

DeepSweep™

VoIP Surveillance Modules

VoIP Controller

VoIP Content

User's Manual

May 2007

Copyright © IP Fabrics, Inc. 2007

IP Fabrics Corporate Headquarters
14964 NW Greenbrier Parkway
Beaverton, OR 97006
Telephone (main line): 503 444-2400
Telephone (FAX line): 503 444-2401
Website: <http://www.ipfabrics.com>

Information in this document is furnished in connection with IP Fabrics products. No license, express or implied, to any intellectual property rights is granted by this document. This document and the software described in it are furnished under license and may only be used or copied in accordance with the terms of the license.

Copyright © 2007, IP Fabrics, Inc. All rights reserved.

Packet Processing Language™, PPL™ and PPL-VM™ are owned and copyrighted by IP Fabrics, Inc.

Microsoft®, Windows® and Windows® XP are registered trademarks of Microsoft Corporation.

Linux® is a registered trademark of Linus Torvalds.

Red Hat® is a registered trademark of Red Hat, Inc.

RedBoot™ is a trademark of Red Hat, Inc.

MontaVista® is a registered trademark of MontaVista Software Inc.

Intel®, XScale® and Pentium® are registered trademarks of Intel Corporation.

Java™ is a trademark Sun Microsystems, Inc.

Table of Contents

1	Introduction.....	4
1.1	Implementation note.....	4
1.2	Overview.....	4
1.3	Cases and Subjects, and Sessions.....	5
2	Browser Pages.....	6
2.1	VoIP Controller configuration.....	6
2.2	VoIP Content configuration.....	9
2.3	T1.678 messages vs. normal DeepSweep “hit” actions.....	10
2.4	SM Statistics.....	11
2.5	T1.678 Messages sent from DeepSweep.....	11
3	VoIP Controller SM Logic.....	13
3.1	Handling SIP message bodies.....	13
3.2	CCOpen.....	13
3.3	CCClose.....	13
3.4	CCUnavailable.....	14
3.5	DTMF Processing in the VoIP Controller SM.....	14
4	VoIP Content SM Logic.....	15
4.1	Content Intercept.....	15
4.2	DTMF Intercept.....	15
5	Time.....	16
5.1	Time format.....	16
5.2	Absolute Time Accuracy.....	16
6	Other CII and CC Interface Considerations.....	17
6.1	CII Messages.....	17
6.2	CC Messages.....	18
7	RFCs Supported.....	19
8	Example VoIP topologies.....	20
8.1	Single port with aggregated SIP+RTP from one source.....	20
8.2	Dual port with aggregated SIP and RTP from separate sources.....	21
8.3	Quad port with aggregated SIP and tapped RTP.....	22
8.4	Quad port with SIP and RTP tapped, separate sources.....	23

Table of Figures

Figure 1.	High level system showing relationship between SM types.....	5
Figure 2.	“VIP1” - Controller SM definition screen.....	6
Figure 3.	“VIP3” - New SubjectID definition.....	8
Figure 4.	“VIP2” - Content definition screen.....	10
Figure 5.	Single input port.....	20
Figure 6.	Dual input ports.....	21
Figure 7.	Three input ports.....	22
Figure 8.	Four input ports.....	23

1 Introduction

This document describes the two Surveillance Modules for CALEA “broadband” use. These are referred to as:

- voip_controller_sm
- voip_content_sm

Detailed information on how to create SMs, SM actions, and creating surveillance assemblies (SAs) can be found in the DeepSweep™ User's manual. This document assumes those concepts are understood by the reader. The reader will also find several example VoIP connection topology diagrams included in Section 8 of this document.

1.1 Implementation note

This document refers to several items that are not supported in the current release of the product. None of these restrictions are expected to limit the adherence of DeepSweep to the required standard but you should consider these aspects prior to initial deployment and work with IP Fabrics to ensure compliance.

- Support for multiple communicating DeepSweep systems.
- Full support for IPv6 addresses.

1.2 Overview

Figure 1 defines the architecture and external user interface of that part of DeepSweep that provides support for lawfully authorized electronic surveillance of SIP-based VoIP traffic. This support is consistent with the ATIS-PP-1000678.2006, version 2 standard. This document will refer to this standard as “T1.678”. DeepSweep is fully compliant with this standard for SIP-based VoIP traffic using the (non-mapped) Direct Signal Reporting (DSR) method. DeepSweep does not support H.323.

DeepSweep provides VoIP CALEA intercept completely self-contained. That is, it does not rely on separate CALEA support in softswitches, SIP servers, routers, etc., nor does it rely on separate probes and mediation/delivery systems.

The capability is provided in DeepSweep by two surveillance modules (SMs) as shown below. The rationale for providing the capability in two SMs is the following:

- The two SMs will likely listen on different networks in some service providers.
- Having multiple SMs allows the work to be spread over multiple PIXLs (Packet Inspection Accelerators).
- In some complex ISP networks, the two SMs may need to reside in different DeepSweep systems (future consideration).

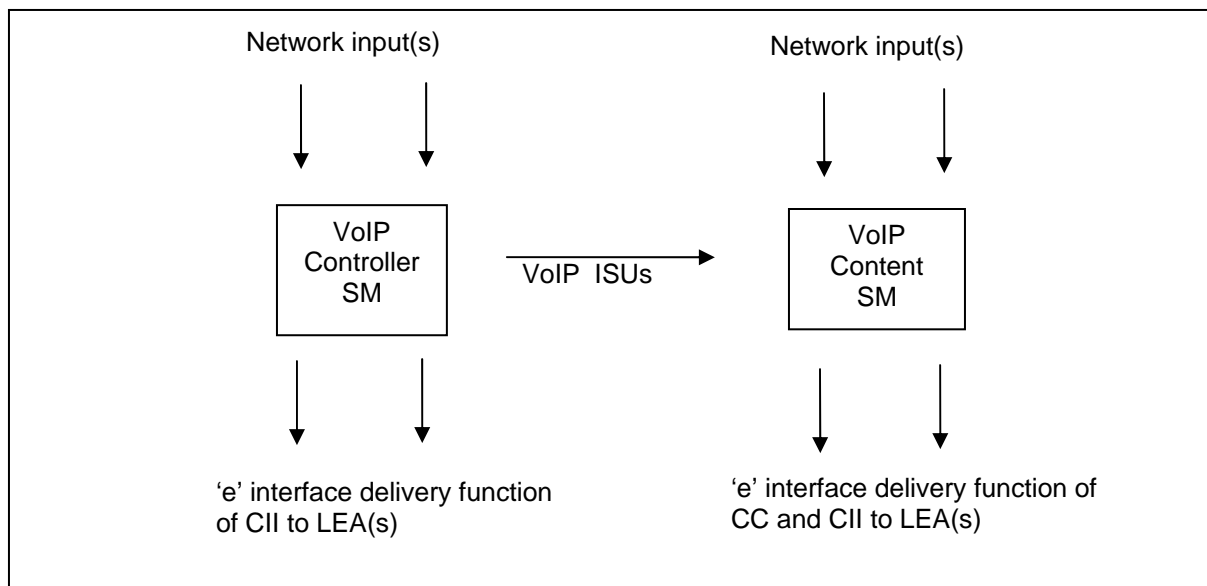


Figure 1. High level system showing relationship between SM types

In a nutshell, the controller SM watches SIP traffic for the active intercept cases. If content surveillance is authorized for a case, the content SM transmits the content (RTP packets) as CC. The content SM needs to be active even if no content surveillance is occurring so that it can detect DTMF signaling in the RTP stream.

