of the network performance; as in any view in Performance Vision, you can drill down to the conversation level by clicking through the graphs.

You can access this view in the graphical interface in Monitoring / Network Performance Chart.

Application performance

Performance Vision provides a series of views on how your applications are behaving. Here is a selection of views you should absolutely use.

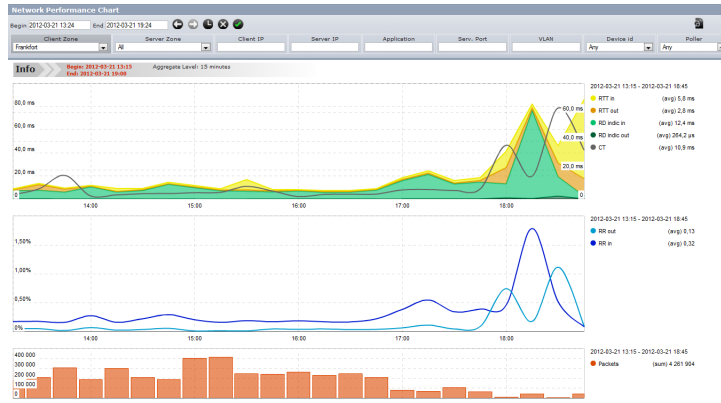


Figure 9.34: Network Performance Chart

Business Critical Application Dashboard

Provided you have configured some critical applications (setting thresholds on quality for a given application), you will get a summary screen of the performance of your most critical applications on this screen. This is an auto-refresh screen, whose data can be integrated in your SNMP based monitoring suite.

By pointing a specific time and link, you can view the origin of a degradation (round trip time, server response time, data transfer time, quantity of transactions).




Figure 9.35: Business Critical Application

You can access this view in the graphical interface in Dashboards / Critical Applications.

Application Performance Dashboard

A simple click from the Business Critical Application Dashboard drives you to the Application Performance Dashboard: it shows you the evolution of the End User Response Time through time (along with the volume of transactions) and its breakdown in Round Trip Time, Server Response Time and Data Transfer Time. At a glance, you can understand the origin of a change in the End User response time.

Underneath this first graph, you find two additional bar charts, which help you understand which server(s) and Client Zone(s) are performing better / worse (and due to what component of the End User Response Time). The servers and zones are always presented from the one that corresponds to the highest volume of transactions to the lowest.

You can drill down and display either the Client Application Dashboard or the Server Application Dashboard by clicking on a specific server or client zone. This drives you to a specific application dashboard focusing on the same application for that specific server or client zone.

This view is available for any TCP application in Dashboards / Application Dashboard.

Application Performance Chart

A more detailed view of the application performance is available here; it will show an even more complete set of metrics: RTT client & server, Server Response Time, Data Transfer Time client & server, retransmission rate, volume of packets.

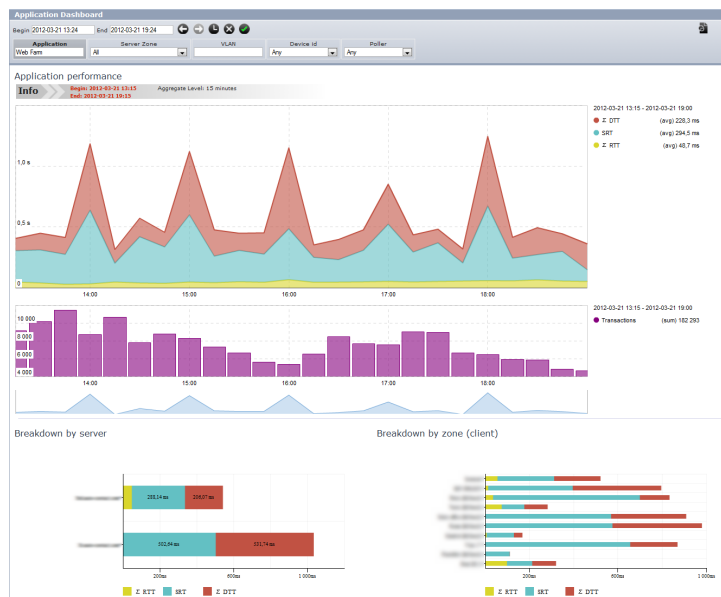


Figure 9.36: Application Performance Dashboard

Using filters you can focus on a specific perimeter and view the evolution of the application performance through time. This view is specifically interesting to link the evolution of data transfer times to retransmission rates and data volumes.




