

Sizing up VoIP listening tools

VoIP traffic analysis products offer an eclectic mix of capabilities and perspective, but no vendor ships a complete tool kit.

■ BY EDWIN MIER, VINCENT BATTISTELLI AND ALAN MINER

With corporate voice-over-IP activity on the rise, network professionals need to know whether network measurement and monitoring tools are available today that can keep pace. The answer, as determined by our test of a half-dozen software and hardware VoIP traffic analysis tools, depends on what specific measuring and monitoring functions you have in mind.

Our objective was to test the growing set of products — nearly two dozen on the market today — that have added VoIP traffic analysis to their repertoire of options. To be included in this test, products had to address the better portion of our seven-point criteria:

- Real-time VoIP traffic monitoring and alarm generation.
- Long-term VoIP activity recording and reporting.
- VoIP traffic generation as used to verify proper call-controller operation, VoIP node availability and VoIP-bandwidth assessment.
- Automated VoIP voice-quality assessment.
- Measurement of VoIP-related quality-of-service (QoS) parameters.
- VoIP traffic and protocol decode.
- Intelligent diagnosis of VoIP service problems.

Acterna, Agilent, Brix, Finisar, NetIQ and Sniffer Technologies, a division of Network Associates, submitted multiple products to address as many of these functional categories as possible. All totaled, more than two dozen discrete products were tested, but none of these combination suites fully addressed all the tasks we defined in our methodology (see Hits and misses chart, page 50).

Products that were either not ready for public testing

or addressed only one or two of these functional areas are discussed in a related story (see www.nwfusion.com, DocFinder: 3423).

The VoIP traffic we applied in this test was real, bidirectional telephone traffic. Phone calls entered and egressed the IP network via several vendors' VoIP gateways or else were soft-phone calls, carrying VoIP end-to-end between laptops equipped with headsets (see How we did it with testbed diagram, DocFinder: 3422).

There are three deployment issues that users need to consider:

- Whether to insert in-line or monitor via a mirrored port.
- Whether to monitor passively or apply active test traffic.
- Whether to measure and monitor from a single point, or multiple remote points.

In-line or mirrored port?

Four of the vendors' products we tested — Acterna, Agilent, Brix and Finisar — offer VoIP-analysis devices that can run in-line, meaning the device is inserted into the direct path of a backbone link so all traffic passes bidirectionally through it.

The alternative is to hang the device off a mirrored switch port, to which a copy of all bidirectional traffic

from a key backbone link is sent. We tested Sniffer Technologies' Sniffer Distributed edition hanging off a mirrored switch port. Acterna, Agilent and Finisar suites also support mirrored-port deployment. Brix offers various flavors of its Verifier monitor modules; the Brix 100 Verifiers can insert in-line. The Brix 1000 and 2500, which generate active test streams, connect to conventional switch ports — not mirrored — the same way as the NetIQ PC nodes, which also actively generate test traffic.

In-line monitoring is more accurate from a timing perspective, and ensures that all traffic in both directions is seen. Mirrored-port monitoring imposes minor timing variations and might miss some traffic under heavy load.

Physically inserting a device in-line needs to be scheduled to minimize backbone disruption.

The devices we tested from Agilent, Brix and Finisar all had fault-tolerant bypass capabilities.

Acterna's DA-3400 unit requires a separate "tap" unit — which we didn't have — to be fault tolerant. Still, Acterna preferred that we test its system in-line, and we encountered no issues.

However, the late beta code for Agilent's Telephony Network Analyzer software crashed periodically during testing, hanging up the in-line monitor device, before the bypass circuitry failed over. The delay was brief, typically seconds, but long enough to drop all active VoIP calls traversing our network.

The bottom line is go in-line if you can, but be aware that there's a risk.

Some of the capabilities we examined — such as monitoring quality of service (QoS) parameters or assessing VoIP voice quality — can be accomplished by passively observing user traffic or by having the devices actively generate and exchange simulated traffic streams.

Agilent's and Brix's products can work both ways. NetIQ's applications only can actively send streams to each other. Acterna, Finisar and Sniffer only work passively.

Brix and NetIQ did the best job of reporting QoS parameters and VoIP voice quality. However, before you start generating traffic between monitor devices over a production network, ensure that network-control safeguards and policies are in place to avoid affecting production traffic.

