STATE OF OPEN-SOURCE DPI: CHALLENGES, OPPORTUNITIES AND ALTERNATIVES

Research Report

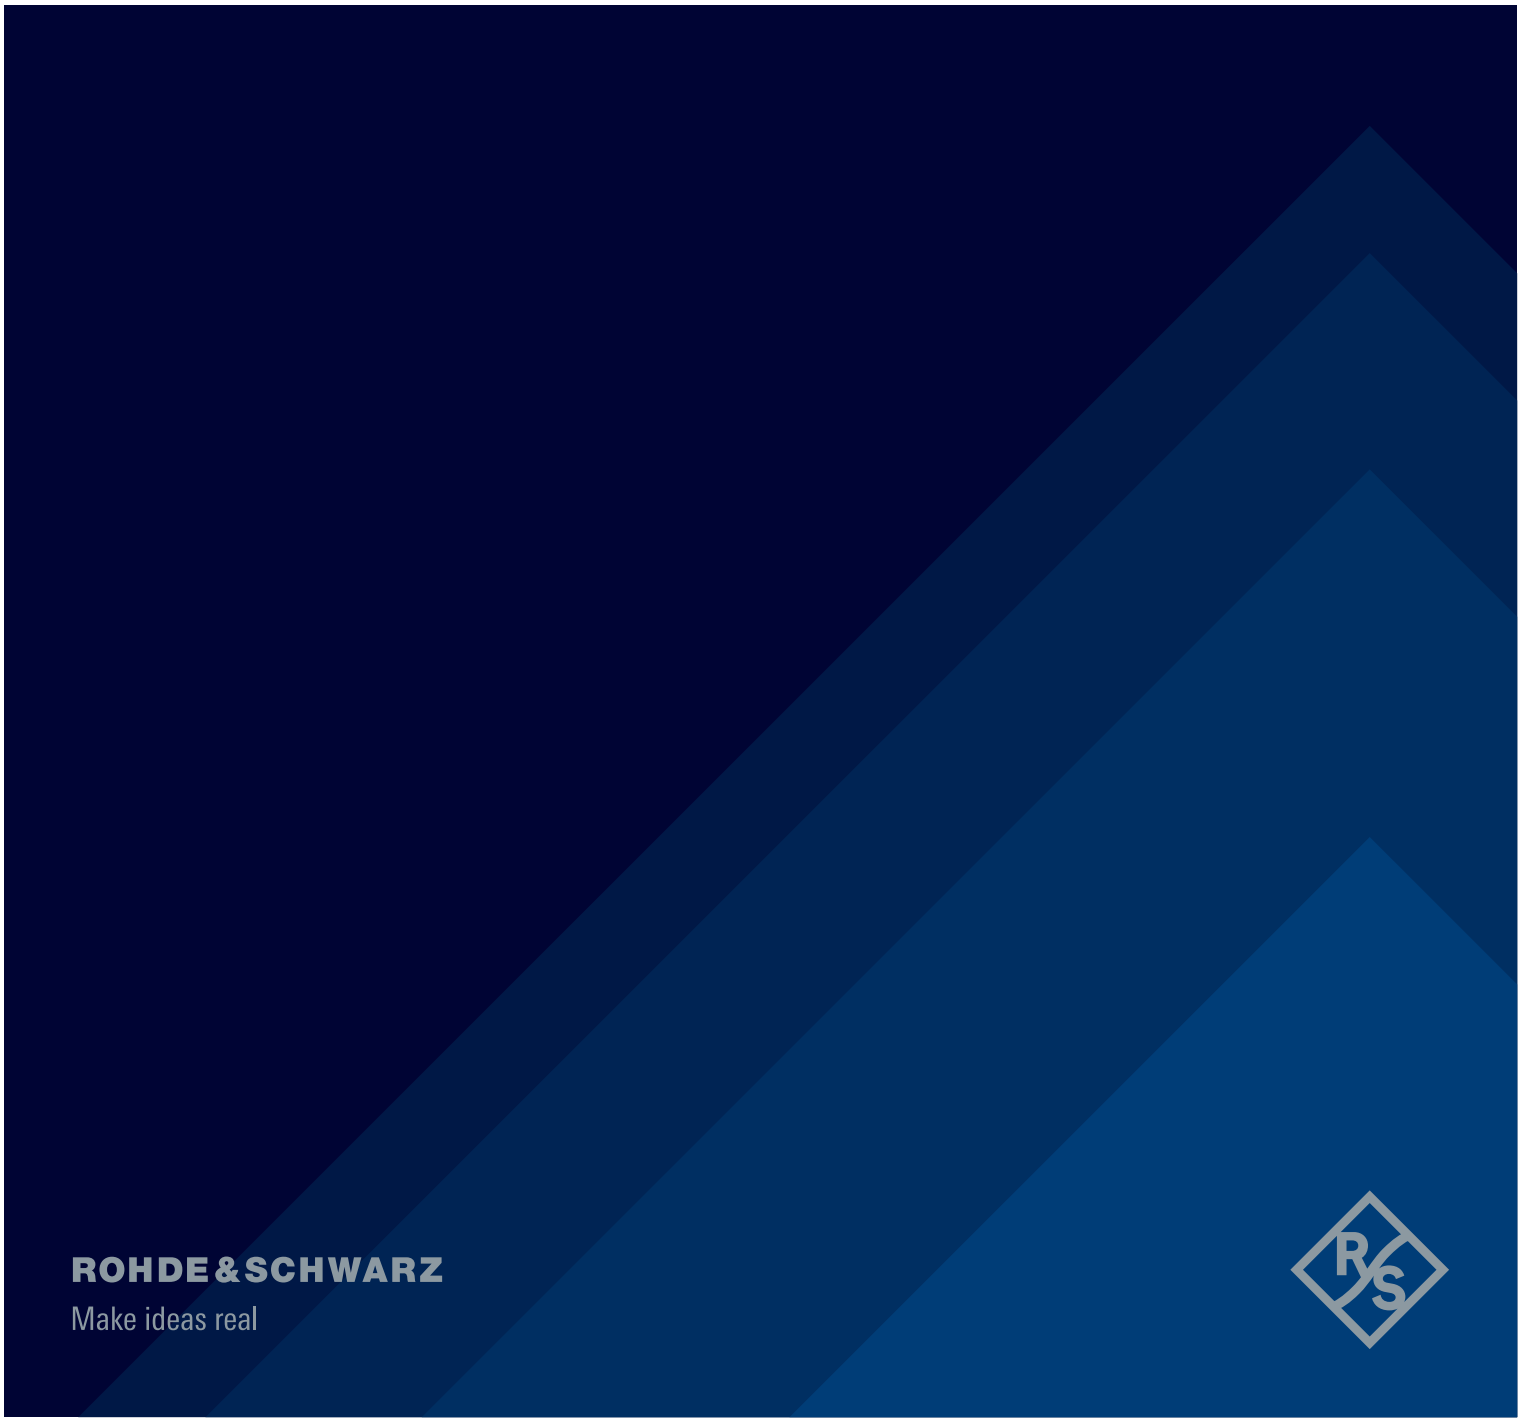

ROHDE&SCHWARZ

Make ideas real

CONTENT

1. INTRODUCTION

Deep packet inspection

Deep packet inspection (DPI) is a traffic detection technology that is used prevalently across today's IP networks to manage and secure traffic flows. DPI involves the identification of packet payloads, which enables networks to ascertain the underlying applications and service types. While shallow packet inspection confines traffic analysis to packet header data, such as port numbers and source URL, DPI deploys advanced analytical techniques and comprehensive signature libraries to deliver accurate application-level insights, enabling transparency through application Layer 7 and beyond. The use of DPI propagates intelligent traffic management via contextual and dynamic policies built on real-time traffic awareness, creating networks that are responsive and efficient. Security-wise, DPI's threat awareness equips cybersecurity vendors with real-time insights into malicious and anomalous traffic flows, mitigating the network's susceptibility to attacks and abuse.

Establishing the role and prospects of open-source DPI

This report aims to study the use of open-source DPI among networking and cybersecurity vendors and understand vendors' perspectives on its efficacy and long-term prospects. Open-source DPI refers to DPI source code or software that is licensed for public use by their developers. The code can be used, edited, customized, repackaged and distributed by vendors as part of their own solutions. Alternatives to open-source DPI are DPI technologies developed in-house and commercial DPI solutions provided by specialist vendors.

Over the years, the availability of various open-source DPI software has contributed to the uptake of DPI in general, with vendors capitalizing on their cost effectiveness, flexibility and ease of access. Boasting strong community participation and support from developers, open-source DPI has made DPI accessible to individuals and organizations, regardless of scale and purpose of use. Open-source DPI is commonly deployed in analytics, networking and cybersecurity solutions and telecom / ISP networks. Users of open-source DPI include solution vendors, enterprises, public bodies and research organizations.

Survey: State of open-source DPI

Duration: 03/24-04/24

Participants: 48 networking vendors

Authors: Rohde&Schwarz and The Fast Mode

Focusing on the use of open-source DPI by solution vendors, this report looks into the many pull and push factors that influence its take-up and usage. The report analyzes the capabilities offered by open-source DPI in terms of traffic classification, including its ability to identify emerging applications and traffic flows that are encrypted, obfuscated and anonymized. The ability to cater to specific use cases such as tethering detection and support for integration with other analytical systems such as IPFIX are also investigated. Additionally, the report examines the implication of open-source DPI on the exposure of networks to cyber threats.