Figure 9.37: Application Performance Dashboard

This view is available for any TCP application in Monitoring / Application Performance Chart.

Matrix and bandwidth

Performance Vision provides a set of reports on traffic volumes.

Matrix View

This view shows the mapping of traffic. Using filters you can get this mapping for a specific part of your network or application. It shows the quantity of traffic exchanged from zone to zone. To further in details you can unfold zones to display its sub-zones.

The color code easily shows where the largest traffic can be observed.

You can also look for abnormal traffic through this view, by looking at cells where there should be no traffic in normal conditions (example: Internet to Internet, if you are capturing traffic on your private network).

You can drill down to a bandwidth graph or to detailed conversations with a single click.

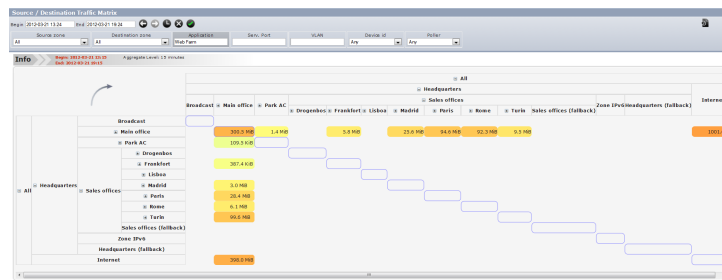


Figure 9.38: Matrix View

This view can be accessed through Monitoring / Traffic Matrix.

Bandwidth Graph

You can graph the evolution of bandwidth through time.

From there, you can drill down to detailed conversations to display the main contributors of a peak of traffic for example.




Figure 9.39: Bandwidth Graph

This view can be accessed through Monitoring / Bandwidth chart.

Top reports

You can easily get the top clients, servers, applications for any traffic (all or a specific application, zone, etc...)

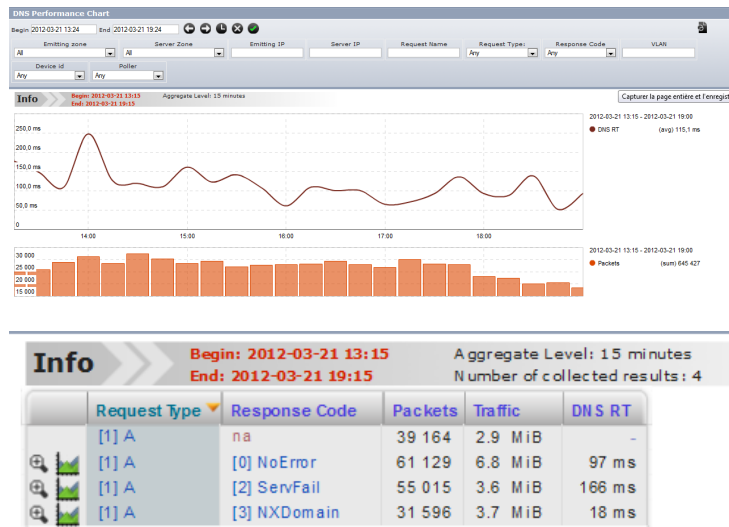
You can sort each top on the most adequate criteria (volume, sessions, SYNs, etc...)

This view can be accessed through Monitoring / Top Reports.

DNS performance

Performance Vision provides in depth view of name resolution events and performance (for DNS, Netbios, mDNS, ...) When conducting a troubleshooting, this view can display:

- The evolution of the DNS activity (an excessive peak may reveal a misconfiguration / infection)
- The evolution of DNS response times (which impacts the quality of experience of end users).
- Unexpected name resolution protocols (are you still using Netbios/WINS, when you thought you only rely on pure DNS? Have you got more DNS requests in error than successful ones?)
- Are some of my hosts trying to resolve out of abnormal servers? (Rest of migrations, misconfigurations, infections).
- Can I see hosts with abnormal request volumes? (infection, misconfiguration)



- Have I got some configuration issues (short TTL values, lack of caching)? Look at the DNS conversations with the largest number of transactions.

This view can be accessed through Diagnostic / DNS.

