

internet evolution

Most Recent Comment

"I don't think the ISPs are concerned about illegal content going through their networks (as in moral issues). They are concerned that traffic is going through their networks and they are not earning anything on it. P2P is today's headache because there's soooo much traffic but ISPs just ..."

Mr. Roques on P2P Taste Test

DISCUSS PRINT



DIGG



DEL.ICIO.US



REDDIT



EMAIL THIS



TWEET THIS

P2P Taste Test

Written by [Carsten Rossenhövel](#)

7/9/2009 11 comments

Introduction

Any mention of Deep Packet Inspection (DPI) these days is bound to generate heated discussion. Everybody in our industry has a strong position, either viewing peer-to-peer (P2P) filtering as mandatory and the only way to manage consumer traffic in Internet service provider (ISP) networks and/or to enforce copyright law, or as evil because a service provider shall treat all customers equally and the media industry should accept the end of its business model.

In the end, we found that mentioning money quieted opponents and made the business case for ISPs obvious. None of the P2P users we met (in a non-representative internal survey) was ready to pay for premium, 24x7, wire-speed, broadband Internet access. And few service providers seem to be ready to spend money on a huge, unbeatable consumer network. But protocol filtering to manage the vast amounts of data from P2P power users (more than 50 percent of all Internet traffic from less than 10 percent of the users, according to one [study](#)) is the only way to achieve both goals today. *Quod erat demonstrandum*. Three protocol filtering vendors participated in the test and showed off great results.

We were not so sure about the content filtering part of the story. The media industry keeps fighting, specifically in European countries. The recent Pirate Bay trial and the discussion going in France about the HADOPI law ("HADOPI" being the acronym for the government agency of the *Haute Autorité pour la Diffusion des Œuvres et la Protection des droits sur Internet* created by this bill), which regulates and controls Internet usage in order to enforce compliance with copyright laws, are only the most prominent examples. It is also interesting to monitor what has happened in countries where courts already ruled in favor of the media industry a while ago, like [Ireland](#) and [Belgium](#).

In this followup to our [original test of P2P solutions](#) published in March 2008, we tested three vendors' products between November 2008 and February 2009: [Cisco Systems Inc.](#) (Nasdaq: CSCO)'s SCE-8000, [ipoque's](#) PRX-10G Traffic Manager, and [Procera Networks Inc.](#) (Amex: PKT)'s PacketLogic 10014.

Unlike the March 2008 test, which was paid for by SNEP, the French recording industry association, each vendor in this latest test paid to participate. Vendors had only a generic veto right to opt out, not specific editing rights related to individual test cases. Vendors got the right to review their results and our interpretations related to their results; we did not share the full report with them, nor the results or names of the other participating vendors.

And in contrast to our first test, we also increased the testbed scale to 25 Gbit/s at the P2P application layer and added the latest P2P protocols.

Unfortunately, not a single content filter vendor was ready to publish its results in our test. Why? According to our sources, some pretty amazing functionality has been implemented, but it [just doesn't scale yet](#). [Audible Magic](#) rejected participation in our test, unfortunately. Other vendors, like [SafeMedia Corp.](#), had [publicly expressed interest](#) to join our previous test, but when we invited them to the current campaign, nobody responded. Yet another content filter vendor expressed interest, but then claimed that all its resources were completely consumed by a trial in Australia (although we patiently waited for more than three months).

Therefore, we cannot report that media content filtering solutions (as opposed to protocol-based filtering) would be ready for today's scale of broadband Internet networks based on our testing.

Contents of the Report:

- **Page 2: [How & What We Tested](#)**
- **Page 3: [Testbed Details & Methodology](#)**
- **Page 4: [P2P Clients Tested](#)**
- **Page 5: [Test Results](#)**
- **Page 6: [Throughput Claims Tested](#)**
- **Page 7: [P2P Detection & Regulation](#)**
- **Page 8: [Filtering & ISP Constraints](#)**
- **Page 9: [Conclusions](#)**
- **Page 10: [Testing Notes & Specifications](#)**

— Carsten Rossenhövel is *Managing Director of the [European Advanced Networking Test Center AG \(EANTC\)](#), an independent test lab in Berlin. EANTC offers vendor-neutral network test facilities for manufacturers, service providers, and enterprises. Carsten heads EANTC's manufacturer testing and certification group and is responsible for the design of test methods and applications.*

Next Page: [How & What We Tested](#)

How & What We Tested

The goal of our new series of tests was to verify the functionality and performance of P2P filter devices -- also known as Deep Packet Inspection (DPI) devices -- with regard to P2P traffic detection/policing functionality and the suitability of deployment in ISP environments. DPI devices inspect and take action based on the contents of the packet (commonly called the "payload") rather than just the packet header.

In a first step we verified the forwarding performance of the devices under test for a mix of stateful application-layer traffic (HTTP and P2P) to analyze the detection accuracy under load. Next, we evaluated the P2P regulation and filtering performance and accuracy. Finally we emulated specific ISP traffic conditions, asymmetric routing, and encapsulated traffic, and tested the devices' performance under these special conditions.

The DPI devices acting in the P2P world can be divided into two categories: one group doing traffic (pattern) classification, usually protocol/signature based; and the other classifying the

exchanged content of P2P sessions, usually based on hash values bound to specific content, like an audio MP3 file.

With our new campaign we tested protocol-based solutions. The signature library of DPI devices doing protocol-based classifications incorporates protocol transport layer (TCP/UDP) ports, strings of protocol messages, numeric properties of protocol sessions, and behavioral and heuristic parameters of protocols into the signature calculation.