Acterna, Finisar and Sniffer products were tested with one monitoring device plugged into the network, observing all network activity from one point of view. Brix, NetIQ and Agilent products — except for Agilent's Network Analyzer — had devices, including PCs with special client software, plugged in at two or more points on the network.

The products that had multiple points of network presence tended to report results more thoroughly. Take QoS measurements, for example: Neither Acterna, Finisar nor Sniffer report one-way latency, which is a key metric in assessing VoIP connection quality.

Now onto the detailed test results for each product package.



Acterna DA-3400 and VoIP Analysis Software

Acterna addresses all the functional areas pretty well, doing the best job in real-time VoIP monitoring and QoS measurement of VoIP parameters. And a few areas, such as VoIP traffic generation, are minimally addressed.

The main VoIP-monitoring application yields VoIP screens that offer increasingly detailed views of VoIP activity, focusing on VoIP call quality, throughput, jitter and packet loss. The panes are easy to view and understand, with one exception: Jitter and packet loss are reported differently on two different screens, which was confusing — and misleading.

Strong features of the Acterna package are accurate accounting of VoIP calls — both completed and active calls — and the ability to drill down for details on particular calls. The package accurately reported whether silence suppression is employed per VoIP stream. The graphical jitter and VoIP-bandwidth displays are excellent. Voice quality is shown on a simple, but effective "good," "fair" and "poor" scale, as set by the user. There are some nice, built-in report-generation capabilities, too.

On the downside, there is no association retained between end user-to-end user VoIP streams, and their call-control protocol. We ran Session Initiation Protocol (SIP) and H.323 calls in the test bed, but Acterna could not report the activity of either call-control environment. Without this perspective, monitoring call performance or effectively diagnosing, for example, call-setup problems is more difficult.

Bottom line: Acterna

DA-3400 Data Network Analyzer loaded with Ethernet Analysis Software 1.1, and optional VoIP Analysis software

www.acterna.com

Price as tested: Hardware/software platform(s), \$22,575; additional software, \$3,115.

Best suited for: Real-time VoIP monitoring and QoS measurement.

Agilent Telephony Network Analyzer, Voice Quality Tester

The half-dozen products Agilent offers in its VoIP-analysis repertoire collectively cost more than \$100,000, ranking it the second most-expensive product package we tested, behind Brix.

The centerpiece is the Telephony Network Analyzer (TNA) software, which runs on the Network Analyzer platform. The latter is a portable, Windows-based chassis that accommodates up to two

VxWorks-based LAN or WAN-specific interface modules. We tested a late-beta release of Version 1.2 of TNA and found it unstable. When the software crashed frequently, it took the Network Analyzer platform and all active VoIP calls with it. Agilent says it has fixed the problem.

TNA provides a real-time view of all VoIP sessions in a large table, which is difficult to maneuver. This single table, with about 20 columns that you have to scroll through to view, provides copious details of each call from the Real-Time Protocol (RTP) layer down. However, you can't associate an RTP stream with its higher-layer, call-control protocol within this application.

The table contains a lot of raw data, but you must know what the displayed values mean to make sense of the data. That applies to the Mean Opinion Score (MOS), a 1 to 5 scale widely applied to represent relative VoIP voice and connection quality — and QoS ratings that the software assigns to each call. These automated ratings weren't working correctly, either, during the evaluation.

The Expert Commentator, a component of the base Network Analyzer software system, is a real-time monitoring tool that shows SIP and H.323 call-control details. Again, there's a lot of raw data but not much readily useable information. SIP and H.323 call data is represented inconsistently, making it difficult to compare details between the two protocol environments. Also, you can only observe full details on a call after the call has been completed and a Call Detail Record written.

Agilent's Voice Quality Tester (VQT) runs on a different hardware platform. It is designed to generate VoIP call set-up sequences to exercise the user's H.323 or SIP equipment, and generate VoIP streams to a remote VQT station to assess call quality. Setting up this package is difficult and tedious. Also, some help in interpreting the call-quality assessment ratings is sorely needed.

We had better luck with Agilent's Application Analyzer, which is PC software that runs on two or more nodes on the network. This product measures QoS parameters and worked very well in our testing.



Bottom line: Agilent

Network Analyzer 1.2; Telephony

Network Analyzer with Predictive MOS 1.2 (late beta); Application Analyzer 2.0, with Voice Analysis option; Report Center 1.0; IP Telephony Report Software 12.01; and Voice Quality Tester <http://xpl.comms.agilent.com>

Price as tested: Hardware/software platforms, \$58,000; additional software, \$45,500

Best suited for: VoIP traffic generation, QoS measurement, and VoIP traffic and protocol decode.