TCP events

Performance Vision provides in depth view of TCP anomalies and events. When conducting a troubleshooting, this view can display:

- TCP conversations where the sessions are not ended correctly (Time outs, RSTs...) This may help you understand when you can observe disconnections, if the client or server side is responsible for it.
- Bad transmission rate: if the Data transfer is slow for a specific application, it may of course be due to network congestion, retransmission issues, but also to TCP errors like 0-Windows. By looking at specific conversations, you can view whether the TCP window is being reduced and by whom (client / server).
- Abnormal behaviors: by sorting the TCP events by number of SYN packets, you can easily view which machines are generating a very high volume of TCP session start, which eventually do not drive to a complete TCP session setup. If you see, machines with large volume of SYN packets and few / no session setup, this machine is either misconfigured or infected.

This view can be accessed through Diagnostic / TCP events.

ICMP errors

Performance Vision provides in depth view of ICMP errors. When conducting a troubleshooting, this view can display: ICMP errors will report the volume of flows which cannot be setup (either because the network, host, or port is unreachable). This can reveal:

The figure shows a screenshot of a 'DNS Messages' table. The table lists various DNS messages with columns for Time, Filter, VLAN, Device ID, Requester IP, Requester Zone, Server IP, Server Zone, Request Name, Request Type, Response Code, Packets, Bytes, DNS RT, and Pops. The table is sorted by the number of packets in descending order.

Time	Filter	VLAN	Device ID	Requester IP	Requester Zone	Server IP	Server Zone	Request Name	Request Type	Response Code	Packets	Bytes	DNS RT	Pops
2012-03-21 13:14:39	none	-	enb3	172.16.1.119	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	45 847	4.9 KiB	-	-
2012-03-21 13:14:39	none	-	enb3	172.16.1.131	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	38 777	3.4 KiB	-	-
2012-03-21 13:14:37	none	-	enb3	172.16.1.132	Main office (Belgium)	172.16.1.255	Main office (Belgium)	[1]	[32] WILDCARD	na	38 722	3.4 KiB	-	-
2012-03-21 13:14:33	none	-	enb3	172.16.1.132	Main office (Belgium)	172.16.1.255	Main office (Belgium)	[1]	[32] WILDCARD	na	38 281	3.4 KiB	-	-
2012-03-21 13:14:33	none	-	enb3	172.16.1.131	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	38 035	3.3 KiB	-	-
2012-03-21 13:14:31	none	-	enb1	172.16.1.91	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	28 856	1.9 KiB	-	-
2012-03-21 13:14:33	none	-	enb1	172.16.1.97	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	19 516	1.7 KiB	-	-
2012-03-21 13:14:36	none	-	enb1	172.16.1.119	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	15 813	1.4 KiB	-	-
2012-03-21 13:14:30	none	-	enb3	172.16.1.138	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	15 807	1.4 KiB	-	-
2012-03-21 13:14:32	none	-	enb1	172.16.1.138	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	11 940	1.0 KiB	-	-
2012-03-21 13:14:30	none	-	enb3	172.16.1.137	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	10 110	900.0 KiB	-	-
2012-03-21 13:14:29	none	-	enb2	172.16.1.91	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	7 750	702.5 KiB	-	-
2012-03-21 13:14:31	none	-	enb3	172.16.1.113	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	6 856	433.4 KiB	-	-
2012-03-21 13:37:18	none	-	enb1	172.16.1.113	Main office (Belgium)	172.16.1.24	Main office (Belgium)	updat	[1]	[2] ServFail	6 854	433.4 KiB	-	-
2012-03-21 13:37:18	none	-	enb2	172.16.1.132	Main office (Belgium)	172.16.1.255	Main office (Belgium)	[1]	[32] WILDCARD	na	5 864	535.8 KiB	-	-
2012-03-21 13:16:37	none	-	enb2	172.16.6.108	Twin Belgium	172.16.1.13	dnstool	dnstool	[1]	[2] ServFail	5 140	225.8 KiB	-	-
2012-03-21 13:14:30	none	-	enb2	172.16.1.135	Main office (Belgium)	172.16.1.255	Main office (Belgium)	dnstool	[32] WILDCARD	na	4 458	403.6 KiB	-	-
2012-03-21 17:18:33	none	-	enb2	172.16.1.113	Main office (Belgium)	172.16.1.24	Main office (Belgium)	dnstool	[1]	[2] ServFail	3 328	237.4 KiB	-	-
2012-03-21 17:18:33	none	-	enb2	172.16.1.113	Main office (Belgium)	172.16.1.10	dnstool	dnstool	[1]	[2] ServFail	3 328	237.4 KiB	-	-