ISPs and carriers

DPI is currently one of the hottest issues for ISPs and carriers, not only for P2P traffic regulation. For ISPs and carriers seeking new sources of revenue by deploying services like triple play, DPI devices are essential for:

- Analyzing subscriber statistics on the application layer for network design planning
- Setting up global/regional application control at peering points where service providers purchase/sell bandwidth from/to other service providers
- Setting up per-subscriber service-level agreements (SLAs) or policies, in order to enforce smarter services, institute volume/duration-based billing, be more competitive, provide better QoE (quality of experience), and increase average revenue per user (ARPU)

Our tests positively verified participating vendors' performance claims. Or, in flashier marketing terms, "Hey service providers, we have proven multi-gigabit-per-second application filtering rates! These boxes are ready for the backbone network or peering points now, and will in all likelihood decrease network production costs!"

The basic P2P detection functionality and accuracy test showed very good results for all three devices, all being accurate within 2 percent of the possible maximum rate. For this test all products were exposed to around 6.25 Gbit/s of application traffic per interface pair, at maximum, and where the device's interface number allowed, 25.0 Gbit/s. The application traffic consisted of a mix of HTTP and the three most popular P2P applications: [BitTorrent Inc.](#), eDonkey, and Gnutella.

We exposed the participants to a more complex protocol mix of HTTP and P2P traffic for the advanced tests of detection, regulation, and filtering/blocking. In the detection section, all three devices had little problem detecting the popular protocols, however, with the more obscure protocols some updates and adjustments for Cisco and Procera were necessary to enable them to perform as expected.

For regulation, all three devices showed good accuracy in throttling BitTorrent, eDonkey, and Gnutella; the worst-case deviation we recorded was around 8 percent. Some protocols were more difficult to throttle than others; Direct Connect was the most challenging.

Last but not least, for filtering/blocking, all three devices were able to perform with high accuracy -- only 1.8 percent of the offered P2P protocol traffic was able to slip through. Again, the update of the signature definitions in Procera's case led to very good results reported here. Ipoque's blocking accuracy was unmatched, with barely any detectable P2P traffic transmission.

Our test plan of special ISP conditions initially included an asymmetrical routing scenario and traffic encapsulation. We removed this from the test later on, as vendors had limited support and, more importantly, raised the issue that service providers do not intentionally plan for asymmetric routing scenarios so they would only appear transiently.

Next Page: [Testbed Details & Methodology](#)

Testbed Details & Methodology

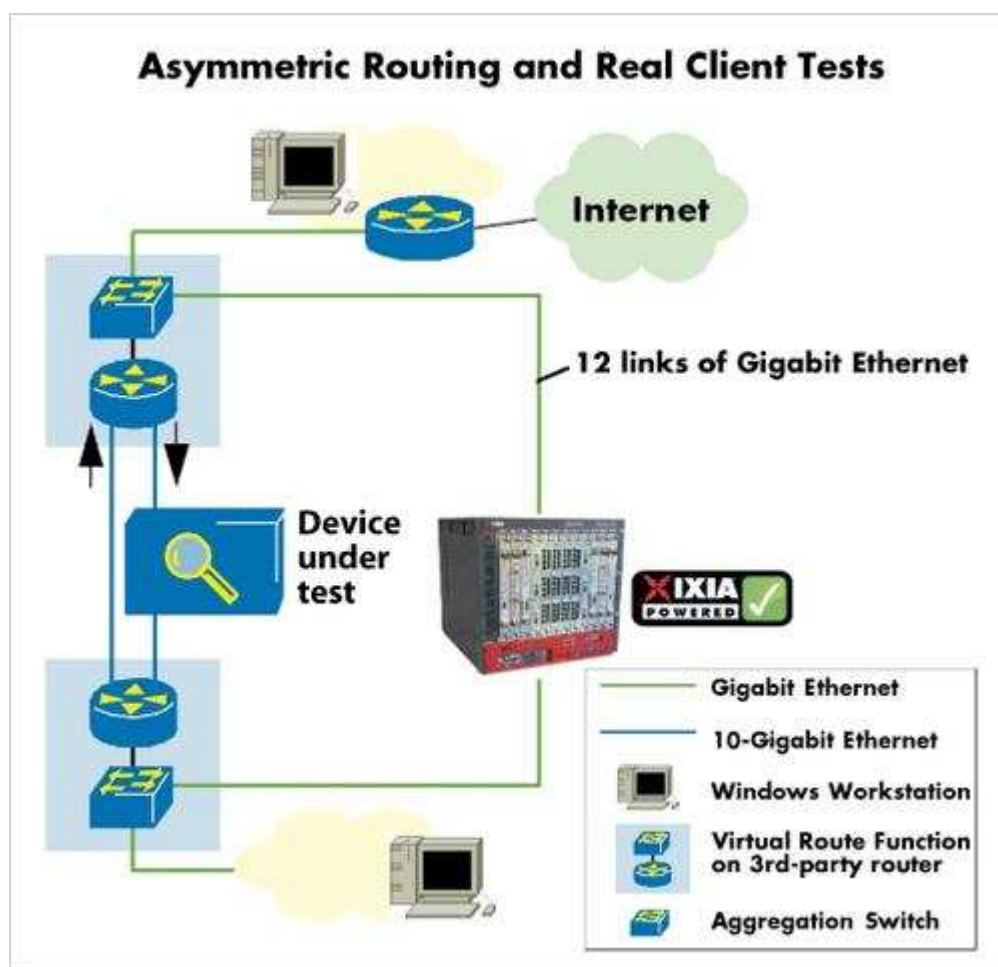
As the following two figures show, we used [Ixia](#) (Nasdaq: XXIA) test equipment to set up our

peer-to-peer test bed and to emulate ISP network conditions. We used an Ixia XM12 chassis and a total of eight 10-Gigabit Ethernet load modules (ASM1000XMV12X-01 Application and Stream Load Module).

For the application layer performance and the P2P detection/filtering tests, Cisco's SCE-8000 was connected to a total of four 10-Gigabit Ethernet load modules (using the maximum number of ports Cisco offered on this device), whereas ipoque's PRX-10G Traffic Manager and Procera's PacketLogic 10014 were connected to a total of eight 10-GigE modules. The devices under test were operated in transparent bridging mode, so transparently configured into the Ethernet link acting as a simple wire, forwarding, limiting, or blocking the traffic without intervening with the higher protocol stack layers.