Thus the chain that the controller SM is in must be connected to networks containing the SIP traffic. The chain the content SM is in must be connected to networks where the subjects' content traffic (e.g., voice) flows. There can be multiple content SMs per controller SM, and they can be in the same and/or different DeepSweep systems than the controller SM. (Implementation note: Multiple communicating DeepSweep systems feature is not available in the current release.)

The architecture allows multiple intercepts to be active simultaneously; multiple intercepts on behalf of different LEAs (law enforcement agencies), including multiple intercepts on the same subject (aka subscriber or target); and adding or deleting intercept cases without interrupting ongoing intercepts.

1.3 Cases and Subjects, and Sessions

It is important first to understand the relationships among *cases*, *subjects* and, *subject ids*.

Case	Typically a court order authorizing surveillance, typically of a single subject.
Subject	A term used loosely herein. Typically a person to which a case applies. In T1.678 the terms subject and subscriber are used interchangeably. There can be more than one case that involves the same subject.
Subject ID	A specific network identification of a subject. A subject can be known by multiple IDs, and thus a case can typically define multiple subject IDs. A subject ID can be associated with the person (e.g., user@hostname) or with equipment associated with the person (e.g., phone number). Because subjects can be in multiple cases, so can subject IDs.

2 Browser Pages

In DeepSweep all functions of the Surveillance Modules are configured via browser pages. The following describes the setup details for the pair of SMs for VoIP.

2.1 VoIP Controller configuration

Figure 2 shows the main page for the controller SM. The left shows the list of cases that have been defined. The bottom left side shows the information about the selected case. The right side contains attributes about the controller SM as a whole.

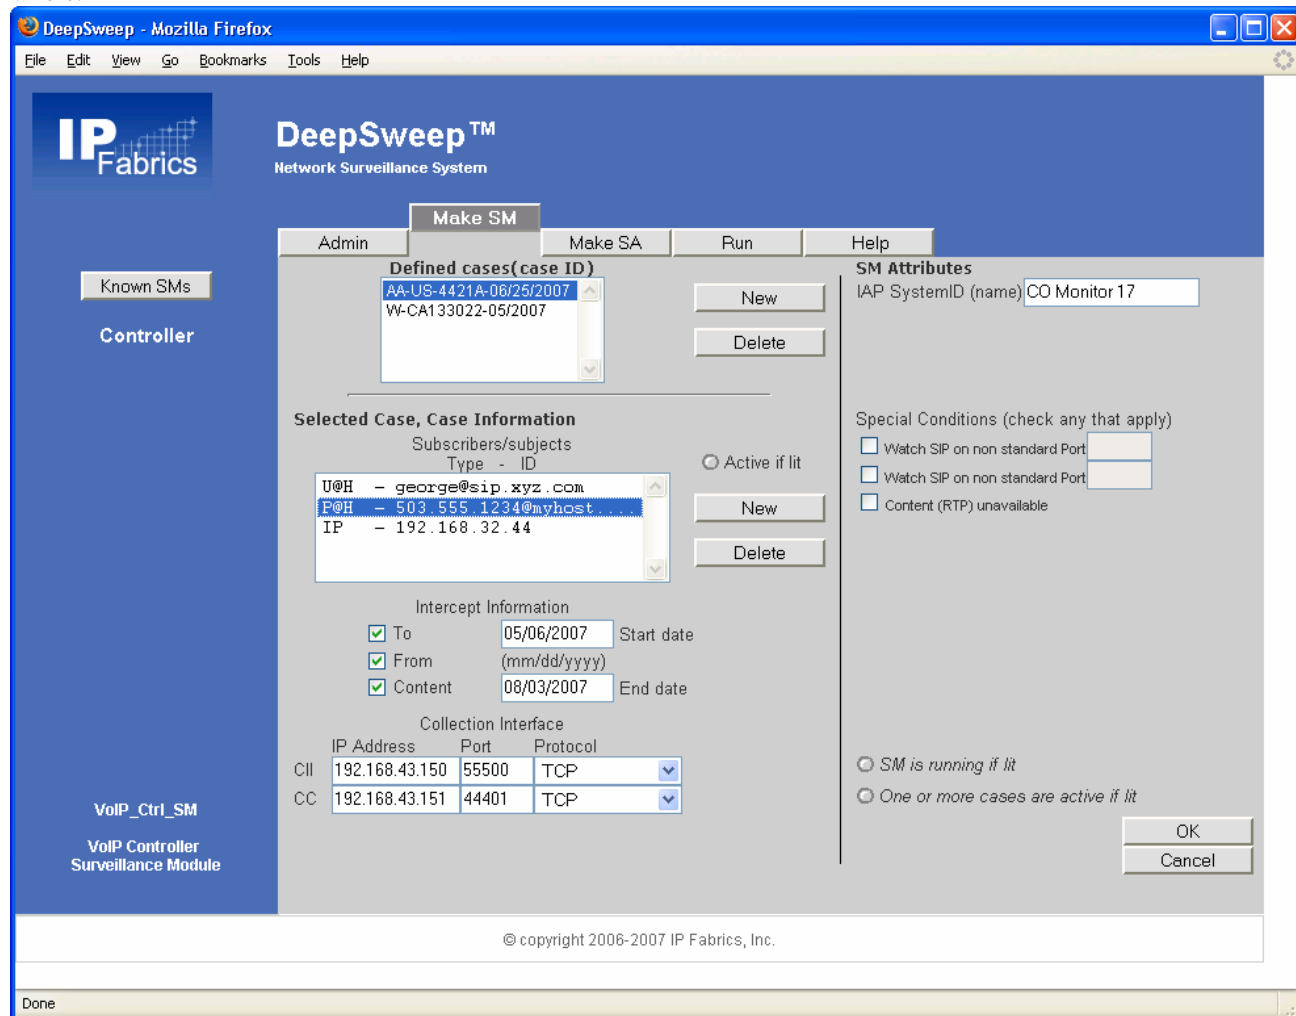


Figure 2. “VIP1” - Controller SM definition screen

The upper left has a scroll box that shows the name(s) of all cases that have been defined to the DeepSweep system. A case is identified by a case ID, which is a 1-25 character string. Depressing the NEW button brings up a box that asks you for a 1-25 character case ID. Providing that the case ID is different from all existing case ID's, depressing OK in that box brings you back to VIP1, where you can then describe the case beneath. Depressing DELETE next to the case scroll box causes the entire case to be deleted. If the case is not active, the system will prompt the user for confirmation. If the

case is active, the system will prompt the user with stronger wording, because deleting a case while it is active is unusual and serious.¹

The lower left side of the screen shows the definition of the selected case, or in the case of a newly created case, is blank. An indicator shows whether the case is active, meaning that surveillance is active. By definition, if this indicator is lit, the two indicators on the lower right are also lit.