IP	Application	Traffic	Payload	Packets	Conn. attempts	Conn. established	0-Win	RO indic cl	RO indic srv	Dup.ack	Sess.end	Num.timeout	Client RST	Server RST	
	Web test	1.1 MiB	828.6 KiB	5 042	228	228	255	-	-	-	36	228	94	334	2
39	https	320.5 KiB	192.1 KiB	1 885	126	-	248	-	-	-	0	126	0	247	0
100	https	1.4 MiB	1.2 MiB	4 190	76	-	12	-	< 1 ms	-	0	84	12	46	0
32	https	1.3 MiB	1.0 MiB	4 217	53	-	1	-	< 1 ms	-	0	53	0	53	0
4	NC top	210.1 KiB	160.7 KiB	832	30	30	24	-	-	-	16	30	10	24	6
	https	200.4 KiB	157.9 KiB	606	22	22	21	-	-	-	1	22	19	21	0
	NC top	124.9 KiB	99.9 KiB	424	12	12	12	-	3 ms	-	2	12	7	12	0
	https	29.9 KiB	23.3 KiB	96	6	-	10	-	-	-	1	5	1	10	0
4	https	445.2 KiB	407.9 KiB	699	14	14	9	-	2 ms	-	4	14	1	9	0
4	NC top	2.4 MiB	2.2 MiB	2 819	8	8	9	-	-	-	4	8	3	8	0
9	https	25.1 KiB	18.9 KiB	101	7	7	7	-	-	-	0	7	7	7	0
0	http	724.7 KiB	656.7 KiB	1 127	45	45	7	1 ms	-	-	11	45	1	7	0
0	http	724.7 KiB	656.7 KiB	1 127	45	45	7	1 ms	-	-	10	45	1	7	0
4	https	62.3 KiB	61.8 KiB	166	7	7	7	-	-	-	1	7	8	7	0
9	https	38.2 KiB	29.0 KiB	145	9	9	7	-	-	-	0	9	7	7	0
1	http	129.8 KiB	121.3 KiB	143	2	2	6	-	-	-	0	2	0	6	0
23	https	177.0 KiB	159.9 KiB	335	12	12	6	-	-	-	7	12	0	6	0
23	http	803.2 KiB	725.7 KiB	1 277	56	0	-	-	-	-	18	56	0	6	0
4	https	123.5 KiB	105.2 KiB	301	12	12	6	-	-	-	7	12	0	6	0
6	CRM SPOC	210.9 KiB	184.2 KiB	594	6	6	6	-	-	-	20	6	1	6	0

IP	Application	Traffic	Payload	Packets	Conn.attempts	Conn.established	0-Win
39	Web test	1.1 MiB	828.6 KiB	5 042	228	228	255
35	https	320.5 KiB	192.1 KiB	1 885	126	-	246
32	ftp	12.4 KiB	1.3 KiB	187	42	-	82
100	https	4.7 MiB	4.4 MiB	4 640	0	-	51
100	https	4.2 MiB	3.9 MiB	4 487	0	-	51
100	NC top	210.1 KiB	160.7 KiB	832	30	30	24

- An unavailable host
- A network which is not reachable (either it does not exist - which reveals a configuration / infection issue on the source host, or it is not available - configuration issue?)
- A port which is not reachable (either the source machine is scanning or it is misconfigured and tries to reach a service which no longer exists or has been migrated).

This view is great to pinpoint configuration and infection issues.