Migration to commercial DPI

While the report aims to highlight the role and contribution of open-source DPI, it acknowledges the inherent need among fast-growing vendors to scale their analytical and monitoring capabilities by transitioning to commercial DPI solutions. The report consequently explores the availability of migration tools and their influence on vendor transitioning decisions, and presents the benefits of commercial DPI solutions.

This report is produced jointly by leading provider of next-gen DPI, ipoque, a Rohde&Schwarz company and The Fast Mode. It is based on a global survey that was conducted from March to April 2024, involving 48 leading networking, analytics and cybersecurity vendors who are experts in IP traffic management and security. They either have prior or current experience in using open-source DPI in their solutions and products.

2. TAKE-UP AND USAGE OF OPEN-SOURCE DPI

Open-source DPI enjoys the same level of popularity as commercial DPI

The growing importance of DPI in managing large-scale networks and complex traffic visibility use cases has resulted in the proliferation of a wide range of DPI solutions, encompassing hardware and software offerings. In terms of procurement, vendors can choose to install open-source software, license a commercial solution or build their own DPI technology in-house.

The survey assessed vendors' familiarity with different models of procurement. According to the results, open-source DPI ranks highest, with 41.7% of vendors being very familiar and another 25.0% being quite familiar. The remaining 33.0% say that they are somewhat familiar. Open-source DPI, which is based on freely available software and is supported by user communities, can be executed by internal teams or through paid support services from the maintaining developer. For most vendors in the initial stage of using DPI, open-source software provides a viable option for experimenting with DPI and assessing its effectiveness for their solution portfolio.

Vendors surveyed also show a high degree of familiarity with commercial DPI solutions, with 35.4% and 31.3% of vendors agreeing that they are very familiar or quite familiar, respectively. In commercial DPI, specialist providers lend their expertise to networking and cybersecurity vendors via ready-to-use DPI solutions that can be embedded and customized according to their requirements. Commercial DPI requires a paid license.

The survey shows a lower level of familiarity among vendors across in-house DPI solutions, with only 27.7% and 25.5% of vendors stating that they are very familiar and quite familiar, respectively, with such implementations. In in-house models, vendors leverage their own teams to develop DPI methodologies that are built natively for their solutions, to cater for custom use cases and architectures. More than one fifth (21.3%) of vendors are not familiar at all with in-house DPI, compared to only 10.4% for commercial DPI.

DIAGRAM 1 Familiarity with DPI solutions

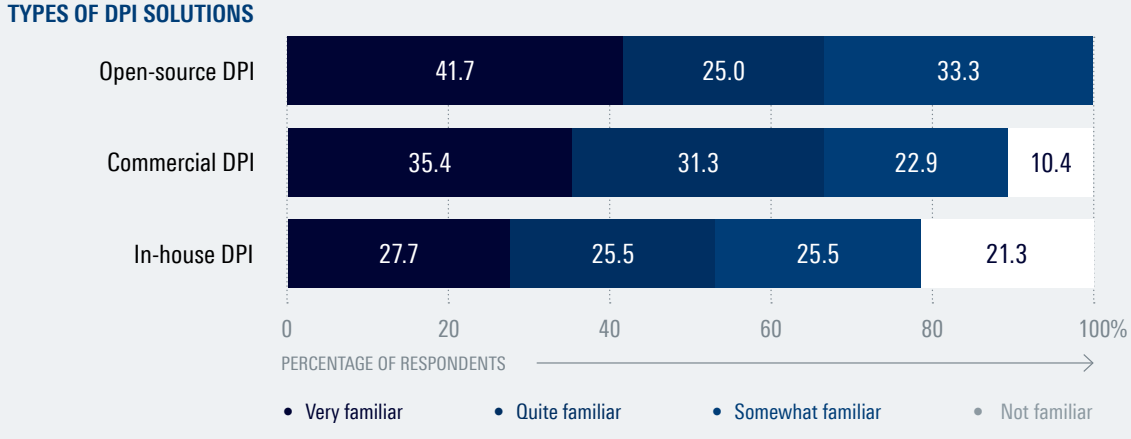

Vendor neutrality and low initial cost two biggest factors driving open-source DPI

Open-source DPI boasts several key benefits, primary of which are vendor neutrality and low initial costs. Vendor neutrality facilitates interoperability and grants vendors the freedom to develop the best mix of features and tools for their analytics systems. Based on the survey, the majority of 62.5% of vendors strongly agree that open-source DPI promotes vendor neutrality.

The next biggest benefit is the low initial cost of open-source DPI, with more than half (52.1%) of vendors strongly agreeing to this, and another quarter (25%) moderately agreeing. With zero licensing fees, open-source DPI results in substantially lower upfront expenditure.

Another key advantage of open-source DPI is support from a wider community. However, only one third of vendors (33.3%) strongly agree to this, while 41.7% of vendors moderately agree. Support from a wider community, via online networks, user forums and ground events, provides users of open-source DPI accessible avenues to seek advice and guidance.

The survey also finds that less than a third of vendors (31.3%) strongly agree that open-source DPI enables fast deployment times due to codes being instantly accessible. Another half of vendors (50%) moderately agree.

Software form-factor offered by open-source DPI is another benefit that is cited by vendors. More than a quarter (27.1%) of vendors strongly agree, and close to half (47.9%) of vendors moderately agree to DPI's software form-factor being an advantage. Software modules can be integrated into any architecture – physical, virtualized or cloud-native – making open-source DPI compatible in all environments.

Less than a quarter (22.9%) of vendors strongly agree that open-source DPI offers a lightweight implementation, for example, by using mature statistical models instead of complex machine learning algorithms. More than a third (39.6%) of vendors moderately agree to this.

DIAGRAM 2 Key benefits of open-source DPI

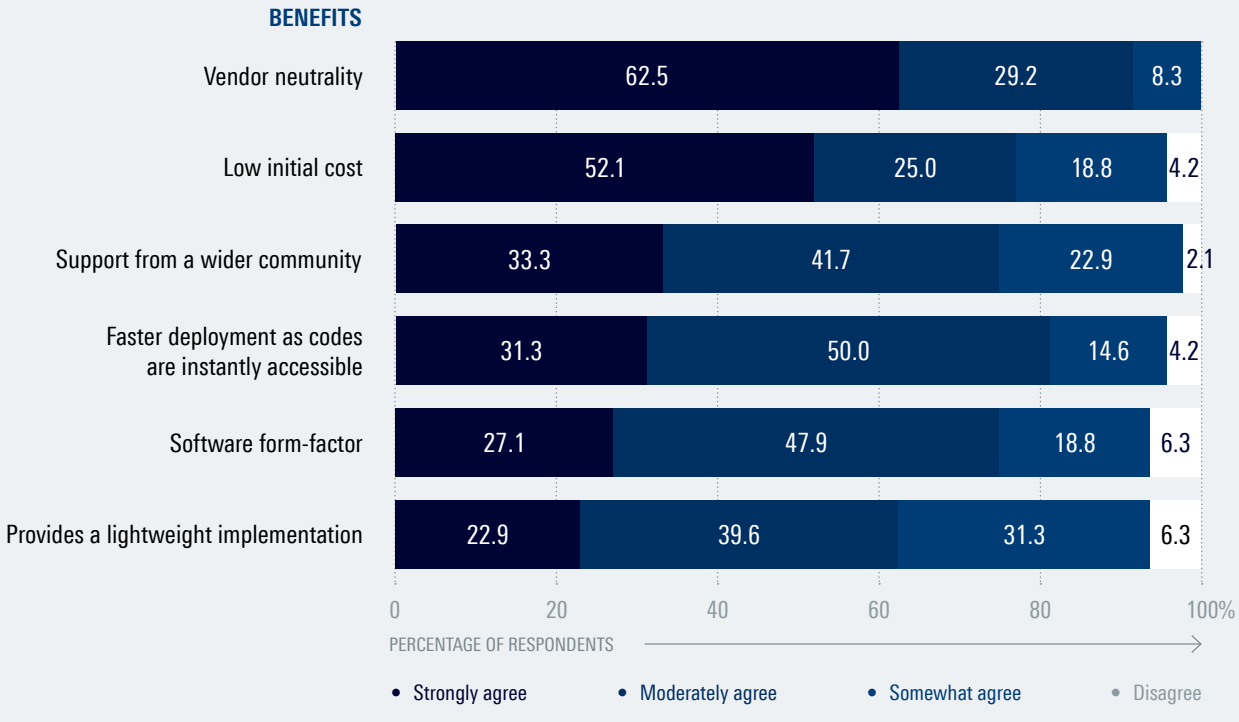

3. ENABLING GRANULAR TRAFFIC ANALYSIS WITH APPLICATION AND THREAT AWARENESS

Protocols are classified comprehensively with open-source DPI; visibility gaps persist across applications, service types and threats

A robust DPI tool combines different layers of monitoring to deliver deep insights into traffic flows. Application awareness is a focal point of analysis in DPI and is executed via advanced methods such as statistical and behavioral analysis. Application awareness enables granular classification of traffic, for example, by protocol, application and service type. Metadata extraction is another form of analysis in DPI, where packet header information is used to understand traffic flows. Metadata extraction, when combined with application awareness, delivers fine-grained application-level and packet-level insights covering various transmission (e.g. session duration and bandwidth) and performance (e.g. speeds and latency) metrics.