A scroll box shows the subject IDs that are part of the case. DeepSweep provides for any number of subject IDs per case because a subject may be known to the network in multiple ways. Subject ID can be a number of things. Although the system could deduce its meaning from the format of the entered ID, it requires the user to express an ID type. This allows the system to do a better job of flagging incorrect IDs. The following table shows the ID formats and how they are interpreted.

ID Type	Subject ID format	Matches
U@H	user@hostname	sip:user[:password]@hostname[:][?] sips:user[:password]@hostname[:][?]
U@I	user@ip_address	Same as above with IP address in place of host name. IP address may be IPv4 or IPv6 (dot notation, colon notation, compressed notation, IPv4 within IPv6 notation)
P@H	phone_number@hostname	Same as first case with phone number in place of user name
P@I	phone_number@ip_address	Same as second case with phone number in place of user name
P@?	phone_number	Same as P@H with any host name or P@I with any IP address ²
HOST	hostname	sip:hostname[:][?] sips:hostname[:][?]
IP	ip_address	Same as above with IP address in place of host name
TEL	phone_number	tel:phone_number

The following notes apply:

User	If <i>user</i> is all lower case, user part match is case insensitive. If one or more upper case letters in <i>user</i> , must match case exactly. ³
Ip_address	DeepSweep accepts IPv4 “dot” notation, IPv6 colon notation, IPv6 compressed notation, and mixed notation (IPv4 address within an IPv6 address).
Hostname	Host or domain name (e.g., comcast.net). Upper/lower case insensitive.
Phone_number	Series of numeric characters (0-9), optionally started with “+” and optionally containing the visual separator characters hyphen, dot, open parenthesis, and closed parenthesis [-.().]. I.e., +15034442444 and +1(503)444-2444 are equivalent.
Password	Ignored in the CII if present (ignored for purposes of matching the subject ID).

¹ An active case is one for which surveillance is currently underway. To be active, the case must have at least one subject ID, must have its collection interface(s) defined, and the current date must be greater than the start date (if the start date is not blank) and not greater than the end date (if the end date is not blank).

² The P@? ability to describe a subject by just the phone-number part is particularly useful for incoming (to the subject) calls where the host name or IP address may not be statically known.

³ In SIP URIs, unlike other URIs, the user part of the URI is supposed to be case sensitive. DeepSweep provides both means, except that it provides no case-sensitive comparison if it is trying to match a subject with all lower case user letters.

- Info beyond host** If the CII contains SIP URI information to the right of the host (port number, uri-parameters, and/or headers), they are ignored for purposes of matching the subject ID
- % HEX HEX** The % character followed by two characters in the range 0-9,A-F may appear at any point. This is used to insert UTF-8 or other encodings. For instance the name "x%C3%B3y" puts the two hexadecimal byte values C3 and B3 between the characters x and y.
- tel:** Signifies a telephone URI (e.g., tel:911).

Depressing the NEW button next to the subjects scroll box brings up a box that asks you for the subject ID. Providing that the subject ID is different from all existing subject ID's in this case and that its format is consistent with the previous table, depressing OK in that box brings you back to VIP1. Adding a subject ID to an active case will cause that subject ID to be active immediately.