ID	Endpoint Zone	Packets / Bytes	Type	Code	Src event ID	Src event name	Dest event ID	Dest event name	Event type / Protocol
225	Main office (Network)	8 827 / 391.3 KiB	Time Exceeded (0)	Time to Live exceeded in Transit (0)	172.16.1.20	Main office (Network)	172.16.1.254	Rome (Network)	161 udp
225	Main office (Network)	8 804 / 390.0 KiB	Time Exceeded (0)	Time to Live exceeded in Transit (0)	172.16.1.20	Main office (Network)	172.16.1.254	Rome (Network)	161 udp
51	Main office (Network)	195 / 122.2 KiB	Destination unreachable (0)	Host unreachable (0)	172.16.1.48	Main office (Network)	16.16.1.1	Headquarters (Network)	no icmp
254	Main office (Network)	776 / 53.2 KiB	Time Exceeded (0)	Time to Live exceeded in Transit (0)	172.16.8.200	Digipolice (Network)	172.16.7.254	Rome (Network)	no icmp
254	Main office (Network)	776 / 53.2 KiB	Time Exceeded (0)	Time to Live exceeded in Transit (0)	172.16.8.200	Digipolice (Network)	172.16.7.254	Rome (Network)	no icmp
225	Main office (Network)	702 / 49.1 KiB	Time Exceeded (0)	Time to Live exceeded in Transit (0)	172.16.1.52	Main office (Network)	172.16.7.254	Rome (Network)	no icmp
161.49	SRV LEO	602 / 247.1 KiB	Destination unreachable (0)	Port unreachable (0)	172.16.1.24	Main office (Network)	162.168.161.49	SRV LEO	137 udp
161.49	SRV LEO	686 / 245.4 KiB	Destination unreachable (0)	Port unreachable (0)	172.16.1.24	Main office (Network)	162.168.161.49	SRV LEO	137 udp

This view can be accessed through Diagnostic / TCP events.

INDEX

A

Aggregation, [11](#), [77](#)
Aggregation period, **83**
Alerting, [39](#)
Application, [6](#), [32](#), [65](#), [70](#), **83**
Application NC, **83**
Application Port Range, **83**
Application Signature, **83**
Autocap, [56](#)

B

BCA, [34](#), [43](#)
BCN, [35](#), [44](#)
Browser, [77](#)
Business Critical Application, [34](#), [43](#)
Business Critical Network, [35](#), [44](#)
Byte, [5](#)

C

Client, [8](#)
Connection Time (CT), **83**
Conversation, [7](#), [77](#), **83**

D

Dashboard, [43](#), [44](#), [48](#), [51](#)
Data Transfer Time, **83**
Deduplication, [17](#), [78](#)
Delta sessions, **83**
Destination, [8](#)
Device Identifier, **83**
Distributed Architecture, [19](#)
DNS, [75](#), [78](#)
DTT, [61](#), [77](#)

E

Email, [39](#)
End User Response Time, **83**
EURT, [50](#), [60](#)

F

Fallback, [5](#), **84**
Flow, **84**

I

ICMP, [72](#)

Initial Sequence Number, [78](#), **84**
IP merging, [7](#)

J

Jitter, [47](#), **84**

K

KiB, [5](#)

L

Language, [29](#)
License, [25](#)

M

Matrix, [10](#)
MiB, [5](#)
Mirroring, [14](#), [17](#)
MOS, [45](#), [47](#)

O

Observation period, **84**
Open Source, [79](#)

P

Packet Analysis, [56](#)
Packet Loss, [47](#)
PCAP, [56](#)
PDF, [36](#)
Promiscuous mode, [24](#)
Protocol, [15](#), **84**
Pulsar, [26](#), [27](#)

R

Report, [36](#)
Reset, [71](#)
Restore, [27](#)
Retransmission, [71](#), **84**
Retransmission Delay, **84**
Retransmission Duplicate ACK, **84**
Retransmission Rate, **84**
Retransmission Total, **84**
RFC
 RFC 1034, [75](#)
 RFC 1035, [75](#)
 RFC 1918, [5](#)

RFC 3261, [45](#)

RFC 3435, [45](#)

RFC 3550, [45](#)

RFC 3551, [45](#)

RFC 3605, [45](#)

RFC 3927, [5](#)

RFC 4193, [5](#)

RFC 4291, [5](#)

Round Trip Time, [84](#)

RST, [71](#)

RTCP, [44](#)

RTP, [44](#)

RTT, [60](#)

S

Server, [8](#)

Server Response Time, [84](#)

Session, [79](#), [84](#)

Shell, [26](#)

Signature Dynamic port, [84](#)

Signature Web application, [84](#)

SIP, [44](#)

SNMP, [39](#)

Source, [8](#)

SRT, [60](#), [77](#)

Subnet, [85](#)

Support, [28](#)

T

TAP, [15](#)

TCP, [54](#), [78](#), [79](#)

TCP Handshake, [85](#)

Tcpdump, [56](#), [78](#)

Time To First Byte, [85](#)

Timeout, [85](#)

U

Upgrade, [25](#)

User, [29](#)

V

VMWare, [22](#), [87](#)

Voice Quality, [45](#)

VoIP, [44](#)

VPN, [28](#)

Z

Zone, [5](#), [30](#), [85](#)