A whopping 78.3% of vendors surveyed agree that open-source DPI can identify protocols (e.g. HTTPS, SMTP, FTP and WebRTC) comprehensively. The remaining 21.7% agree that it delivers moderate levels of identification for protocols.

However, when it comes to applications (e.g. Zoom, Facebook, Telegram, TikTok and Netflix), only 28.3% of vendors agree that open-source DPI delivers comprehensive identification, while 52.2% say that it classifies applications moderately. Similarly, only 28.3% of vendors agree that open-source DPI can classify service types (e.g. messaging, video streaming, file downloads) comprehensively, while 50.0% of vendors agree that its capabilities in this aspect are moderate.

DIAGRAM 3 Classification capabilities of open-source DPI

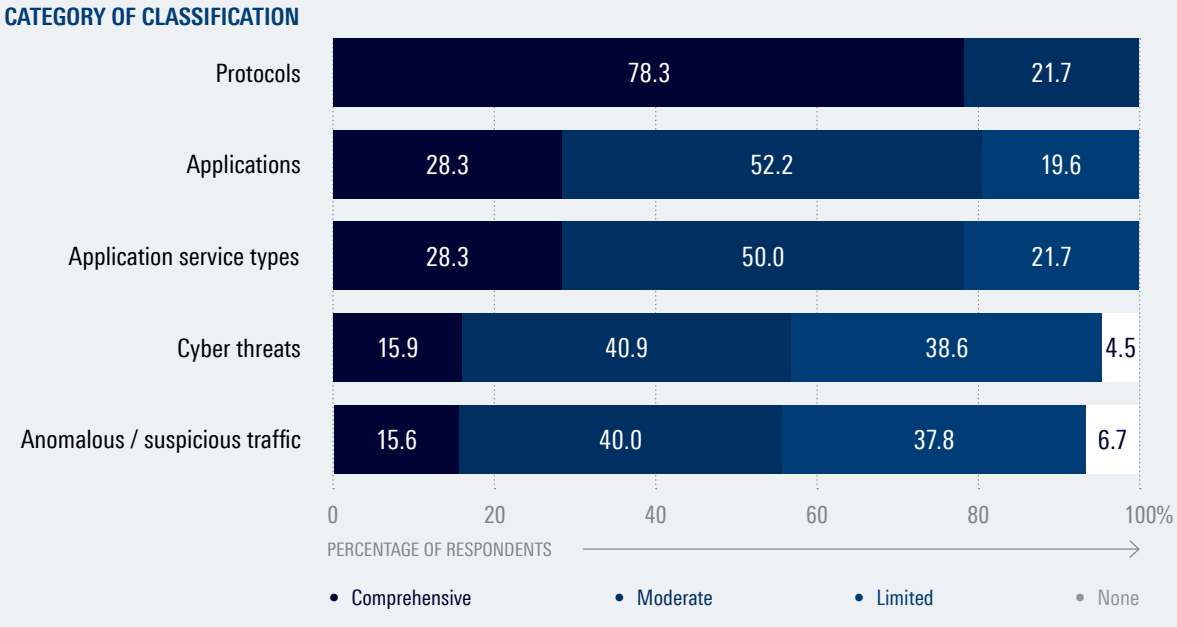

DIAGRAM 4 Classification accuracy of open-source DPI

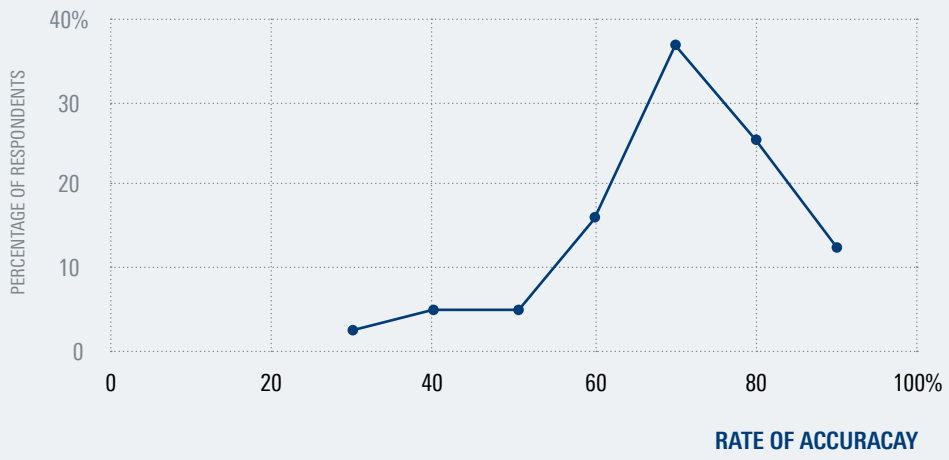

In terms of cyber threats (e.g. spear phishing, DDoS, code injections, malware and DNS tunneling), the share of vendors who believe that open-source DPI delivers comprehensive classification is only 15.9%, while the majority of vendors (40.9%) believe that its level of identification is moderate. The same observation is seen in the detection of anomalous and suspicious traffic (e.g. unknown URLs, unknown devices). Only 15.6% of vendors are of the opinion that the classification provided by open-source DPI is comprehensive, while 40.0% rate it as moderate.

Open-source DPI scores 70.2% on classification accuracy; vendors report varying accuracy rates

To understand the reliability of open-source DPI software, vendors were asked to rate its level of classification accuracy (using a scale of 0-100%) against protocols, applications, service types, threats and anomalies. The findings show an average classification accuracy of 70.2%. The highest rating provided by the vendors is 95%, while the lowest is 30%. These results indicate accuracy issues inherent in open-source DPI and also varying levels of trust among vendors over classification results provided by open-source DPI.

Strong use of statistical, behavioral and heuristic analyses in open-source DPI; limited use of AI-based techniques such as ML and DL

Different DPI tools engage a different mix of techniques to classify traffic, engaging the latest advancements in statistical modeling, computing and networking. A basic type of analysis used in DPI is pattern matching, where sequence of packets and other readily observable attributes, such as packet size, are used to identify the underlying application. Pattern matching is used prevalently in open-source DPI, with 95.5% of vendors agreeing to it being significantly or moderately used.

A more advanced technique used in DPI is statistical analysis. This involves analysis of large, available datasets using statistical inferencing and probability distributions. Statistical inferencing includes techniques such as regression analysis, multivariate correlation and associations, while probability studies quantify the likelihood of the underlying packets representing a particular application. More than 74.4% of vendors report statistical analysis being used significantly or moderately in classifying traffic flows in open-source DPI.

DIAGRAM 5 Classification techniques used by open-source DPI

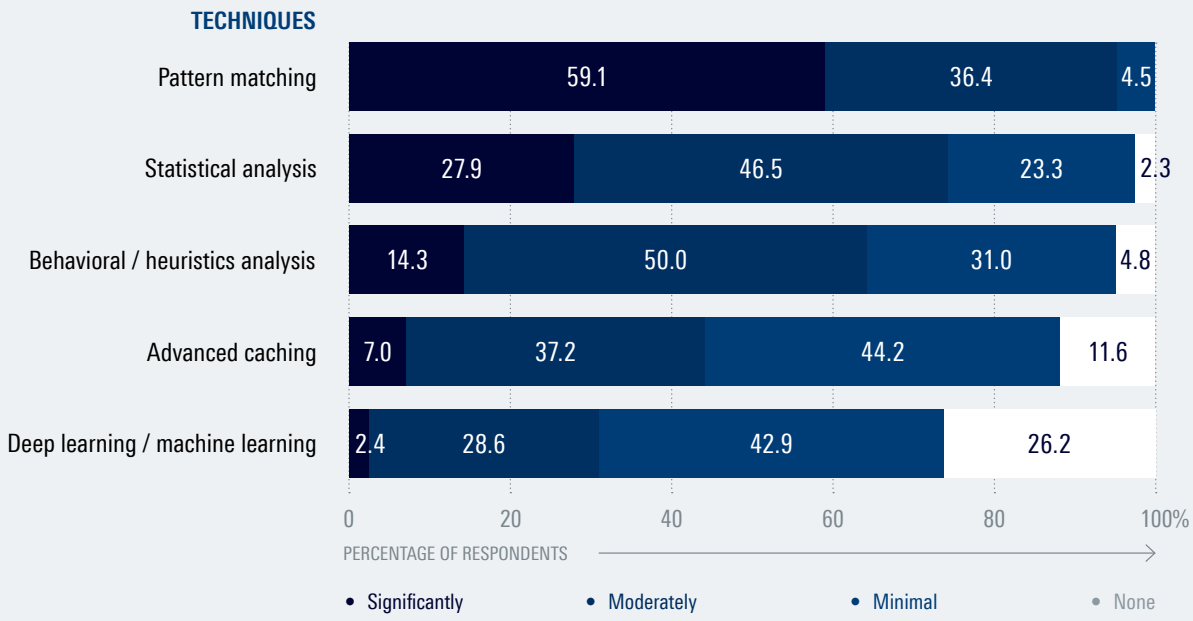

Behavioral analysis looks at how sessions, flows and packets move and behave, end-to-end. This requires the observation of transmission patterns and user interactions. Another method that is particularly useful in traffic classification is heuristics, especially in scenarios where past data is unavailable or insufficient. Heuristics involves trial-and-error, based on the best available knowledge. Heuristic analysis can be used to identify emerging applications and protocols, and also to make educated guesses on threats. Behavioral and heuristic analyses are used significantly or moderately in classifying traffic flows in open-source DPI, according to 64.3% of vendors.