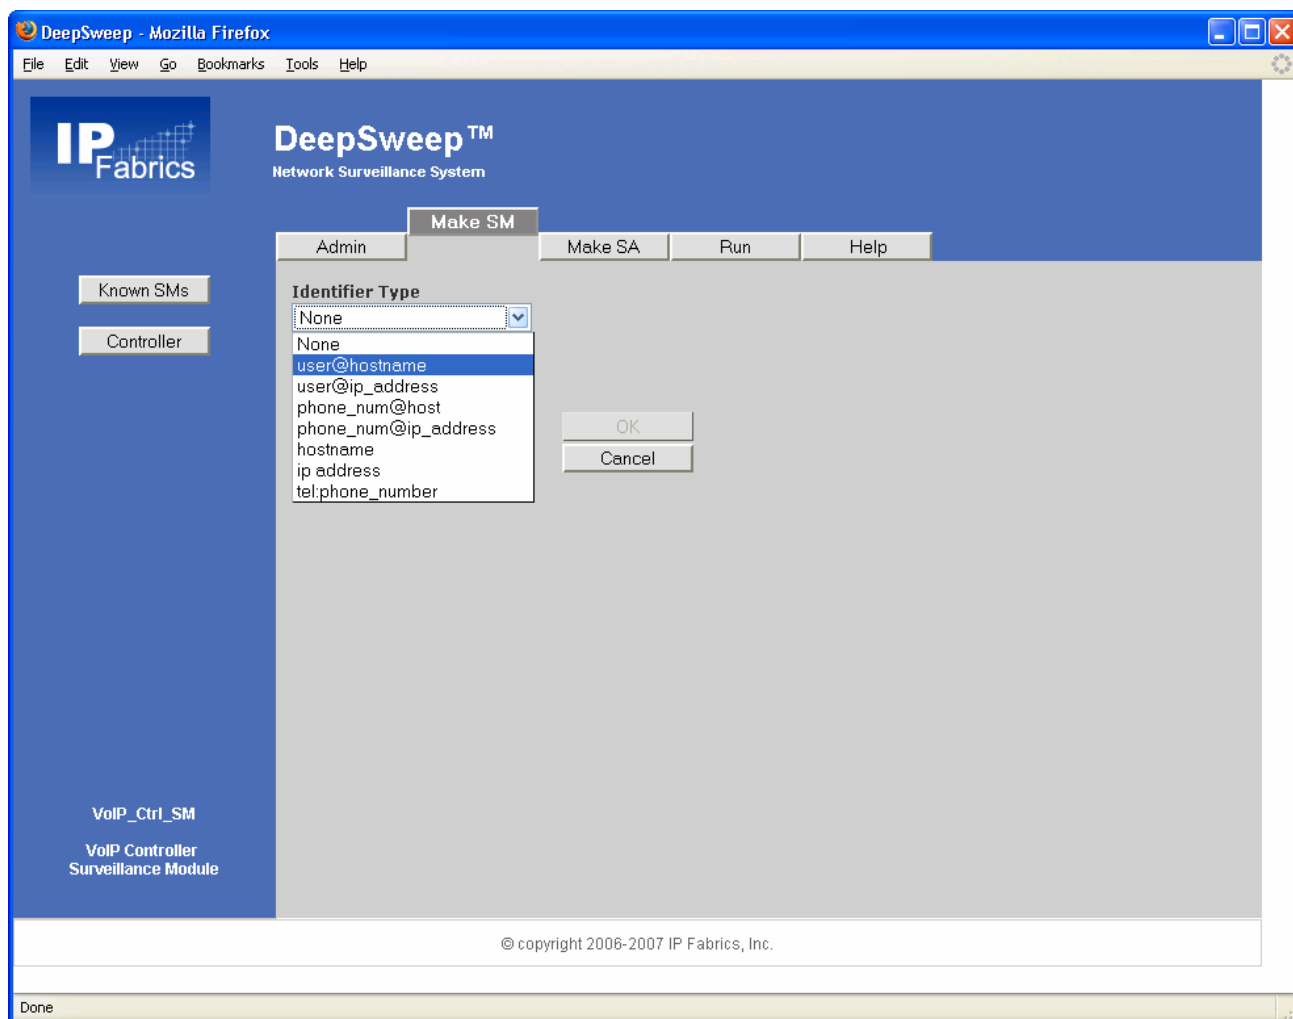


Figure 3. "VIP3" – New SubjectID definition

Returning to page VIP1, depressing DELETE next to the subject scroll box deletes the subject. If the case is active, the system will prompt the user for confirmation. Deleting a subject ID from an active case causes intercept related to that specific subject ID to stop for this case.

The next three check boxes define the type of intercept authorized.⁴ At least one of TO and FROM needs to be checked.

The next two fields are the dates, in the time zone of the DeepSweep system, on which intercept is to start and be completed. If start is left blank, it means “immediately.” If end is left blank, there is no automatic cessation of the intercept.

The last piece of information is the collection interface to be used. The protocol and the CII IP address and port is always required. If CONTENT is checked, the IP address and port of the CC interface is also required.

For the protocol, four choices are provided:

- UDP
- TCP
- UDP with appended message digest
- TCP with appended message digest

The last two are not mentioned at all in T1.678 but the system provides these options for consistency with other DeepSweep Surveillance Modules. If an appended message digest is opted for, the system computes a SHA-1 hash of the whole UDP or TCP payload and then appends the first 96 bits of the 160-bit hash to the end of the payload, thus increasing the payload and packet size by 12 bytes.⁵

If a case is active, one is allowed to delete it (with confirmation) and add or delete a subject ID of the case. One is not allowed to change any of the other information for an active case.

Now to the information on the right. Note that this SM page is unusual in that it indicates if this SM is actually running as the page is viewed. Another indicator shows if any of the cases are currently active (intercept is active). The primary purpose of these are to help the user understand what adding or deleting an intercept will mean.

One piece of information – IAP system ID (name) – exists independent of specific intercepts and is communicated when hits occur. It is a name denoting the intercept access point the DeepSweep system running this SM should be known by.

There are also several special conditions that can be checked. In addition to watching for SIP on port 5060, one can specify one or two additional ports to be watched. Also, if the RTP stream is not accessible even though content intercept may be authorized, one should check the appropriate box. Checking this box causes the CCUnavailable message to be generated when content intercept should otherwise occur. If one checks this box, a VoIP content SM is not needed to be present and assumed not present.

2.2 VoIP Content configuration

VIP2 (Figure 4) is the page for the content SM. This is called the VoIP content SM although the specific protocol it supports is RTP. One must specify the name of the associated controller SM. The system is defined this way so that there can be multiple content SMs associated with one controller SM, possibly on different DeepSweep systems. The IAP system id (name) has the same meaning as on the previous page. If the content and controller SMs are in the same DeepSweep system, they should have the same IAP system names. If not, this is the name of the system in which the content SM resides.

⁴ T1.678 doesn't distinguish between the “to” and “from” cases, but doing so allows the system to distinguish pen-register intercepts from trap-and-trace intercepts.

⁵ 96 bits is used because this is commonly done when producing an HMAC. What is defined here is not an HMAC (it is not computed with a secret key), but it would be a logical next step to take in the future to make it a secure HMAC.