Out of the theoretical maximum of 20 Gbit/s we were able to use up to 6.25 Gbit/s of application traffic per interface pair. We were limited by the application layer performance boundaries on the load module and the very unsymmetrical traffic pattern (the download direction always hitting performance limits first). This way we were able to generate a total of 25 Gbit/s application layer traffic, a massive performance boost, compared with our 2008 test campaign where we were able to generate only 1 Gbit/s application layer traffic.

For our tests emulating specific ISP traffic conditions, we configured a third-party router to set up the asymmetrical routing scenario and to attach workstation PCs and a DSL router to the test bed. A third-party switch aggregated traffic from the load generator and from the real P2P clients. With the latter, we were able to perform tests with the real clients connected to the Internet. The Ixia IxLoad application replay worked for emulated instances between tester ports only.



For the asymmetrical routing scenario an additional 10-GigE link was set up. Both routers were configured to forward traffic to each other in a unidirectional fashion. Depending on the type of the test (upstream or downstream), the device under test was inserted into one of the two links.

The test was designed to precisely emulate current conditions in service provider networks, as well as those that ISPs will face in the near future. We emulated a realistic mix of Web applications and a diverse set of peer-to-peer applications. The P2P applications were selected according to their popularity, difficulty in recognition, and foreseeable future trends. The aggregated traffic model enabled us to assess the filtering performance typically required at large Internet service provider networks.

Even if P2P traffic is one of the most noticeable applications filling up Internet trunks -- typically taking around half of the Internet traffic today -- it is obviously not the only protocol. To make our test more realistic, we added generic HTTP traffic to the mix transmitted in our tests.

Our tests are separated into three phases. First we verified the baseline performance for IP and application layer traffic. In the second phase we focused on detection, regulation, and filtering accuracy under load, using a complex mix of numerous P2P protocols. In the third test phase we emulated special ISP traffic conditions such as asymmetrical routing and encapsulated traffic.

While many application load generators, including Ixia's IxLoad software used in the test, are able to emulate established application protocols such as HTTP, FTP, etc., emulation of P2P traffic is more challenging due to constant improvements in the existing P2P protocols and differences in various clients' implementations, as well as many new emerging ones.

For the emulation of the P2P traffic on a massive scale, we used the "Application Replay" feature of the IxLoad software, which allowed us to emulate thousands of separate P2P clients using previously captured network traffic (PCAP) of a real P2P client. With Application Replay, the load generator emulates a TCP connection by replaying data from the PCAP file. The IP addresses and the port numbers of the communicating parties are dynamically replaced, allowing simulation of multiple hosts. The traffic is generated using a real TCP stack, which allows for realistic traffic behavior and simulation of the typical effects of packet loss and delay.

In addition to the traffic generators, we also used real client PCs to generate peer-to-peer traffic using actual P2P applications. The clients could exchange P2P content between each other as well as with the Internet, in both cases allowing the traffic to pass through the device under test.

Next Page: [P2P Clients Tested](#)

P2P Clients Tested

In our tests we used multiple popular P2P protocols and clients listed in the table below. For BitTorrent and eDonkey, we used two different clients in order to test each device against two implementations of the same protocol with potentially different behavior. For the Gnutella protocol, we used [Lime Wire LLC](#) and Shareaza clients, which implement similar, but slightly differing protocols (Gnutella1 and Gnutella2, respectively).

P2P Protocol	P2P Client	Version
BitTorrent	Azureus/Vuze	3.1.1.0
	µTorrent	1.8.1
EDonkey	eMule	0.49b
	aMule	2.2.2
Gnutella	Limewire	4.18.3
	Shareaza	2.4.0.0
Direct Connect	DC++	0.707
Soulseek	Soulseek	157 NS 13c
BitTorrent (encrypted)	Azureus/Vuze	3.1.1.0
eDonkey (obfuscated)	eMule	0.49b
Ares (encrypted)	Ares	2.0.9.3030
Filetopia (encrypted)	Filetopia	3.04d
MP2P	Manolito	3.0.6

We defined the following composition of Layer 7 traffic to be used for all P2P detection and regulation tests. The total bandwidth of Layer 7 traffic sent through the device under test was approximately the maximum throughput value determined in the test case 4.2. In general, our Layer 7 traffic comprised equal parts of non-P2P (HTTP) and P2P traffic. Only stateful TCP traffic was used in the tests.

The P2P traffic bandwidth was distributed into two categories of P2P protocols depending on their popularity and prevalence in the Internet traffic. The protocols in the "informational" category served the demonstration of general capabilities of the device under test to detect and/or block different P2P protocols and were assigned a small fixed bandwidth. This category included relevant cases of encrypted protocols and different clients of some of the same protocols.

Protocols in the "normal" category included the most popular P2P protocols and were assigned a larger fixed bandwidth. This category served the demonstration of the detection capabilities of the tested device under stress.

The following table details the distribution among the traffic categories and applications. Note that the actual bandwidth achieved in the tests may differ from the desired value. The given load values are per-interface pair of the device under test. Accordingly, the values used for Cisco SCE-8000 (two interface pairs) were twice this value, and those used for Procera PacketLogic and ipoque's PRX-10G (each equipped with four interface pairs) were four times this value.