Another important technique in classifying traffic flows is advanced caching. Caching enables frequently retrieved past results to be stored close to traffic filtering points. Advanced caching updates a cache dynamically based on context, ensures consistency across multi-node caches and provides tiered caching based on application priority. This speeds up classification of commonly used applications. The survey shows that only 44.2% of vendors report significant or moderate use of advanced caching techniques in open-source DPI.

With more traffic being tunneled, encrypted or anonymized, advanced DPI engines are turning to AI-based techniques such as machine learning (ML) and deep learning (DL) to restore lost visibility. ML uses known traffic attributes to develop identification algorithms from large data sets, removing the need for developers to define each application manually. DL goes a step further by using neural networks to extract these attributes automatically. The use of AI in open-source DPI is rather limited, according to the survey, with only 31.0% of vendors agreeing that AI is used significantly or moderately in classifying traffic in open-source DPI.

More than half of vendors think open-source DPI signature libraries are not comprehensive and not updated frequently enough

At the heart of DPI is its signature library. A signature library is developed via continuous analysis of global and local traffic flows across various networks – enterprise, mobile, ISP, IIoT and others. Static and dynamic attributes of packets and flows are correlated with each other to create 'thumbprints' that are used as baselines to classify traffic. The efficacy of a DPI tool essentially depends on the breadth and depth of its signature base.

According to the survey, more than half (52.3%) of vendors think that signature libraries provided by open-source DPI lack comprehensiveness or are not updated frequently enough. This can negatively impact the coverage and accuracy of classification offered by open-source DPI, leading to many blind spots in network analytics.

To provide full coverage of traffic flows, DPI signature libraries must at least encompass commonly used protocols (e.g. HTTPs, SMTP), applications (e.g. Google, Salesforce and X), service types (video, messaging and email) and threats (e.g. spyware, DDoS and ransomware), and be able to identify traffic patterns that are anomalous or suspicious. Signature libraries must also be frequently updated. Protocol upgrades, compression, caching, use of proxy servers, encryption and adoption of multi-cloud architectures lead to changes in traffic signatures, and this must be reflected in the signature libraries.

DIAGRAM 6 Comprehensiveness and update frequency of open-source DPI signature libraries

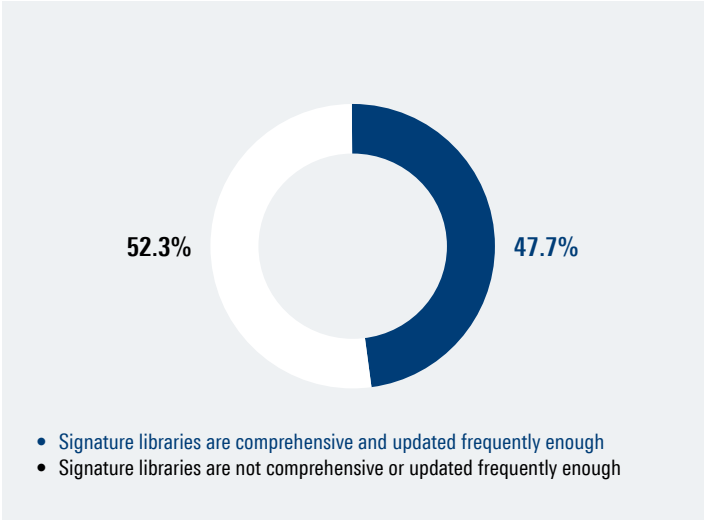

Majority of vendors believe open-source DPI extracts adequate metadata to support common DPI uses cases

Along with traffic classification, DPI delivers metadata extraction. This allows networks to be aware of traffic attributes such as destination and source URL addresses, speeds, jitter, latency, time-to-first-byte, round-trip-time, inter-packet delay and packet size. The ability to log unlimited data points in real-time is key to metadata extraction. Combining this data with traffic classification enables granular traffic analysis of, for example, traffic throughput and latency by application, service type and session.

The survey assessed the adequacy of metadata provided by open-source DPI in supporting common DPI uses cases such as next-gen firewalls, SD-WAN, SASE and network packet brokers. A significant share of vendors (40.9%) somewhat agree that the metadata extracted by open-source DPI is sufficient, followed by another 34.1% of vendors who moderately agree to this. The share of vendors who strongly agree, however, is less than a quarter (20.5%). A small percentage (4.5%) think that the metadata extracted by open-source DPI is inadequate in supporting these use cases.

DIAGRAM 7 Adequacy of traffic metadata provided by open-source DPI in supporting common DPI use cases

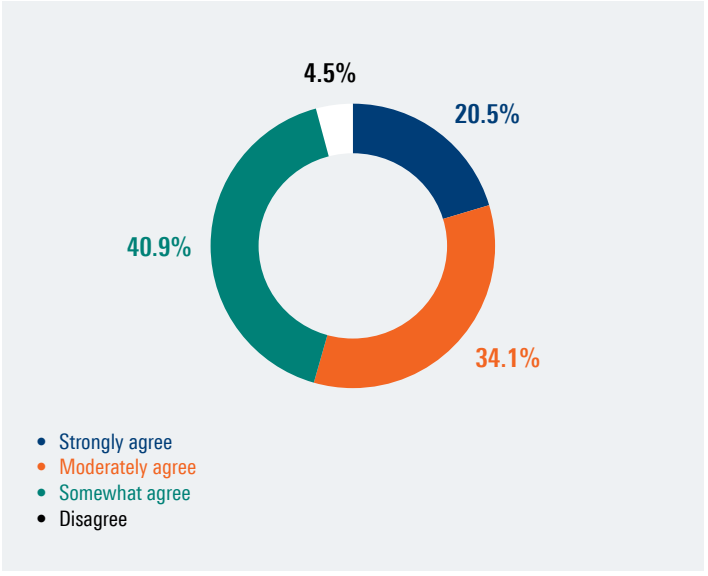

4. CHALLENGES FROM ENCRYPTION, OBFUSCATION AND ANONYMIZATION

New encryption protocols lead to visibility loss in open-source DPI; most challenging protocols are TLS extensions and DoX

The survey assessed open-source DPI in terms of its ability to classify encrypted traffic. Encryption involves the use of complex encryption algorithms to convert plain text into cypher text. Without encryption keys, encrypted applications such as Skype or Microsoft Teams become unreadable. In the context of privacy and confidentiality, encryption safeguards data-in-transit from interception.

Encryption, however, compromises traffic visibility and renders most traditional DPI techniques ineffective in classifying applications. The survey assessed different encryption protocols and how each are handled by open-source DPI.

According to the vendors surveyed, the least challenging encryption protocol is TLS / SSL (TLS 1.0 to 1.3) with 73.3% of vendors agreeing that open-source DPI can detect these flows. This is followed by QUIC and other encrypted Layer 7 protocols (SRTP, S/MIME, DTLS, SSH, SNMP). For both categories, 72.7% of vendors agree open-source DPI is capable of identifying the underlying flows. Other protocols with a similar implication to traffic visibility are VPN protocols (e.g. IPsec, OpenVPN, SSTP), with 71.1% of vendors agreeing that open-source DPI is capable of handling them.