Brix Networks Verifiers and BrixWorx, with Advanced VoIP Test Suite

The Brix system is oriented toward carriers because of its template-based software design, which is structured toward monitoring service-level agreements (SLA). The Sun server-based system has the capability to monitor the ongoing availability of key VoIP nodes, a feature that's needed more in a carrier setting, and can generate VoIP traffic. Pricing is at the high end, too, with a typical Brix monitoring setup exceeding \$100,000.

Brix has no traditional network sniffing and decoding application, so real-time VoIP monitoring and VoIP traffic and protocol decode are not strengths. The user has to specify — by IP address, User Datagram Protocol (UDP) port number range — the particular VoIP node or traffic that is measured according to specified SLA parameters. Brix relies on its own generated traffic streams to perform VoIP voice-quality assessment and QoS measurement, both of which the Brix system does superbly.

The Brix system's real strengths lie in QoS measurement and voice-quality assessment. Brix is one of a growing number of network test equipment vendors that license and integrate VoIP voice quality assessment software from Telchemy. Finisar is another Telchemy partner.

The result is very useful and accurate VoIP voice-quality assessments, which incorporate all relevant factors of jitter, delay and packet loss. We have compared and found the Brix voice-quality assessments consistent with subjective human MOS ratings for the same VoIP connections and impairment conditions.

Bottom line: Brix Networks

100, 1000 and 2500 Verifiers, with BrixWorx Engine 2.2 and BrixWorx Advanced VoIP Test Suite 2.2 software

www.brixnet.com

See VoIP, page 50

VoIP, Continued from page 48

Price as tested: Hardware/software platforms, \$30,000; additional software, \$100,000

Best suited for: VoIP traffic generation, VoIP voice-quality assessment and QoS measurement.



Finisar Surveyor v5.0, THGs monitor/analyzer, and Multi-QoS software

One of the best real-time VoIP-monitoring applications we evaluated was Finisar's Multi-QoS, an optional software package that integrates with the vendor's Surveyor 5.0 software, Finisar's anchor product. We tested the Surveyor software with Finisar monitor/analyzer units, called a THGs (short for 10, 100 and gigabit system).

The Surveyor software, with the optional Multi-QoS application, can run stand-alone on a Win 2000 PC, without the THGs device, and provide the same functionality. Theoretically, performance could be reduced because all traffic has to be viewed through the PC's LAN interface, typically via a mirrored switch port. But this yields a very effective VoIP real-time monitoring system for a fraction of the cost of a system equipped with THGs.

The focus of the Multi-QoS application is a flexible table of all VoIP calls, with a clean, single-line entry for each call. Clicking on a button shows active calls, and clicking on another shows completed calls. Selecting both gives you a view of all current and recently completed VoIP calls.

The table is built from the call set-up protocol. The table captured and represented all our SIP calls and traffic correctly, and did fairly well with our H.323 traffic and calls. Surveyor fully identified and represented one-third of our H.323 calls in the table. However, while the product observed and accounted for the remaining two-thirds of the calls, not all the minor details related to the traffic were reported in full.

Bottom line: Finisar

Surveyor 5.0 with Multi-QoS and Packet Blaster software, and the THG full duplex monitor/analyzer

www.finisar.com

Price as tested: Hardware/software platforms, \$19,390; additional software, \$7,000

Best suited for: Real-time monitoring and alarm generation, and VoIP traffic and protocol decode.

NetIQ Vivinet Software Series and Chariot Advanced v4.3

NetIQ has packaged three of its software tools under the Vivinet brand. They are: Vivinet Manager, which performs call

VoIP traffic analysis packages: Hits and misses

In our testing of VoIP traffic analysis tools, we defined seven functions we felt were necessary in order for these suites to be useful in an enterprise VoIP deployment. Here we outline which product suites hit or missed those functional areas.

Vendor/VoIP analysis function	Acterna	Agilent	Brix	Finisar	NetIQ	Sniffer Technologies
Real-time monitoring, alarm generation	++	Limited *	+	++	None	+
Long-term monitoring, trend reporting	+	+	+	+	+	Limited
VoIP traffic generation	Limited	++	++	+	++	Limited
VoIP voice-quality assessment	+	+	+++	+	+++	None
QoS measurement	++	++	+++	Limited	+++	++
VoIP traffic and protocol decode	+	++	Limited	++	None	+++
Intelligent diagnosis of VoIP problems	+	Limited *	+	Limited	++	+

Key:

+++ = Excellent; fully satisfied or addressed most or all tasks.