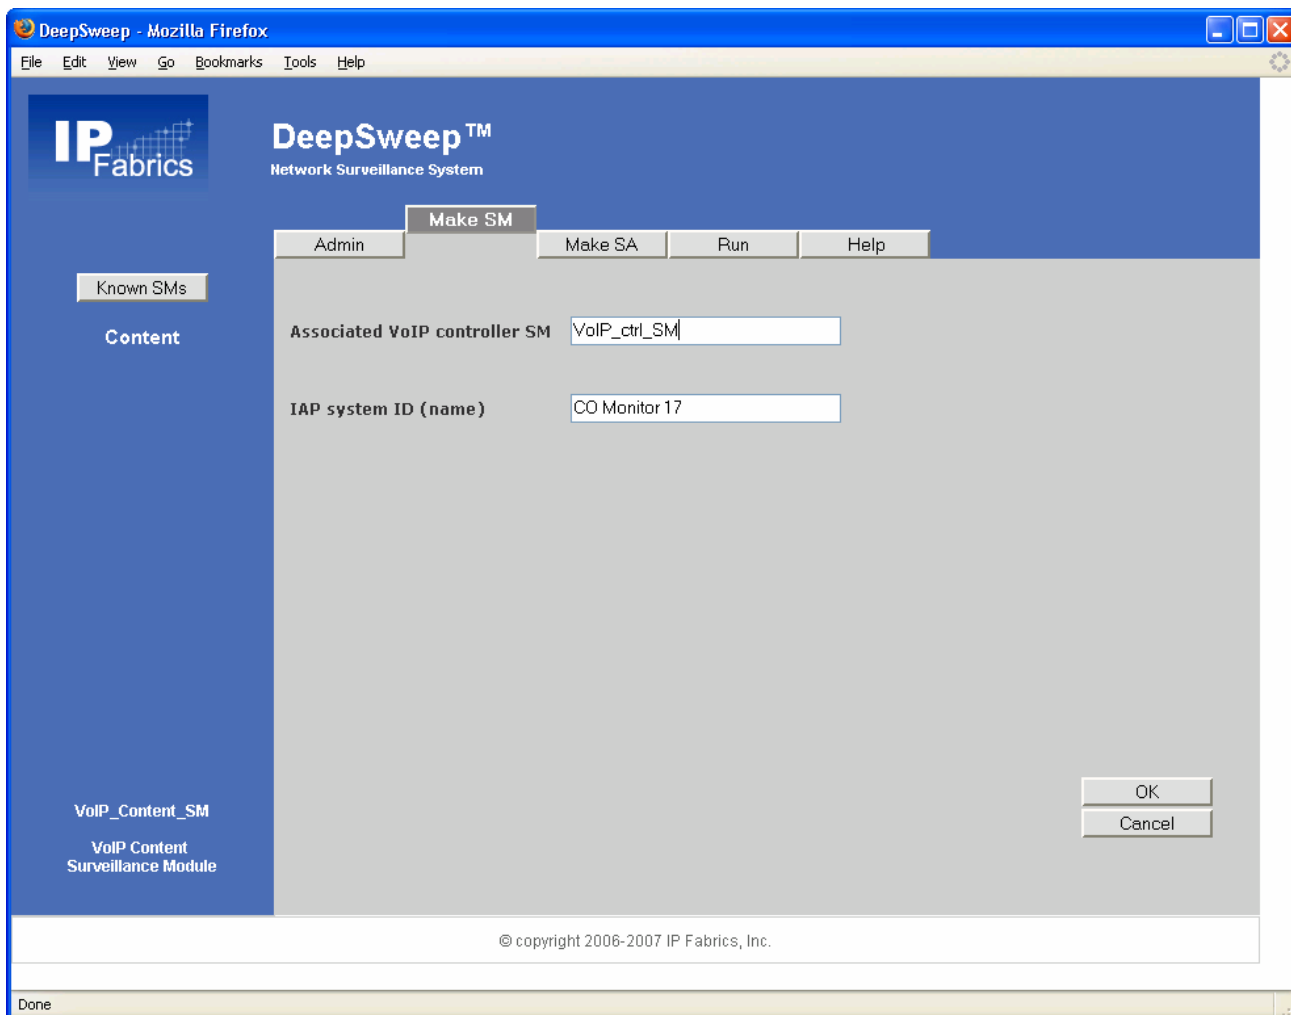


Figure 4. "VIP2" - Content definition screen

The DeepSweep user needs to be careful about the network segments feeding the chains to the controller and content SMs, because VoIP network designs typically use NAT (and this is typically done by the VoIP session border controller).

2.3 T1.678 messages vs. normal DeepSweep “hit” actions

The IAS controller SM doesn't send the detected packet itself as an external message; it sends a T1.678 CII message. Also it sends these messages in situations where there isn't a detected packet (e.g., at startup). The normal SM action options such as record, monitor, SMNP trap, kick to user program and reflect are not appropriate for lawful intercept. That said, Monitor and Record can be useful in confirming that an installation is working properly.

Up to this point, the VoIP SMs have been described more in relation to what T1.678 messages they generate rather than the usual DeepSweep description of “hits.”

It is unlikely that the VoIP SMs will be used in chains with other SMs (other than perhaps using the controller and content SMs in the same chain in the case of a very simple network topology), but it is allowed. Unlike other SMs, the VoIP SMs always pass both hits and misses on to the next SM by default when the chain is being defined and this is be the normal way to use the system.

Although it is valid to have in a system's surveillance assembly a VoIP controller SM without a VoIP content SM (in situations where one is using multiple DeepSweep systems), in most configurations it is an error. Hence, when one says run a surveillance assembly, if the system sees an VoIP controller SM but no VoIP content SM will display a warning message but proceed.

2.4 SM Statistics

The RUN/STATISTICS page (described in the DeepSweep User' Manual) has a statistic for external messages sent. All CII and CC messages sent are counted in this system-wide statistic.

In the DeepSweep architecture, every SM can collect up to four statistical values that are represented on the RUN/STATISTICS page in a 2x2 matrix. The lower right corner is always the number of packets examined by the SM. For the CALEA VoIP controller SM, the statistics should be as follows:

Packets generating one or more CII events.	CII events
	Packets examined

For the content SM, the statistics are:

Packets generating one or more CC events.	CC events
	Packets examined

2.5 T1.678 Messages sent from DeepSweep

T1.678 gives the implementer a basic choice in CII messages (section 6.4). One choice is to map SIP (and H.323, if one is providing that, which DeepSweep does not) into J-STD-025 messages defined for circuit-switched telephony. This is called mapped signaling information, or mapped. The other choice is DSR (direct signal reporting), where most of the signaling is delivered on the CII interface by encapsulating the SIP message. DeepSweep implements DSR. Thus the T1.678 message types that are generated are the DSR and DialedDigitExtraction messages, and, if content intercept is authorized, the CCOpen, CCClose, and CCChange. DialedDigitExtraction comes from the VoIP content SM; the others come from the VoIP controller SM.

Another choice is the format of the CC message (call content). One may use the format defined for T1.678, or the IPCablecom format. DeepSweep implements the former.

There are several other alternatives:

- For correlation (of messages, and of CC with CII), one can use a call identity (e.g., the SIP CALLID) or a unique correlation identifier value. DeepSweep uses the former.
- DeepSweep assumes for each case that there is a distinct IP address and port for the CII and CC (page VIP1).

3 VoIP Controller SM Logic

T1.678 standard contains some optional messages. DeepSweep supports the optional SurveillanceStatus message. When the controller SM starts, it sends one SurveillanceStatus message per active case on the case's CII interface. This message has the case id, the IAP system id, the time, and the status *activation*.

Because the SIP protocol is being watched and because T1.678 DSR messaging protocol is being used, the controller SM does not need to keep much internal state. Note that the system normally watches port TCP, SCTP, and UDP ports 5060 and 5061, and any additional ports indicated on page VIP1.

If a new active case is created, a SurveillanceStatus message is sent at that time as well. If the SA running is stopped, the controller SM sends one SurveillanceStatus message denoting deactivation per active case. Similarly, it sends one if an active case becomes inactive.

A DSR message is sent on the CII interface for every SIP message and response associated with a case/subject, except for "100 Trying" responses. A DSR message contains the following:

- Case id
- IAP system id
- Timestamp
- CALL-ID
- SProtocol specific parameters
- The SIP message or response

3.1 Handling SIP message bodies

If a SIP message contains a message body and the header contains the Content-type field and the type is other than "application/sdp," DeepSweep will not deliver the message body unless content intercept is authorized for the case.

DeepSweep will normally deliver the headers of the MESSAGE message but generally not the body unless content intercept is authorized.

3.2 CCOpen

CCOpen is sent on the CII interface when when the SIP session reaches the point where the media negotiation is completed. Typically this would be when the callee responds with a "200-OK" and specifies the SDP with the actual media parameters chosen. CCOpen contains:

- Case id, IAP system id, timestamp, CALLID
- IP address and port number to where content will be delivered
- CC delivery format (denotes T1.678 versus IPCablecom)
- SDP media information

The other critical thing done at the time of CCOpen is to notify the content SM(s).