Category	Protocol	Application	Encr.	Objective Bandwidth, Mbit/s
Non-P2P	HTTP	IxLoad		3000
P2P Informational	BitTorrent	µTorrent		62.5
	eDonkey	aMule		62.5
	MP2P	Manolito		62.5
	Soulseek	Soulseek		62.5
	BitTorrent	Azureus	✓	62.5
	eDonkey	eMule	✓	62.5
	Ares	Ares	✓	62.5
	Filetopia	Filetopia	✓	62.5
P2P Normal	BitTorrent	Azureus		500
	eDonkey	eMule		500
	Gnutella1	Limewire		500
	Gnutella2	Shareaza		500
	DirectConnect	DC++		500
Total				6000

In our evaluations, we took into account that the device under test combines some of the related protocol statistics into a single statistic. We had to combine the following statistics:

Device under test's statistic	Analyzer's statistics
BitTorrent	Plain BitTorrent (Azureus) Plain BitTorrent (uTorrent)
eDonkey	Plain eDonkey (eMule) Plain eDonkey (aMule)
Gnutella	Gnutella 1 (Limewire) Gnutella 2 (Shareaza)

Next Page: [Test Results](#)

Test Results

The evaluation of the test results, specifically comparison among the three vendors, was challenging, as the tested devices belong to different performance classes and so are designed to handle different loads. On the other hand, all tested devices provide common detection, throttling, and blocking functionality; all three operate as transparent bridges; and all three use 10-GigE interfaces on both customer and Internet sides.

These similarities make it possible to establish a common level for the test result comparison. We normalized the test results presented by absolute values, namely the offered load, achieved throughput, and the observed TCP errors, according to the number of each device's network interfaces. On the other hand, detection accuracy as well as throttling and blocking efficiency are presented as relative values to the expected result.

For example, a 100 percent accuracy would correspond to ideally precise measurement of the protocol traffic. Values below 100 percent indicate that the DPI device has missed some of the protocol's traffic, while higher values are a sign that some other traffic was falsely classified, causing the device to report more traffic than was actually transmitted.

In the throttling and blocking tests, we demonstrate the effect of blocking on the traffic volume transferred during the test by comparing it with the reference test. Here, a value of 100 percent indicates that exactly the same volume was transferred, and values lower than 100 percent indicate that the traffic decreased.

In some cases, we observed the opposite effect; the traffic of some protocols actually increased with filtering enabled, which results in values over 100 percent. This behavior can be usually explained by failure to detect this particular protocol, with the added effect of available bandwidth freed by filtering of other protocols.

In all tests involving application traffic, the device was exposed to a specific offered load and the concurrent connections number. Both of these values were defined per interface pair of the device under test. As result, Procera's PacketLogic and ipoque's PRX-10G devices, both equipped with 4 interface pairs, were exposed to twice the load and concurrent connections number as the Cisco SCE-8000 with its two interface pairs.

The specified load and the concurrent connection number were maintained for the sustained duration of the test: 10 minutes. The ramp-up and ramp-down phases of the test took 90 and 70 seconds, respectively.

In the first test, the devices under test were exposed to a mix of Web and peer-to-peer application traffic. The goal was to verify the throughput of the device when performing the baseline P2P detection.

The traffic consisted of a mix of HTTP and the three most popular P2P protocols -- BitTorrent, eDonkey, and Gnutella -- in a realistic proportion, based on recent surveys of P2P traffic composition in the Internet. The tested devices were exposed to approximately 6.25 Gbit/s of application traffic at 10,000 concurrent TCP connections per device-under-test's interface pair.

In order to pass the test, the device had to be able to detect all protocols used in the mix with accuracy close to 100 percent. After the transmission was finished, we collected the traffic detection statistics from the devices and compared them with the traffic volume actually sent by the load generator. If the device failed to detect part of the traffic because of the overload, the load was lowered in 10 percent increments until detection was accurate.

Next Page: [Throughput Claims Tested](#)