++ = Very good showing; some tasks or criteria not addressed.

+ = Effectively addressed to some degree.

Limited = Addresses minimally; inaccurate results or other problems.

None = No capabilities in this category.

* Based on Agilent's Telephony Network Analyzer 1.2 application, which crashed repeatedly during our testing.
** NetIQ's products address this functional area in ways we could not test, specifically relating to how it applies to Cisco-only VoIP environments.



setup sequences with the user's VoIP nodes, such as gateways and call controllers; Vivinet Assessor, which performs automated call-quality assessment and issues detailed reports; and Vivinet Diagnostics, which performs analyses of VoIP networks through Simple Network Management Protocol polling and various other background processes.

NetIQ also still features its real-time analysis and simulation tool, Chariot Advanced, which the company acquired when it bought Ganymede two years ago.

NetIQ's products are software-only and work pretty much the same way, by generating simulated VoIP traffic. Two or more nodes running end-point versions of the software sync up and send user-specified data streams to each other.

We were impressed with Chariot Advanced 4.3, which sends bidirectional VoIP streams and then reports voice-quality assessment based on VoIP parameters and measured impairments. These same features are integral to the Vivinet Assessor package. The Assessor package adds extensive report-generation capabilities, which the Chariot software lacks.

The automated VoIP voice-quality assessment capability of NetIQ's Chariot Advanced provides the most accurate MOS-estimate assessment of the products

evaluated, based on comparing the NetIQ assessments with human interactive ratings under the same set of VoIP and network conditions.

Bottom line: NetIQ Vivinet Manager 2.1, Vivinet Assessor 2.0, Vivinet Diagnostics 1.0 and Chariot Advanced 4.3

www.netiq.com

Price as tested: Vivinet suite, \$22,500; Chariot Advanced, \$8,000.

Best suited for: VoIP traffic generation, VoIP voice-quality assessment, and QoS measurement.



Sniffer Technologies Sniffer Distributed v4.2, with Sniffer Voice 2.1

The industry-standard Sniffer keeps getting better as Network Associates adds more modules to the product. The latest addition, Sniffer Voice, adds a VoIP-oriented, application-layer perspective to network traffic that's observed or captured by the base Sniffer software. Most of the VoIP-call details are revealed from inside the Expert window of the Sniffer interface.

For reasons that neither our testers nor the onsite Network Associates engineer could determine, the Sniffer package exhibited one problem: The association between monitored VoIP sessions and their higher-level call set-up protocols was lost. As a result, the orientation of the VoIP screens, designed to let users drill down or climb up into any level of VoIP calls, wouldn't work correctly for H.323

calls or SIP calls. The Sniffer team proved this was a bug by plugging in a Sniffer Portable implementation, which showed that VoIP call streams could remain associated with their call setup protocols. Sniffer Technologies promised this would be fixed by publication.

A nice set of VoIP-oriented alarms can be enabled already, and Sniffer hints that a much-expanded set is coming in a release early next year.

Sniffer's forte always has been efficient packet capture, filtering and decode. And it retains this reputation for VoIP, providing extremely detailed yet understandable decodes of VoIP packets, filtering on specific VoIP conversations and exchanges, and even decodes of VoIP call set-up.

But there is room for improvement. For example, the Sniffer could not show the cumulative real-time amount of bandwidth consumed by VoIP traffic. And the system could benefit from a clean, accurate way to show how many VoIP calls are active.

Bottom line: Sniffer Technologies

Sniffer Distributed 4.2 with Sniffer Voice - Release 2.1, Optional Sniffer Voice software v 2.1 and Sniffer Portable (software only)

www.networkassociates.com

Price as tested: Hardware/software platforms, \$11,900; Sniffer Voice software, \$5,250; Portable Sniffer software, \$18,000.

Best suited for: QoS measurement, and VoIP traffic and protocol decode

Mier is president, Battistelli is director of operations, and Miner is director of technology analysis at MierLabs Inc., a network research and product test center based in Hightstown, N.J. They can be reached at emier@mierlabs.com, vbatt@mierlabs.com and aminer@mierlabs.com, respectively.