3.3 CCClose

The CCClose message is sent when call content delivery has ended. Generally this is triggered by a SIP BYE message.

3.4 CCUnavailable

If content-unavailable is unchecked, the system sends a CCUnavailable instead of the above (CCOpen, CCClose) if content intercept is enabled.

3.5 DTMF Processing in the VoIP Controller SM

SIP provides for many alternate ways for DTMF and other signaling to be communicated, including in NOTIFY and INFO messages, as events in the RTP stream (RFC 2833), as specially coded tones in the RTP stream (RFC 4733), and as sound waves in the RTP stream. Signaling that appears in the messages gets communicated as part of the SIP message and thus doesn't need further consideration. Here DeepSweep provides for communication of DTMF as events in the RTP stream consistent with RFCs 2833 and 4733

The detection of the DTMF events is done by the content SM.

4 VoIP Content SM Logic

The VoIP Content SM is simpler. It has a list of IP addresses and ports on which it should watch for UDP/RTP traffic, along with an indication of whether to watch for “to” (destination), “from” (source), or both. It also needs to know, by address, whether it is to intercept content (yes or no) and DTMF events (yes or no).

4.1 Content Intercept

When DeepSweep gets a hit and content intercept is enabled, the system sends the encapsulated RTP packet to the CC collection interface. The packet has an outer IPv4 (or in the future IPv6) header, with the source address being the DeepSweep and the destination address that specified for the CC interface. The packet is sent via UDP or TCP as specified by the user (VIP2) using destination port also specified there. The payload of the UDP or TCP packet consists of

- Case id
- Content SM's IAP system id
- CALLID
- Time stamp
- Packet direction (to versus from subject)
- IP/UDP/RTP packet

4.2 DTMF Intercept

This may occur together with content intercept or separate. When the Content SM is given the IP address and port (from the controller SM) of an RTP stream to watch, it is also given an indication of whether to watch for DTMF events and, if so, the payload type of the DTMF events and whether these are coupled with the redundant payload type or not.

If DeepSweep is watching for DTMF, the system must examine every RTP packet sent from the subject to determine if DTMF events are there. See the examples in RFC 4733.

Detection of one or more DTMF events in a packet means sending a DialedDigitExtraction message from the content SM. Note that it is sent to the CII interface; this is the only case where the content SM sends to the CII interface. This means that the content SM must be told both the CC and CII interfaces. This “DDE” message contains the normal stuff – case ID, IAP system ID, timestamp, call ID, and one or more DTMF “digits.”

5 Time

5.1 Time format

T1.678, like J-STD-025B, uses a timestamp in all messages, and the ASN.1 syntax defines it as a type called GeneralizedTime, and this is a very specific format that ties the date and time to GMT or an offset thereof. Time is defined to the millisecond.

Once DeepSweep detects a surveillance event, it immediately calculates a time stamp to be used in the message. Delay from detection to time-stamp value should be less than 1 ms. The delay from the event detection to the physical transmission of the message from the DeepSweep system must be less than 1 second 95% of the time.

5.2 Absolute Time Accuracy

DeepSweep uses NTP to keep DeepSweep's concept of time accurate to the NTP time standard; DeepSweep's time should be less than 200 ms inaccurate from absolute time.

6 Other CII and CC Interface Considerations

The transport protocol can be UDP or TCP. To provide a way of supporting the data integrity the TCP and UDP checksum is calculated on either interface.

The only thing sent by the DeepSweep system to the LEA collection interfaces are CII and CC messages for the case(s) defined on browser page for VoIP Controller as being associated with that specific interface. Note that when TCP is selected, this means that the normal TCP protocol handshaking packets also pass over the interface.

6.1 CII Messages

The table below summarizes the CII messages that are generated. For all messages, the parameters labeled M and O in the T1.678 standard are always provided, and the parameters labeled C are provided when known or available (see earlier sections).

CII Message	When	Notes
DirectSignalReporting	Occurrence of any SIP message to or from a subject ID of an active case, except for "100 Trying" responses.	DepSweep never delivers the message body if content-type is present and other than application/sdp unless content intercept is authorized, in which case the system always delivers the message body.
CCOpen	SIP signaling is at a point where the media stream is being enabled, and content intercept is enabled. Normally occurs on a "200 OK" message with the media parameters.	DeepSweep always include the optional encapsulated signaling message parameter. Delivery identifier is CCAddress. CCDeliveryFormat is CCDeliveryHeaderModuleOID.
CCClose	For content intercept being enabled and the intercept being in the open state, occurs when the media stream is being disabled. Normally triggered by a SIP bye.	Has subset of parameters of the matching CCOpen
CCChange	Not used. Although this is listed as a message used with DSR, 6.1.2 of the standard says that it is not required if changed media info is reported another way. The system always reports it in the DirectSignalReporting message above.	
CCUnavailable	Instead of CCOpen (and without CCClose) if the SM notes that no content intercept is available.	
DialedDigitExtraction	When the content SM decodes an RTP payload and discovers one or more DTMF signals encoded as events in a telephone-event payload type.	

All CII messages contain the case ID and IAP system ID, both of which come from information provided on page VIP1. All CII messages also contain a timestamp, which is produced in real time. All the CII messages used contain a CorrelationIdentifier, such that all CII messages related to the same intercept have the same value. DeepSweep uses the SIP CALLID for this.

6.2 CC Messages

The system uses the CCDelivery APDU described in section 6.2.2 of the standard. DeepSweep always includes a CCDelivery header. The standard says this can be omitted if every stream is delivered to a unique collection port and if the original content transport includes timing and sequence (e.g., RTP), but with the ability for cases to be added on the fly there is no foolproof way to ensure that every intercept stream will be going to a unique IP address and port.

The delivery header contains the case ID, content SM IAP ID, call-id, time, packet direction, and a sequence number. The remainder of the message is the encapsulated IP/UDP/RTP packet.

Another content message in the standard is UUContent for content sent in the signaling plane. Section 6.3.1 of the standard allows us to omit this because its use would be redundant with the inclusion of content in the DirectSignalReporting message per the SIP content-type field as described earlier.

7 RFCs Supported

768	User Datagram Protocol
793	Transmission Control Protocol
2198	RTP Payload for Redundant Audio Data
2327	SDP: Session Description Protocol
2373	IP Version 6 Addressing Architecture
2396	Uniform Resource Identifiers (URI): Generic Syntax
2486	The Network Access Identifier
2806	URLs for Telephone Calls
2833	RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals
2960	Stream Control Transmission Protocol
3174	US Secure Hash Algorithm 1 (SHA1)
3261	SIP: Session Initiation Protocol
3264	An Offer/Answer Model with Session Description Protocol (SDP)
3265	Session Initiation Protocol (SIP) – Specific Event Notification
3515	The Session Initiation Protocol (SIP) Refer Method
3550	RTP: A Transport Protocol for Real-Time Applications
4733	RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals

8 Example VoIP topologies

8.1 Single port with aggregated SIP+RTP from one source

Figure 5 is a single PIXL product in the minimal configuration.