DIAGRAM 8 Open-source DPI support for classification of encrypted traffic

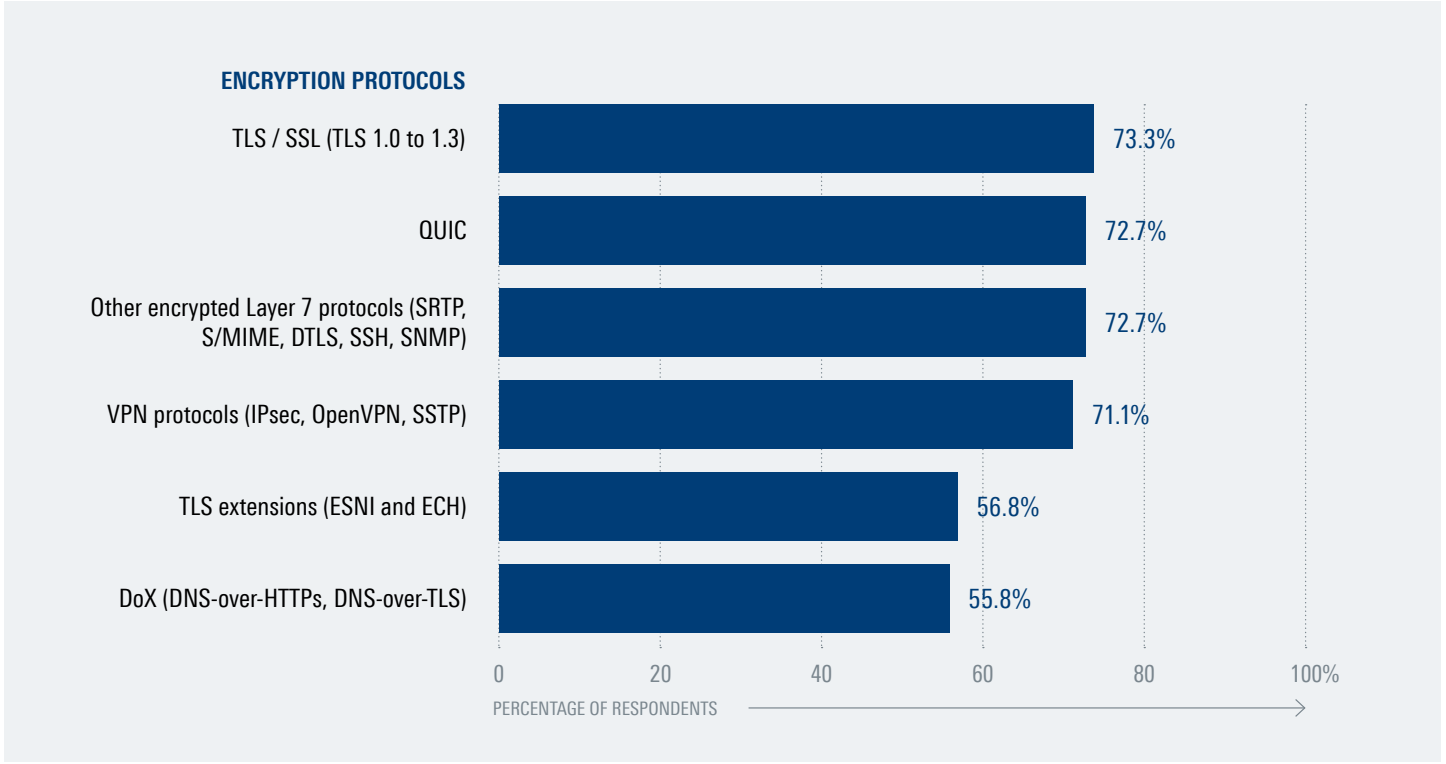

Two protocols, which are found to be more challenging, are TLS extensions (e.g. ESNI and ECH) and DoX (e.g. DNS-over-HTTPs, DNS-over-TLS). The share of vendors who agree that existing open-source DPI software provides adequate visibility into these protocols is much lower, at 56.8% and 55.8%, respectively. These results indicate several gaps in open-source DPI in identifying encrypted traffic.

The introduction of new encryption protocols progressively erodes more information from network monitoring tools. ESNI for example, hides the name of the server, while ECH hides all handshake information. As a result, open-source DPI technologies that are not continuously enhanced with advanced inspection techniques to read encrypted flows can lead to significant gaps in traffic monitoring and analyses. Not only does this cripple any form of application-based traffic steering and resource allocation, it also increases the risks from encrypted threats such as encrypted botnets and malware.

Traffic obfuscation and anonymization worsen transparency woes in open-source DPI

Apart from encryption, techniques such as obfuscation and anonymization are also frequently used to protect and mask data flows. Anonymization is used to protect the identity of the user. The use of anonymization techniques such as CDN, VPN and Tor enables users to conceal their IP addresses and access content or sites that are otherwise prohibited or not accessible in certain jurisdictions. Obfuscation, on the other hand, involves techniques that disguise traffic. Examples are domain fronting, tunneling and mimicry.

Only half (50.0%) of vendors agree that open-source DPI provides visibility into anonymized traffic. The corresponding share for obfuscated traffic is even lower, at 40.5%. These results indicate an obvious shortfall in open-source DPI's ability to tackle anonymization and obfuscation, which worsens existing visibility issues. Consequently, networks may have to revert to blanket policies for traffic routing and forwarding. It may also lead to networks either abandoning application-based security routines or becoming overzealous with security rules, thus affecting application performance and incurring higher overall overheads for the network.

DIAGRAM 9 Open-source DPI support for classification of anonymized and obfuscated traffic

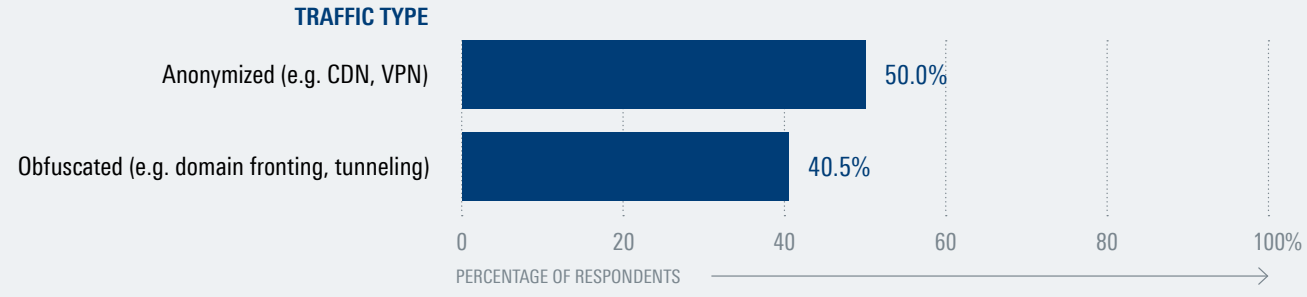

5. ADDED FEATURES AND CAPABILITIES IN OPEN-SOURCE DPI

Majority of vendors expect more in terms of technical capabilities offered by open-source DPI

DPI's biggest value proposition is its ability to classify and analyze large volumes of traffic in real-time. Vendors typically embed DPI in their proprietary hardware and software solutions, such as routers, switches or next-gen firewalls (NGFW). DPI is also deployed as a standalone function, integrated on virtualized or cloud-native networking architectures, for example, in the mobile core network or a data center network. As a standalone function, DPI analysis can be fed into multiple applications or network functions. Dedicated DPI hardware is also available but less popular.

Regardless of how it is packaged and delivered, a DPI solution must ensure performance and scalability, while catering for the reporting requirements of the use case and the traffic type it serves. These technical requirements apply to all forms of DPI, including open-source DPI.

Based on the survey, 8.7% and 52.2% of vendors rate memory efficiency offered by open-source DPI as excellent and quite satisfactory, respectively. The corresponding values for its throughput and speeds are 13.0% and 41.3%. Close to one third of vendors rate both capabilities as somewhat satisfactory, and the remaining 6.5% rate these as poor.

DPI must also cater for dynamic process upgrades which ensure live updates / weekly software releases are integrated during runtime without any interruptions. This feature is rated excellent by 13.3% of vendors and quite satisfactory by another 35.6% of vendors. Similarly, support for VNFs and CNFs is rated excellent and quite satisfactory by 9.1% and 36.4% of vendors, respectively. Both these capabilities however, are rated poor by close to one fifth of vendors (17.8% and 18.2%, respectively).

DIAGRAM 10 Technical capabilities of open-source DPI

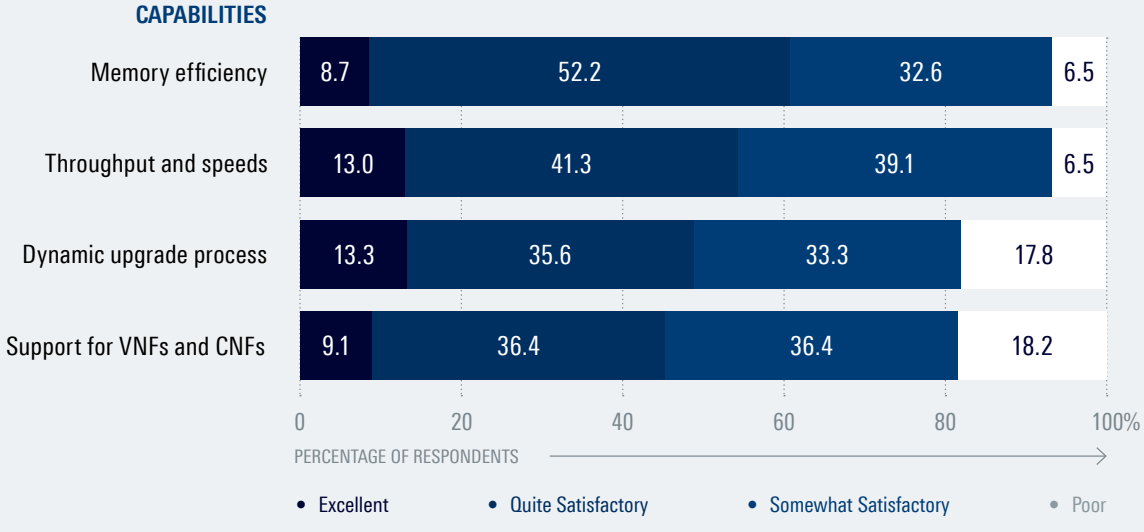

With approximately half of vendors rating each of the listed capabilities as somewhat satisfactory or poor, the findings point to a shortfall between the expectations of vendors and the actual performance of open-source DPI. The findings also point to open-source DPI's potential limitations in complex use cases which require intensive filtering and analysis.