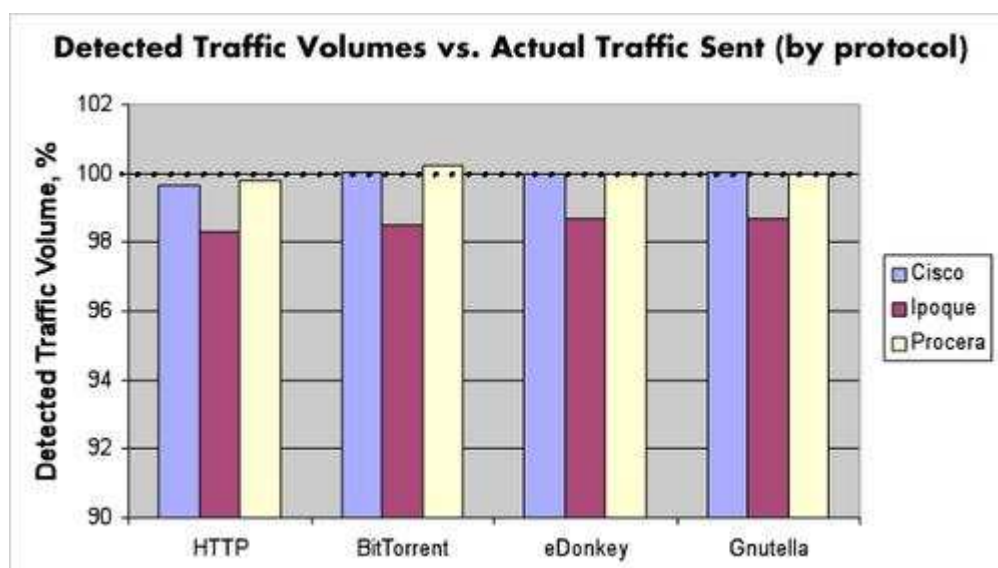
Throughput Claims Tested

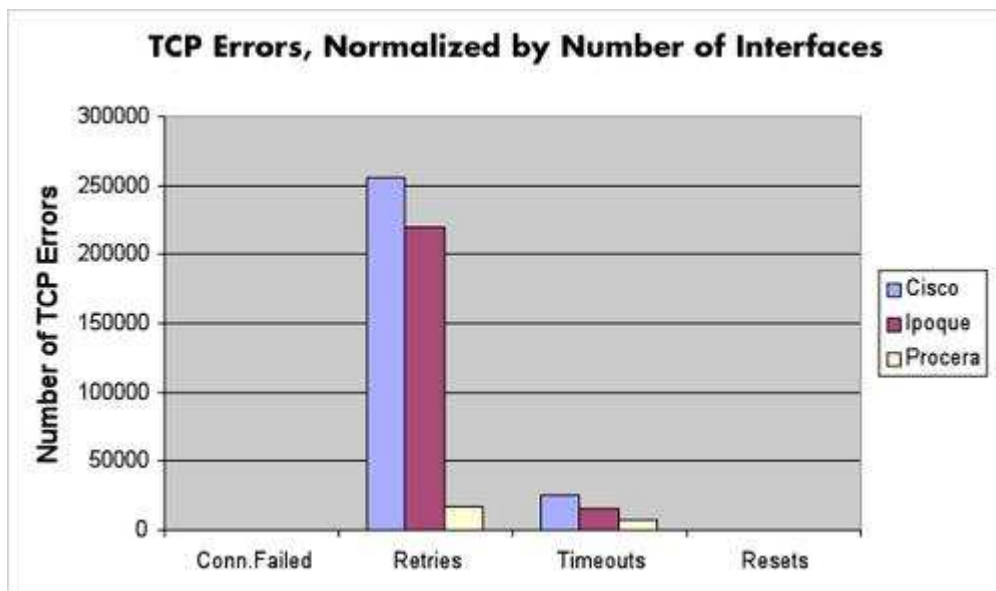
The following graph shows the maximum application layer traffic rate achieved with each device, based on the test bed configuration (two or four channels). The Cisco SCE-8000 was attached with two channels (four ports), the Procera PacketLogic and ipoque's PRX-10G were connected with four channels (eight ports). In a perfect world, we would have had a sufficient number of Ixia emulator ports plus a large 10-GigE switch to aggregate the load for each filtering system port up to wire-speed; however these resources are expensive and could not be supplied. We still had one of the largest-scale P2P test beds worldwide; it is not a trivial task to implement a test solution scaling to tens of gigabits of real application layer traffic.

All three vendors claimed much higher P2P throughput; we were not able to confirm their figures on the P2P layer due to the limited scale of the test bed.

The successful completion of this test provided an individual performance baseline to be used in all following tests. The chosen P2P protocols should be easily detectable by the device, and would only result in lower accuracy in case of an overload.

Based on the statistics produced by the tested devices, and the statistics from the analyzer, we calculated the accuracy of the detection, presented in the chart below. All three devices were able to detect all protocols. The reported traffic volumes for each of the protocols did not deviate by more than 2 percent from the actual volume transmitted by the load generator. The somewhat higher deviation of ipoque's PRX-10G happens because the first few packets of newly opened TCP streams go unclassified until the device tracks P2P protocol data. The SCE-8000 and PRX-10G caused a relatively high number of TCP retransmissions during the test, which however did not cause a major impact on the throughput.





Next Page: [P2P Detection & Regulation](#)

P2P Detection & Regulation

In the next series of tests, we exposed the participating devices to a more complex protocol mix of traffic to challenge their detection capabilities for a wide variety of peer-to-peer protocols.

The device under test was exposed to approximately 6 Gbit/s of application traffic at 9,000 concurrent TCP connections per interface pair. The offered traffic was chosen to be slightly lower than in the previous test, in order to ensure that the devices under test would operate well within their performance limits.

The goal of this test was to determine detection accuracy for many P2P protocols sent simultaneously in a complex mix. The presence of multiple protocols with closely related behavior patterns might lead to misdetection or misinterpretation by the device under test and result in statistical deviation.

This time, we included several popular protocols: BitTorrent, eDonkey, Gnutella 1 and 2, and Direct Connect, each with approximately 500 Mbit/s of bandwidth per interface pair. In addition, many other P2P protocols, including several encrypted, were added to the mix at a lower bandwidth. HTTP was added as a non-P2P protocol in order to verify that the P2P filtering does not affect common protocols running in parallel.

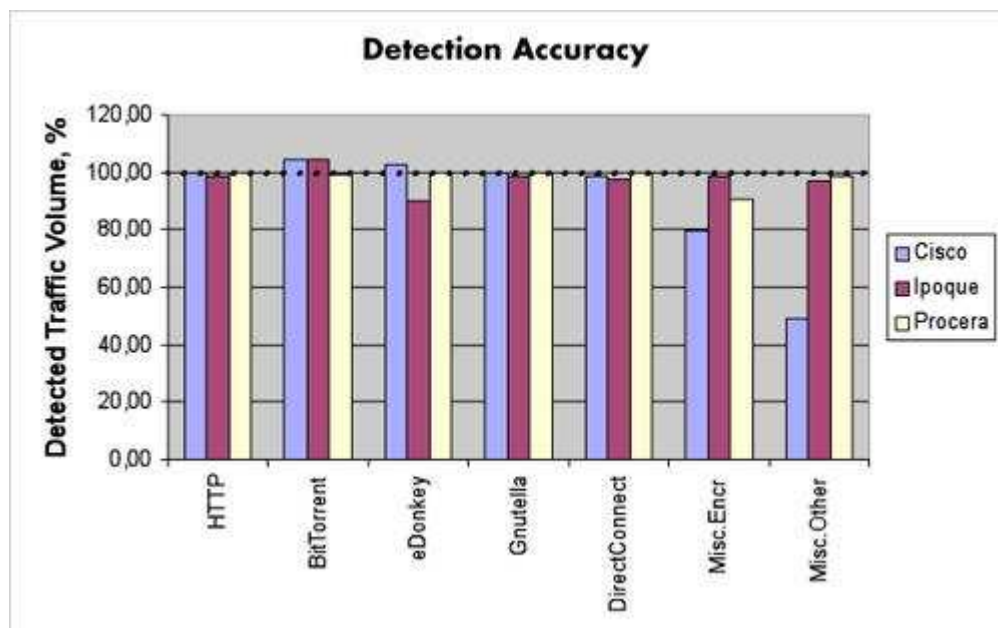
All three devices detected the popular P2P protocols BitTorrent, eDonkey, Gnutella, and Direct Connect in the high-bandwidth group with high precision.