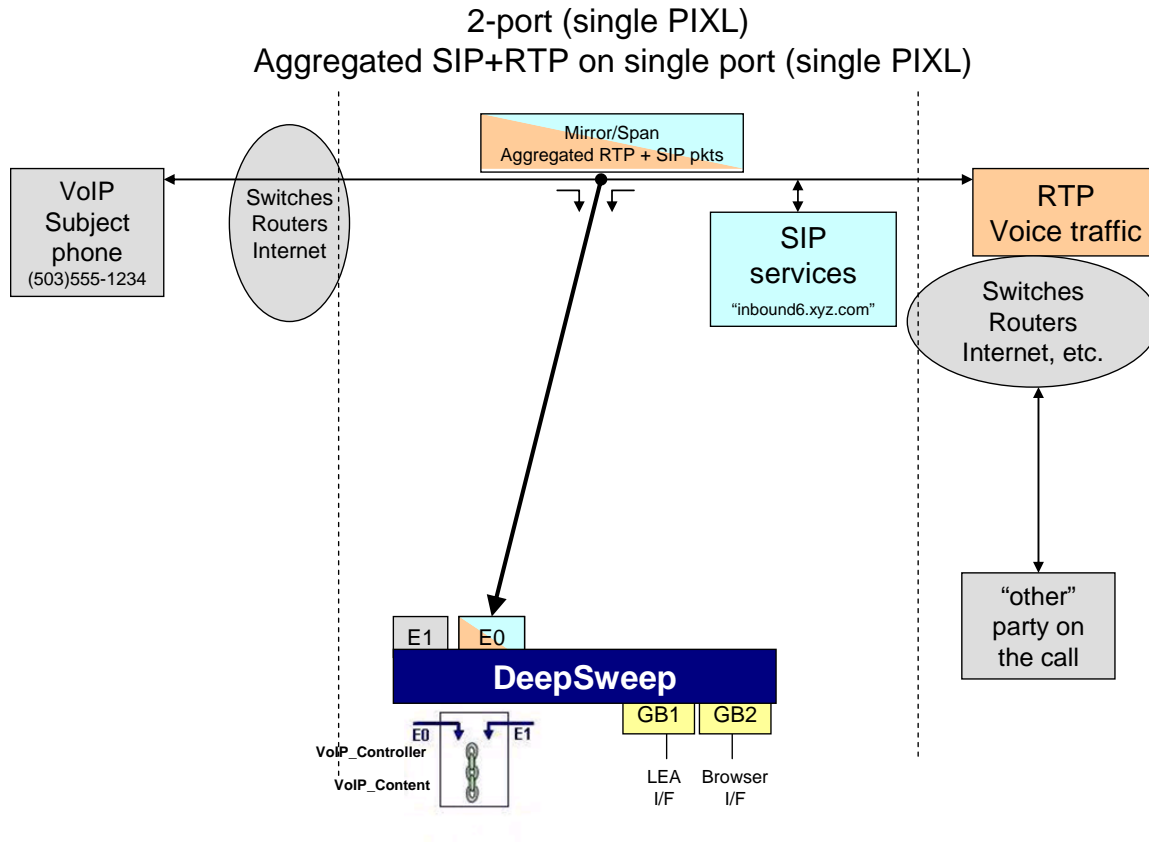


Figure 5. Single input port

8.2 Dual port with aggregated SIP and RTP from separate sources

Figure 6 is a single PIXL product but improved performance over above configuration since the traffic is separated prior to being processed by DeepSweep.

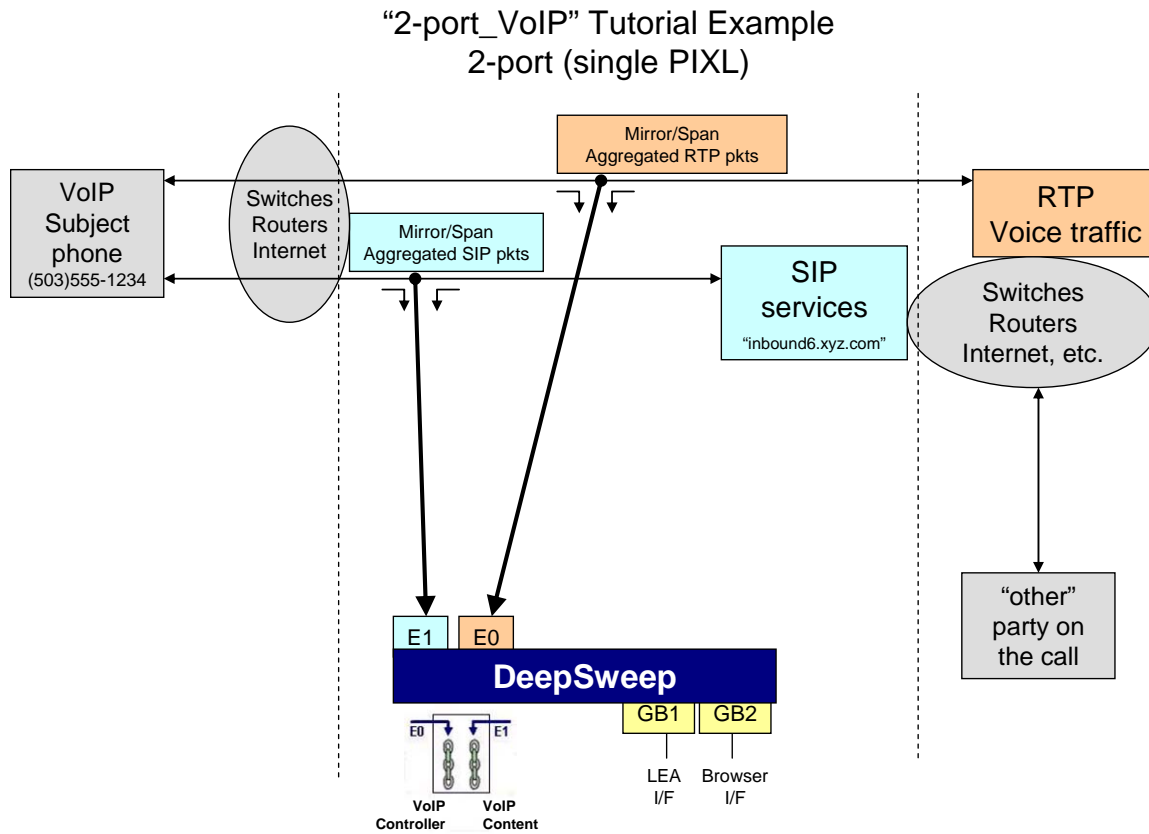


Figure 6. Dual input ports.