Open-source DPI offers critical add-on features including custom classification, first packet classification and IPFIX integration

Apart from traffic classification and metadata extraction, DPI solutions typically include a number of added features that ensure their adaptability and effectiveness in various environments. According to 77.8% of vendors, one of such features in open-source DPI is custom classification. Custom classification refers to users' ability to add their own signatures based on internal reporting requirements and application usage patterns.

Another capability offered by open-source DPI, as stated by 77.8% of vendors, is first packet classification (FPC). FPC promotes consistency in how packets in a flow are treated, with policies implemented from the first packet. FPC also enhances network performance by allowing network devices to commence processing instantaneously and cuts down filtering overheads from full-fledged inspections.

One of the most valuable add-on features in DPI is seamless integration with other data collection tools, such as Netflow or IPFIX. This feature is critical in unifying traffic analysis across different networking devices. Three quarters (75.0%) of vendors surveyed claim that open-source DPI offers this integration.

To gauge vendors' take on other supplementary functionalities offered by DPI, participating vendors were asked about tethering detection. Close to half (45.2%) of vendors say that open-source DPI can detect tethering. This allows mobile operators to identify unauthorized sharing of network capacity and is delivered by analyzing related protocols and device data.

Open-source DPI offers various levels of customization; more than one third of vendors say it enables full customization

Customization of a DPI tool to the host solution, organization, region or industry is important in ensuring complete traffic coverage and an optimized implementation that balances performance and costs.

According to 44.7% of vendors surveyed, open-source DPI offers full customization by solution types. This enables it to be applied to any networking, cybersecurity and analytics use case, across varying traffic intelligence requirements, making it viable for a wide range of end solutions such as routers,

DIAGRAM 11 Critical add-on features supported by open-source DPI

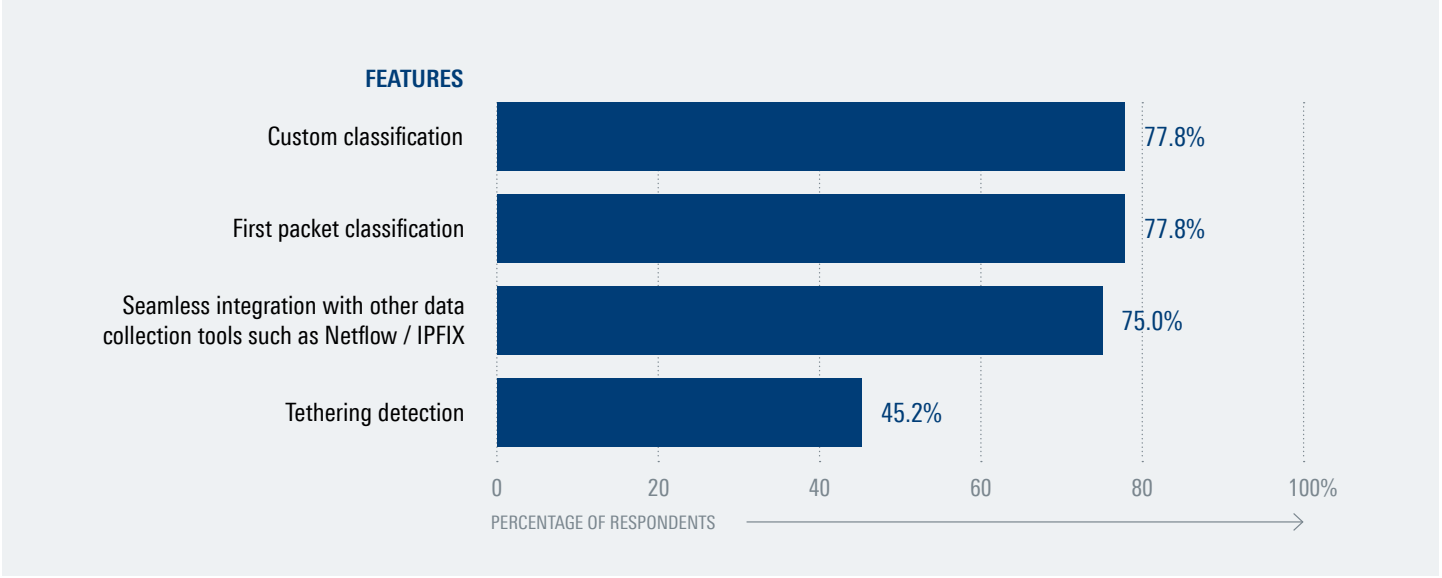

policy control engines, secure access service edge (SASE), zero-trust network access (ZTNA) and IP probes.

Similarly, 38.3% of respondents think that open-source DPI provides customization by the user organization, based on their network size, traffic volumes, monitoring KPIs and analytics requirements.

More than one third of vendors (36.2%) agree that open-source DPI enables full customization by region, therefore catering for local traffic trends.

Open-source DPI also provides full customization by industry types, as stated by 34.0% of respondents. This enhances its application across industries with specialized requirements. Examples are industries that deal with ultra-low latency applications and huge traffic loads such as telecoms, and industries handling niche applications, for example, smart manufacturing.

DIAGRAM 12 Degree of customization offered by open-source DPI

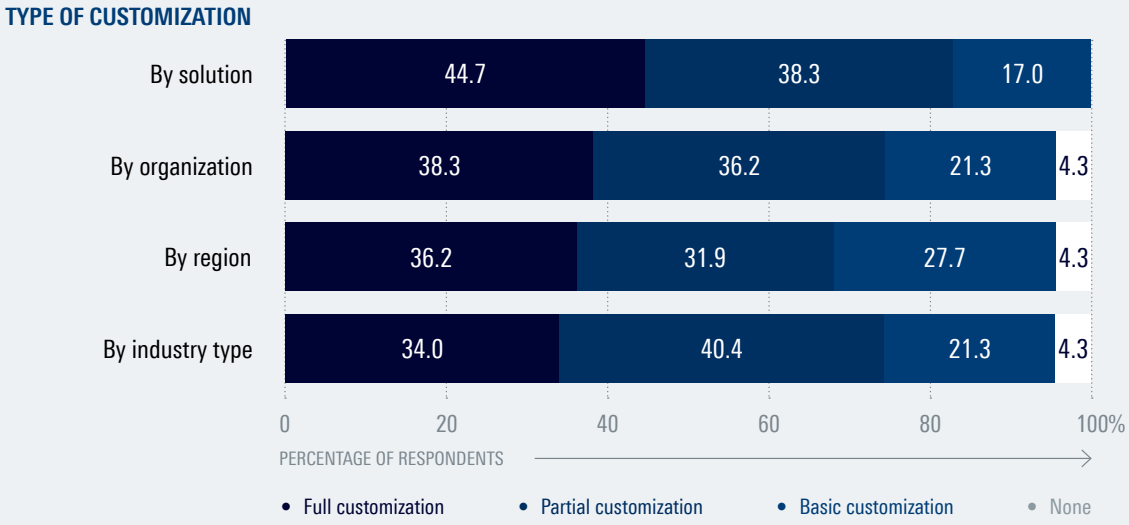

6. FACILITATION AND LONG-TERM SUPPORT

Open-source DPI boasts easy integration, API stability and long-term continuity but performs poorly in terms of support, customer service and global presence

Deployment facilitation is important in ensuring fast and seamless execution of DPI. Given that open-source initiatives are targeted primarily at promoting a technology and spurring innovation, the level of deployment facilitation provided depends heavily on the involvement of the maintaining organization, the developer community and its network of users.

According to the survey, more than a third of vendors (34.8%) rate open-source DPI as excellent in terms of ease of integration across any solution or architecture, while another 37.0% rate this feature as moderate. API and API stability offered by open-source DPI is also rated well. More than a quarter (26.1%) of respondents rate this feature as excellent, and half (50.0%) of respondents rate it as moderate.

In terms of continuity in the long run, open-source DPI performed moderately well, with 21.7% and 52.2% of vendors surveyed rating it as excellent and moderate, respectively.

Continuity in the long-run refers to an active initiative that continues to exist for 10 years or more.

The findings of the survey also reveal a few downsides to open-source DPI in terms of deployment facilitation. Open-source DPI does not have adequate global presence, according to the vendors surveyed with only 21.7% and 26.1% of vendors rating it as excellent or moderate, respectively. The remaining 52.2% rate its global presence as limited or poor.

Likewise, the survey finds support and customer service severely lacking in open-source DPI. A significant majority (71.8%) of vendors say that support and customer service provided by open-source DPI is limited or poor. Only 10.9% and 17.4% of vendors think that open-source DPI provides support and customer service that is excellent or at least moderate.