Procera's PacketLogic initially failed to correctly detect Gnutella2 from the group of high-bandwidth protocols, as well as Manolito, Soulseek, and Filetopia from the low-bandwidth group. Procera was able to quickly analyze the detection problems and provide an update to its signature definitions within one day. In the repeated tests, PacketLogic was able to detect all these protocols with high precision. Both Cisco SCE-8000 and Procera PacketLogic showed a similar deviation in detection of encrypted BitTorrent traffic, failing to detect a part of the traffic. We determined the highly synthetic nature of the emulated traffic as the cause and could show in a test with the real client that the detection of the encrypted BitTorrent worked perfectly with both devices.

Both these devices also showed a misinterpretation of Soulseek protocol as a different P2P-based video streaming protocol, PPlive. Initially, SCE-8000 was also able to detect the Ares protocol only partially, but the detection could be easily improved to 100 percent in a

test with the real client by adjusting the algorithm parameters.

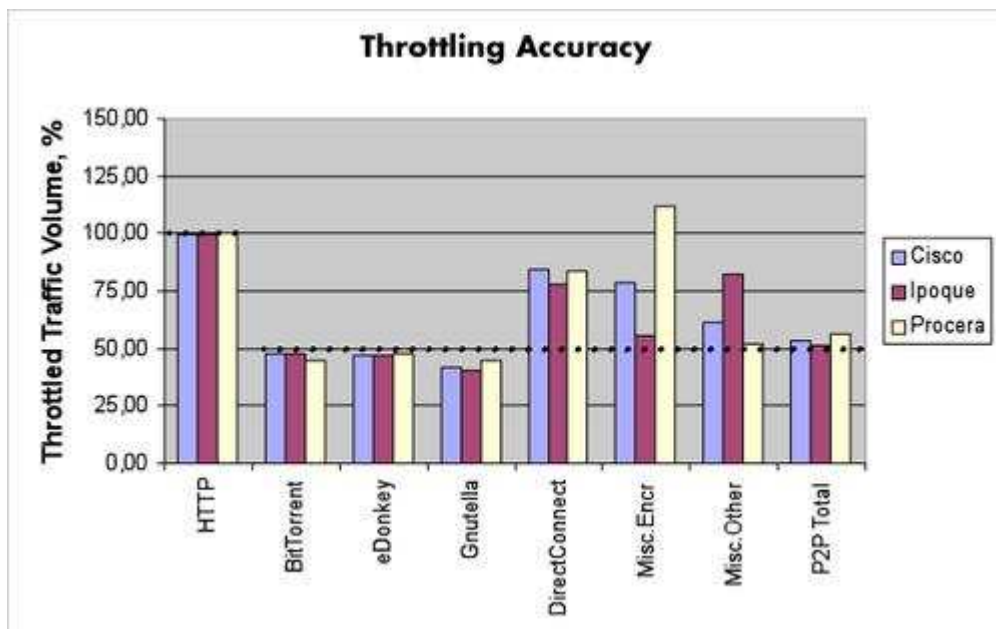
All three devices had little problem detecting the popular protocols. Also the more obscure protocols were detected properly, albeit with some updates and adjustments in the case of SCE-8000 and PacketLogic. In comparison, ipoque's algorithms performed consistently from the beginning.



Further, we tested the selective traffic throttling feature available in all three devices. In this test, the DPI devices had to limit the P2P traffic to 50 percent of the reference value, while HTTP traffic must be left unaffected.

All three devices showed similar behavior with the popular protocols and were able to throttle BitTorrent, eDonkey, and Gnutella with good accuracy. The initial deviation in Procera's case was caused by failure to detect some of the protocols and was fixed with the signature definitions update. Although all three participants correctly detected Direct Connect in the previous test, they experienced a similarly high deviation in the throttling test, letting more than 75 percent of Direct Connect traffic through.

The miscellaneous low-bandwidth protocols were throttled with mixed results; however, the total throttling accuracy for the complete P2P traffic was very close to the desired 50 percent, with the largest deviation at 56.2 percent by PacketLogic. HTTP traffic was not noticeably affected in any case.

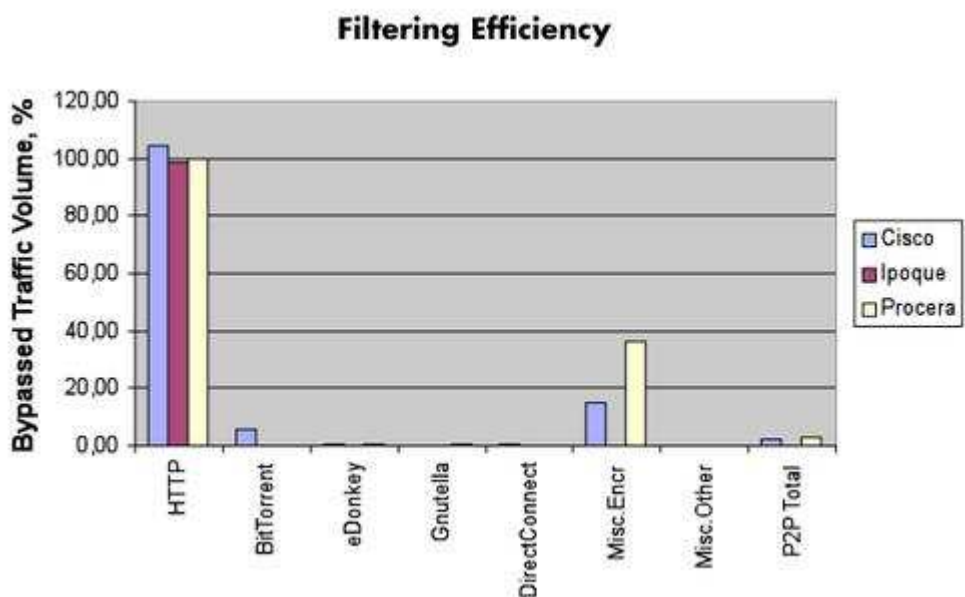


Next Page: [Filtering & ISP Constraints](#)

Filtering & ISP Constraints

Finally, using the same traffic profile, as in the previous test cases, we configured the devices under test to block every protocol in the P2P class. The devices had to block explicitly recognized protocols only and not merely block P2P-like behavior. By comparing the traffic results with the reference test, we could determine the blocking accuracy.