8.3 Quad port with aggregated SIP and tapped RTP

Figure 7 is a dual PIXL product.

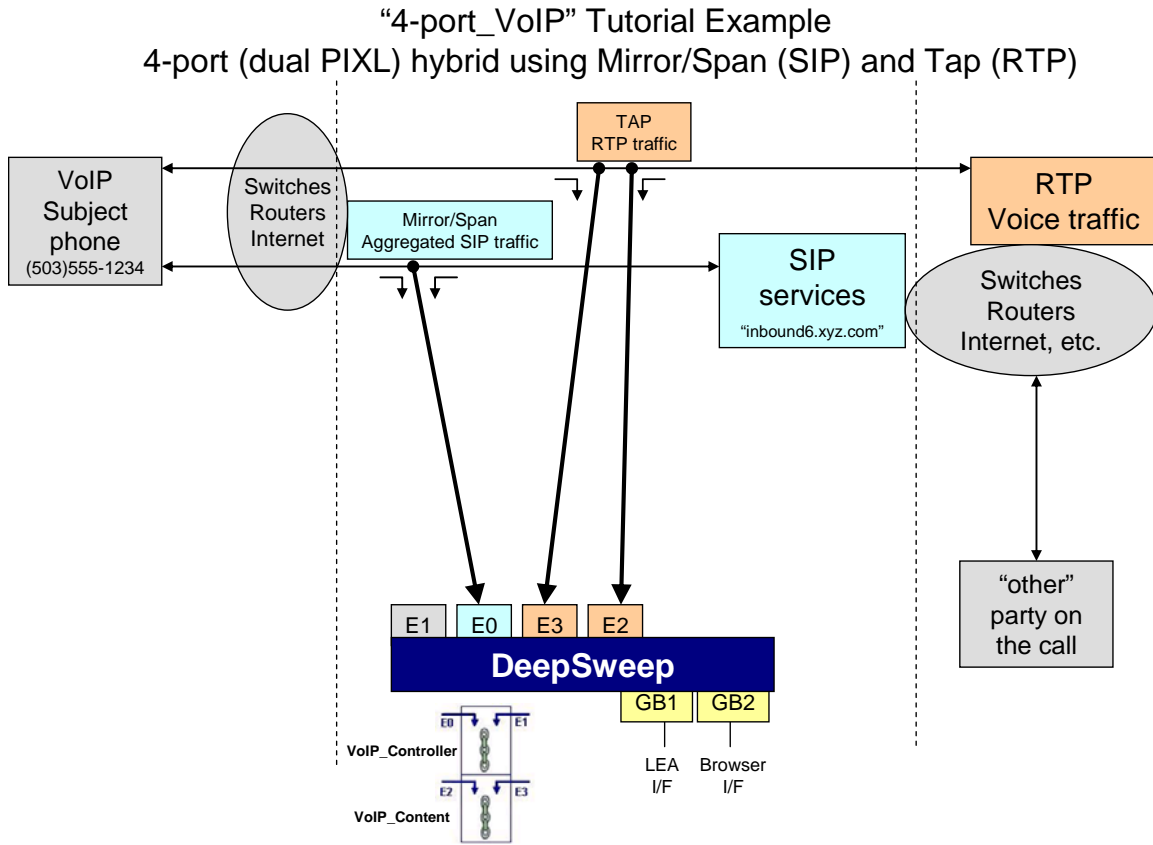


Figure 7. Three input ports.

8.4 Quad port with SIP and RTP tapped, separate sources

Figure 8 is a dual PIXL product.

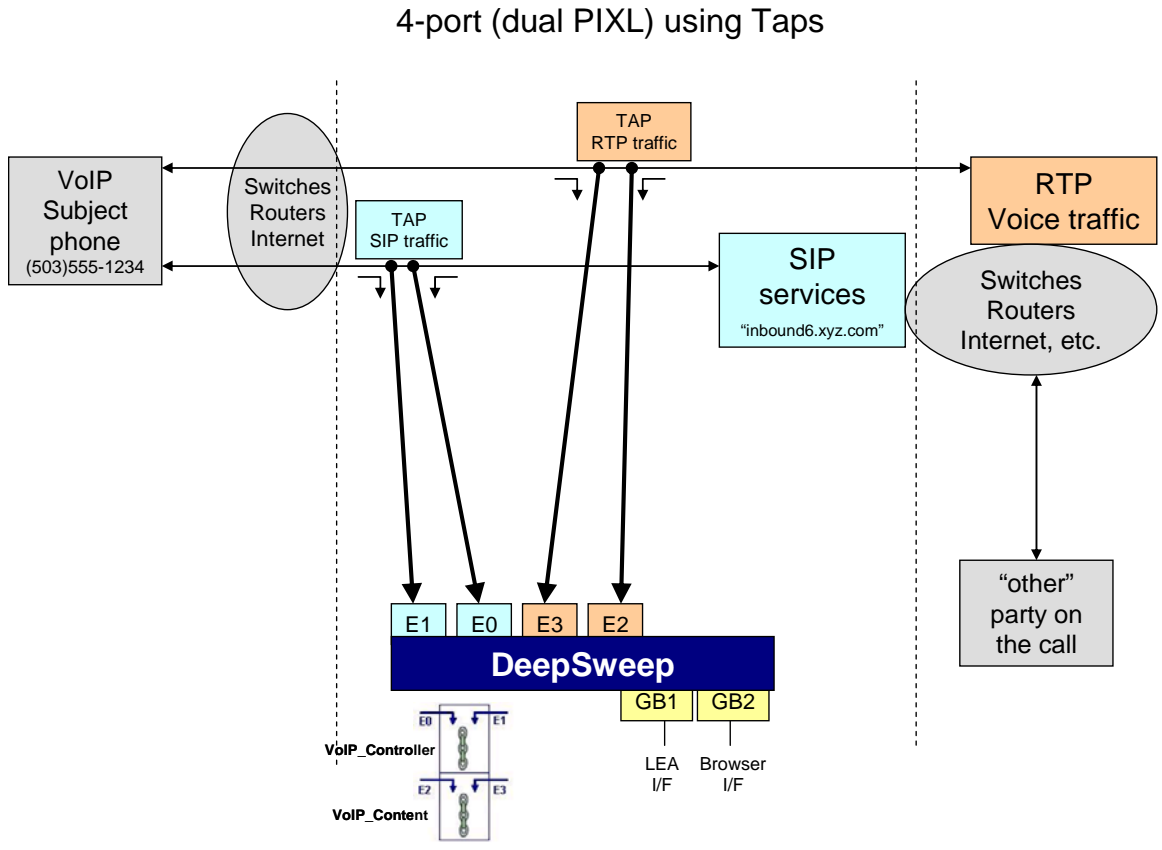


Figure 8. Four input ports.