DIAGRAM 13 Open-source DPI deployment advantages

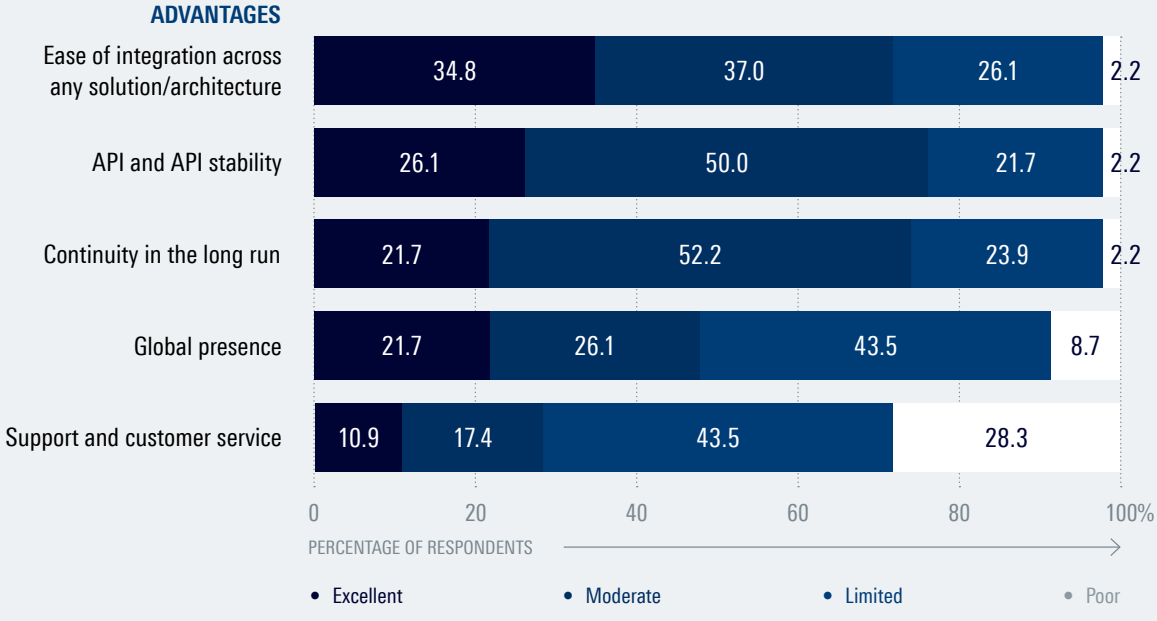

7. COST AND SECURITY IMPLICATIONS

Losses from issue resolution delays and extensive customization costs can outweigh initial savings in open-source DPI

One of the strongest merits of open-source DPI is that it does not involve an upfront licensing fee. This is a strong influencing factor for its take-up. The survey assesses the overall cost implications of using open-source DPI by evaluating all direct and indirect costs that come with its use, not only at inception but also in the long-run.

The survey results show issue resolution delays in open-source DPI to be a major concern among vendors, cost-wise. Backed by lean setups, open-source DPI initiatives are typically limited by resources which results in a lack of dedicated support. Consequently, users end up relying on user forums, support networks and independent consultants to troubleshoot and fix issues, which is time-consuming and can result in business losses. According to the survey, half of the vendors surveyed (50.5%) expect delays from limited support to result in high or very high costs.

One of the major outlays post installation in open-source DPI is customization costs. These come from extensive configurations and modifications to cater to the underlying use cases, applications and organizations. For example, a load balancer and an application performance monitoring solution have very disparate visibility requirements and would require different resource and output configurations for an optimal implementation. Similarly, different industries and enterprises have different reporting frequencies, priorities, analytical requirements and transparency policies. Close to half (46.8%) of vendors admit that customization costs in open-source can be high or very high.

Of a lesser concern in terms of cost are operational costs. According to the survey, only 29.8% of vendors find operational costs in open-source DPI high or very high. Close to half (44.7%) find these costs moderate. Operational costs relate to software management and maintenance, which include resource allocation, IT stack optimization and performance monitoring.

DIAGRAM 14 Cost of open-source DPI in the long-term

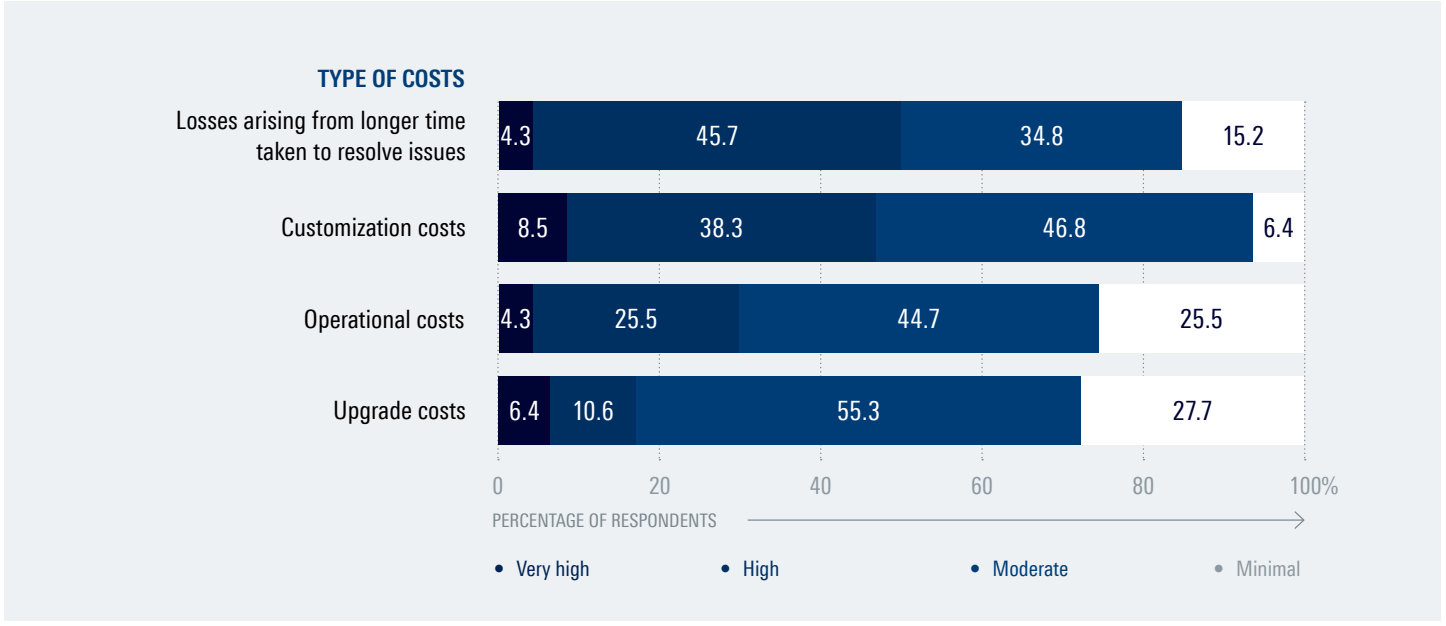

Upgrade costs are found to be least concerning among vendors. The survey shows that the share of vendors who find these costs high or very high in open-source DPI are only 17.0%, while the majority of vendors (55.3%) think that these costs are moderate. Upgrade costs refer to expenses incurred in migrating to the latest version of a software.

Loopholes created during internal customizations and undiscovered bugs – errors are major security gaps in open-source DPI deployments

Deploying a new software can expand an organization's attack surface and introduce new security risks. This can happen during installation and execution, code modification, changes in access and usage rights and configuration of APIs, and during day-to-day operations. In an open-source DPI deployment, risk exposure is influenced by the vulnerabilities in the software and the host solution (e.g. an SD-WAN gateway), as well as the effectiveness of the security measures (e.g. testing, quality control and feedback mechanisms) put in place by the maintaining organization.

The survey aimed to understand the security gaps in open-source DPI deployments. The findings show that loopholes created during internal customizations are the biggest concern, based on 17.8% of vendors who think it is a major issue, and 35.6% of vendors who see it as quite an issue.

This is followed by software bugs or errors that remain undiscovered, with 13.3% of vendors finding it a major issue and 40.0% agreeing it is quite an issue.

Adoption of unverified or untested versions is another contributor to increased risks, according to 11.1% of vendors who admit that it is a major issue, and another 40.0% who say that it is quite an issue.

Loopholes created during the integration process also increases the organization's security risk exposure. According to 8.9% of vendors, this is a major issue, while 37.8% think that it is quite an issue.

Surprisingly, codes being openly accessible to anyone does not pose much of a risk. More than half (52.2%) of vendors say that it is not an issue, and 23.9% find it only a minor issue. The share of vendors who see it as a major issue or quite an issue is only 23.9%.