All three devices were able to perform blocking with high accuracy. The amount of P2P traffic able to bypass the filtering was only approximately 2.4 percent with the Cisco SCE-8000 and 2.9 percent with the Procera PacketLogic. Ipoque's blocking accuracy was, however, unmatched, with less than 0.01 percent of P2P traffic transmission compared to the reference test. Just as in the regulation test case, HTTP traffic was not affected in any of the cases.



In this series of tests we also analyzed how the device under test reacts to the less favorable conditions in an ISP infrastructure. Several circumstances may make inspection of traffic and detection and filtering of P2P protocols more difficult. In the following test cases we recreated various scenarios of these difficulties and tested the performance and accuracy of

the device.

Asymmetrical routing: This scenario is imaginable if a DPI device is used in a ring topology. This is an especially challenging situation, as the device under test will not be able to monitor the complete traffic flow, but instead only one direction of the traffic.

Two of the three participating vendors were ready for our test; one of the two did not approve the publication of the results, so we think it would be unfair to only present the results of the remaining vendor, which, besides, were not very accurate.

P2P detection in encapsulated traffic: In today's ISP networks, most traffic will be transmitted using some kind of encapsulation. Many networks are using VLAN or Q-in-Q traffic to differentiate among numerous aggregated networks. MPLS also becomes very popular.

For this test, we sent the same mix of HTTP and P2P traffic as in the detection test, but the Ethernet frames were encapsulated using 802.1q VLAN tagging. The device under test was configured in the detection mode and was expected to show the same throughput performance and detection accuracy as for the non-encapsulated traffic.

All three tested devices were able to process the VLAN-tagged traffic and showed no noticeable deviation from the normal detection test. Both performance and detection accuracy was preserved.

Next Page: [Conclusions](#)

Conclusions

This latest test cycle confirmed that DPI devices implement the requisite features in detection, regulation, and filtering, and offer the necessary performance and accuracy required in today's service provider environments.

In other words, whether we talk about intelligent management of consumer traffic or about freeing bandwidth by throttling the massive amount of P2P traffic, the devices we tested are ready to be rolled out in service provider backbones. The testing also revealed detailed background information useful for service providers in upcoming RFI/RFP scenarios, pointing to areas of questions and further investigations (e.g., protocol specific regulation/filtering or asymmetrical routing).

Unfortunately, we cannot repeat our hymns of praise for the content filtering part of the story. Our lesson learned regarding media content filtering solutions (as opposed to protocol-based filtering) is that vendors still shy away from public testing at today's scale of broadband Internet networks. However, we have indications to expect significant improvements in the next year's time frame and hope to include some of the media content filtering solutions in our next round of testing. Protocol-based filtering devices may be used in conjunction with content-based solutions to offload traffic and thus to ensure sufficient performance and accuracy to eliminate the very detrimental false positive incidents.

Performance figures and detection rates are *the* most important differentiators in the competitive DPI solution market. Our interviews with DPI vendors also indicated that device management, available traffic statistics, subscriber-based reporting features, and, last but not least, the flexibility in modifying and applying new application-based signatures are very important key requirements. From a design perspective, software solutions optimized for standard hardware compete with solutions based on specialized hardware. The software vendors claim that their pricing is more flexible and, in particular, profitable, comparing basic prices of the two solutions, whereas the hardware argument is speed.

Although we increased the application layer performance of our last year's test bed by a

factor of 25, we were surprised that we still did not hit the performance limits of the devices under test. The performance of DPI devices increased massively in the last year, retaining very good detection rates.

For additional performance boosts in future large-scale P2P test campaigns, test equipment vendors are already striving to increase the effective application layer traffic load per test module interface. This effective interface performance is very important from a test lab perspective, to avoid adding massive scale and very expensive aggregation devices into the test bed. New test areas can include performance and accuracy tests in the following areas: logging, subscriber database performance tests, billing and accounting, service prioritization, and bandwidth modification on demand.

In the near future, the integration of DPI functions into other network components is a challenge for vendors; a special focus is on mobile backhaul networks. Due to the restricted frequency spectrum in 2G/3G mobile networks, intelligent data traffic management will be an important factor for profitability, as more and more customers use dongles to connect their P2P-enabled PCs at home or laptops via mobile networks.

Next Page: [Testing Notes & Specifications](#)

Testing Notes & Specifications

Some additional notes and explanations on our terms, methods, and participating vendors:

Bandwidth & traffic types

For Layer 3 traffic tests we specified the bandwidth as the link bandwidth, which included the Ethernet header (14 bytes); Ethernet frame check sequence (4 bytes); preamble (8 bytes); and interframe gap (12 bytes). Layer 3 packets were used for stateless traffic (like multicast streams of video applications) without emulation of the TCP state machine.

For Layer 7 traffic we specified the bandwidth as the application level throughput as defined in RFC-2647 ("Benchmarking Methodology for Firewalls"). The Layer 7 generator was able to emulate individual TCP stacks per multiple sessions. This allowed application emulation containing TCP techniques like packet retransmission and stateful session handling.

Traffic direction

We differentiated between upstream and downstream traffic in our test bed setup and test configurations. The upstream traffic was transmitted to the device-under-test's network interfaces pointing to the customer side, while the downstream traffic went toward the device's interface pointing to the Internet. For common protocol traffic such as HTTP, all clients were situated on the customer side, and the servers on the Internet side, so most traffic flowed in the downstream direction. For the P2P protocols, approximately equal amounts of traffic flowed in each direction.

We defined bidirectional throughput as symmetrical flows, downstream and up. We used the term "total throughput" for asymmetrical traffic describing the sum of upstream and downstream traffic. For example, 10 Gbit/s of bidirectional traffic thus corresponds to 20 Gbit/s of total traffic, but also signifies that the traffic flow is (roughly) symmetrical.

Device-under-test operation mode

A device under test was characterized by the way it was integrated into the provider's network for purposes of monitoring or filtering. One device, ipoque's PRX-10G Traffic Manager, was operated in the so-called transparent bridging mode, where the device under test was transparently configured into the Ethernet link, forwarding or blocking the traffic without intervening into Layer 7 and transport protocols. In the regulating mode, the chosen traffic type may be either blocked completely or limited to certain bandwidth.

The device under test didn't act as an active IP router, but it used two integrated Ethernet

switches on each side for load balancing. For the surrounding infrastructure, the use of such a device is mostly indistinguishable from a direct link, apart from the traffic being filtered. The presence of the switches, however, should be taken into account and requires proper configuration of the device under test for a specific LAN environment.

The presented solution implements a protocol- and behavior-based detection of P2P traffic, with the former functionality based on protocol-specific signature analysis and, for some protocols, on heuristic analysis.

Traffic volume calculation

The evaluation of the detection accuracy and throttling and blocking efficiency in our tests was always based on the transferred traffic volume. We compared the traffic volume as reported by the load generator with the traffic volume statistics produced by the tested device in order to calculate the detection accuracy. For the calculation of the throttling and blocking efficiency, we compared the load generator statistics of the reference test and a test with the DPI device. These calculations were performed for each protocol separately, as well as for the total P2P traffic.

The statistics on the device under test were always reset before beginning the test and allowed to settle after the test end, in order to clear any stale connections and ensure that all transmitted traffic will be recorded in the results.

The load generator statistics had to be converted to the same units used in the device's statistics in order to allow direct comparison. All devices calculated the traffic statistics either on Layer 2 (Ethernet frames) or Layer 3 (IP packets). In order to convert the results of the load generator (calculated on Layer 7 as the number of bytes transported in the TCP connections) to the same format, we separately calibrated statistics for all used protocols.

Finally, here's a more detailed overview of product features and limitations, as reported by each vendor:

General Information

	Cisco Systems	ipoque	Procera
Device Name	SCE-8000	PRX-10G Traffic Manager	PacketLogic 10014
Software Versions	Version 3.1.6 build 014s	Firmware Version 2.8	Firmware Version 12.1.1.37
	Protocol Pack PP15 build 6	Signature Definitions 2.8	Signature Definitions S-3-20081215
	Classifier 3.1.6 B3s	Client Control Software 2.8	Client Control Software PLClient v12.1

Supported Physical Interfaces

Supported Feature	Cisco Systems	ipoque	Procera
1-Gbit/s Ethernet interfaces	None	None (optionally 12 instead of 10GigE)	8 ^{1,2} (Copper)
10-Gbit/s Ethernet interfaces	4	12 ³ (6 on each side)	8 (SFP+ Ports)
Maximum Throughput Claimed	40 Gbit/s	60 Gbit/s	80 Gbit/s
Maximum DPI processing capacity claimed	15 Gbit/s	60 Gbit/s	80 Gbit/s
Maximum Number of L4 flows claimed	2 million unidirectional flows	240 million	48 million unidirectional flows
Maximum Number of subscribers	250,000	5 million	5 million

1. Up to 8 further 1GigE interfaces can be added via SFP ports instead of 10GigE interfaces

2. Not used in the tests

3. Only 8 interfaces were used in the tests

Supported Networking Features (Vendor Claims)

Supported Feature	Cisco Systems	Ipoque	Procera
MPLS encapsulation	Yes	Yes	Yes
GRE encapsulation	No	Yes	No (planned for future release)
L2TP encapsulation	Yes	No	No (planned for future release)
802.1q VLAN tagging	Yes	Yes	Yes
Q-in-Q VLAN tagging	No	No	Yes
802.3ad Link aggregation	Yes	Yes	Yes
Transparent forwarding in case of failure	Yes	Yes (using external optical bypass)	Yes

Insertion Mode

Supported Feature	Cisco Systems	ipoque	Procera
Out-of-band monitoring	Yes	Yes	Yes
Transparent bridging	Yes	Yes	Yes
In-band monitoring	Yes	Yes	No
Proxy	No	No	No

See "Device-under-test operation mode" for detailed description.

P2P Detection Features Tested

Supported P2P Detection Mode	Cisco Systems	ipoque	Procera
Protocol recognition	Yes	Yes	Yes
Behavioral analysis	Yes	Yes	Yes
Content recognition by hash code	No (planned for future implementation)	Limited ⁴	No
Content recognition by media analysis	No	No	No
P2P client detection	No	No	Limited
P2P traffic filtering	Yes	Yes	Yes
Bandwidth regulation	Yes	Yes	Yes

4. The device under test supports selective blocking of the BitTorrent tracker requests based on the info_hash value in the request.

Supported P2P Protocols Tested

P2P Protocol	Enchr.	Cisco Systems			ipoque			Procera		
		Detection	Regulation	Blocking	Detection	Regulation	Blocking	Detection	Regulation	Blocking
BitTorrent	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
EDonkey		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Gnutella1		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Gnutella2		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DirectConnect		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IMesh		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MP2P		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FastTrack		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
WINMX		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Soulseek		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BitTorrent (encrypted)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
EDonkey (obfuscated)		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Arcs		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Freenet		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Filetopia		Yes	Yes	Yes	Yes	Yes	Yes	No	No	No

— Carsten Rossenhövel, *Managing Director*, [European Advanced Networking Test Center AG \(EANTC\)](#)

Channel: Consumer Internet, Digital content & entertainment, Enterprise IT, Personalization & privacy, Security, Telecom infrastructure, Telecom services

Tags: [Access technologies](#), [Data center/storage](#), [Digital-Physical Convergence Tutorial](#), [IP ...](#)

[DISCUSS](#)

[PRINT](#)



[DIGG](#)



[DELICIOUS.US](#)



[REDDIT](#)



[EMAIL THIS](#)