DIAGRAM 15 Security gaps arising from the use of open-source DPI

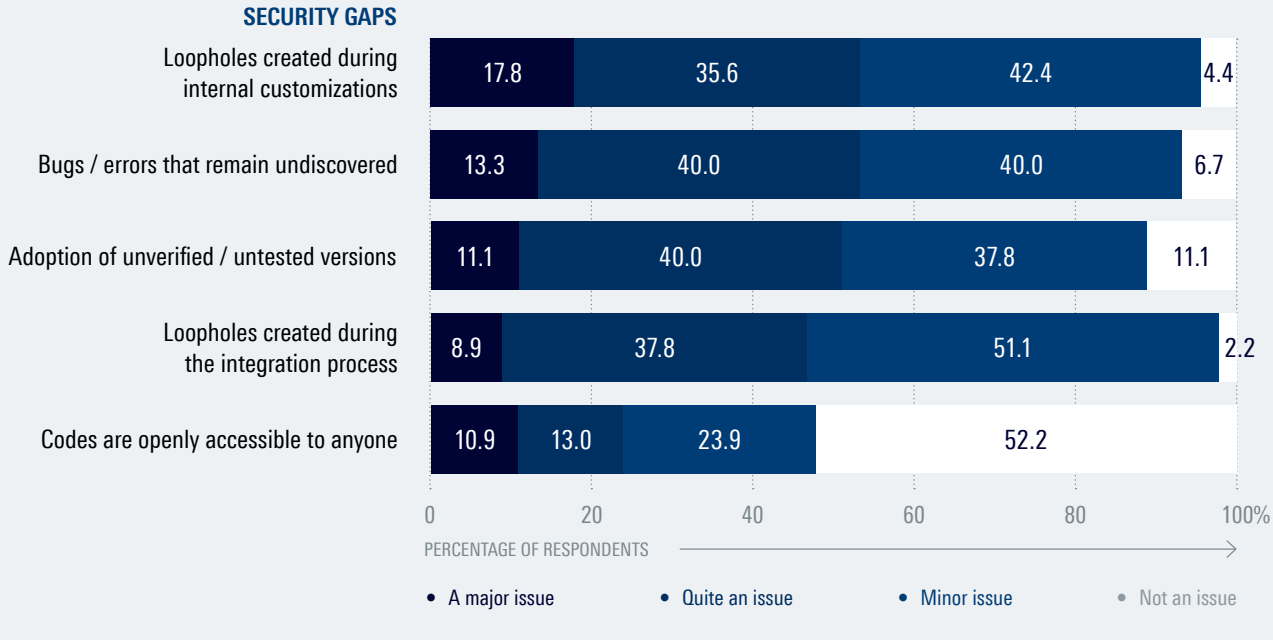

8. MIGRATION TO COMMERCIAL DPI

As organizations grow, data needs can become increasingly complex. Managing and delivering thousands of applications, files and databases from one end to another leads to burgeoning traffic loads that need to be monitored and secured at all times.

These rapidly growing visibility and informational requirements are further exacerbated by multi-cloud architectures, increasingly elusive threat vectors and the prevalence of ultra-low latency and bandwidth-intensive applications. As a result, the demand for robust and reliable traffic intelligence tools has become stronger than ever.

For networking, analytics and cybersecurity vendors, enhanced traffic intelligence not only improves their processing speeds and accuracy, it also enables them to introduce new features across their solutions. For example, vendors can offer analysis into encrypted threats for zero-trust network access or cloud access security broker. They can also provide real-time insights on the impact of popularly used video-based applications on network QoS for network monitoring tools. Additionally, vendors can introduce higher SLAs for their solutions and command a premium, leveraging improved levels of accuracy and speeds enabled by superior analytics.

These challenges and opportunities are driving vendors currently using open-source DPI to explore other alternatives to boost their traffic detection capacity and capabilities. In this scenario, the next logical step is to switch to commercial DPI. Commercial DPI tools provide organizations, specifically fast-growing enterprises, powerful and scalable traffic analytics solutions that are fully backed by specialist providers with deep expertise and experience. This allows vendors to seamlessly scale up their filtering capabilities, while driving optimality and cost effectiveness in the long-run.

Higher traffic volumes, need for a comprehensive library and increasing traffic complexities are pushing vendors to switch to commercial DPI

Migrating to commercial DPI, however, requires vendors to take into account evolving traffic conditions, financial implications, technical requirements and the level of support that is needed. The survey assessed a number of factors that are likely to play a major part in vendors' decision to switch to a commercial DPI solution.

The most common factor is growth in traffic, which was voted by 62.2% of vendors, followed by the need for a really comprehensive and frequently updated signature library, which was voted by 57.8% vendors.

Commercial DPI tools provide organizations, specifically fast-growing enterprises, powerful and scalable traffic analytics solutions that are fully backed by specialist providers with deep expertise and experience. This allows vendors to seamlessly scale up their filtering capabilities while driving optimality and cost effectiveness in the long-run.

DIAGRAM 16 Factors pushing vendors to switch to commercial DPI

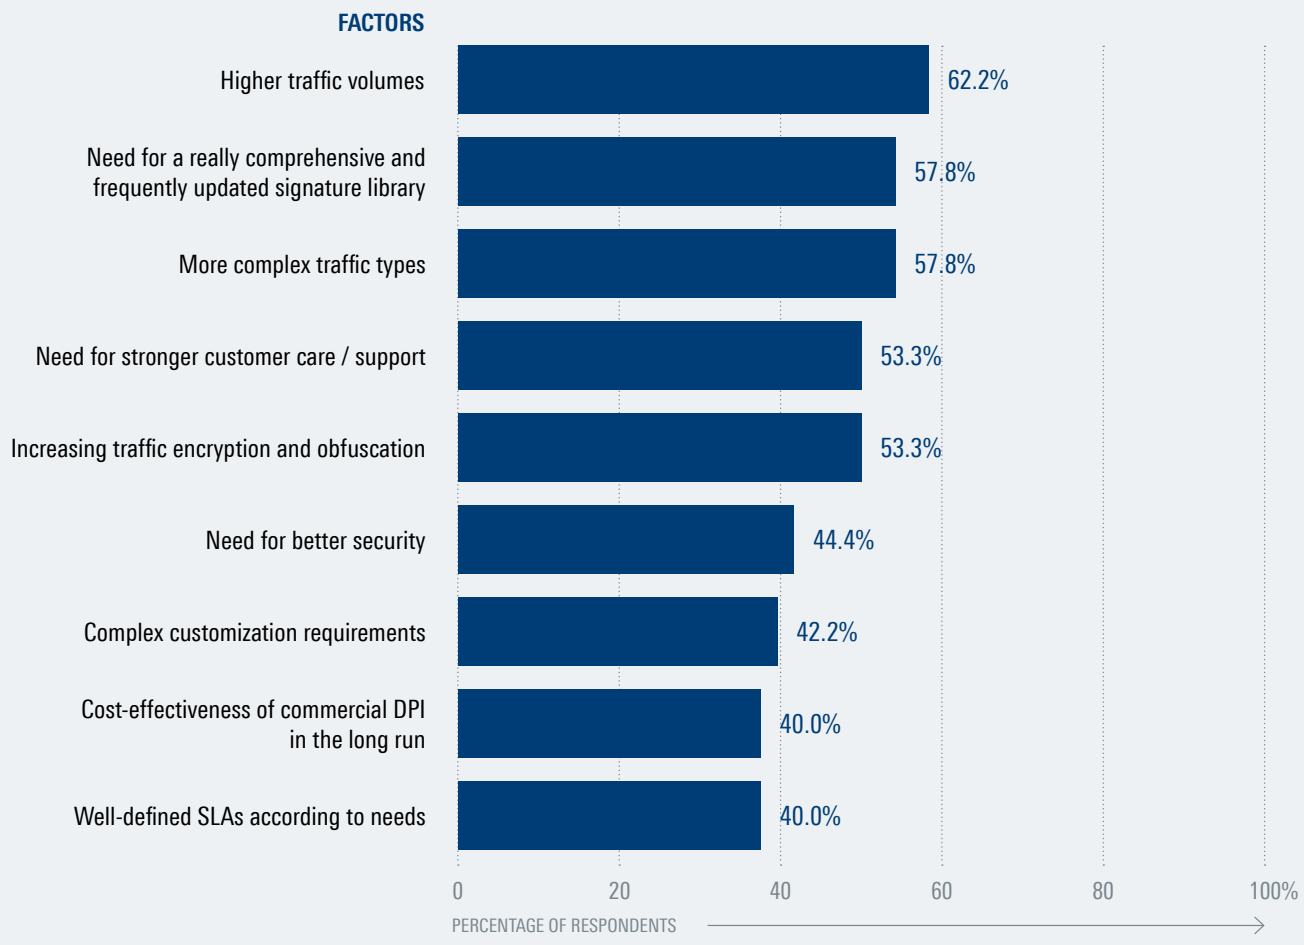

This is followed by the increase in complex traffic types, which is cited by 57.8% of vendors, and the need for stronger customer care and support, as stated by 53.3% of vendors. The next most influential factors are the increase in traffic encryption and obfuscation, as noted by 53.3% of vendors and the need for better security, as cited by 44.4% of vendors.

Complex customization requirements are cited by 42.2% of vendors, while cost-effectiveness of commercial DPI in the long run is mentioned by 40.0% of vendors. A similar share (40.0%) of vendors also agree that well-defined SLAs according to customer needs, offered by commercial DPI solutions, play a role in vendors' decision to switch.

Migration tools play an important role in encouraging the switch to commercial DPI

Replacing one DPI solution with another requires substantial planning, integration, reconfiguration, testing, operational management, communication and retraining. Depending on the host solution, for example a cloud access security broker or a routing device, and the level of analytics that is required, transitioning to commercial DPI may result in various costs and complexities. This is expected to hold back many decisions to scale up existing analytics systems, leading to legacy DPI solutions being retained permanently.

Bridging old and new

Custom migration tools can greatly facilitate existing open-source DPI users who intend to upgrade to commercial DPI solutions. Migration tools can be installed alongside the main DPI program. A migration tool synchronizes with existing open-source DPI software and translates current information structures and configurations to enable classification results from commercial DPI to be reproduced in the same format. This allows advanced analytics from a commercial DPI solution to be seamlessly integrated into the host solution and its components, without the help of third-party integration tools or extensive code edits.

The survey assessed how the availability of migration tools influences vendors' decision to move from open-source DPI to commercial DPI. Close to a fifth (17.8%) of vendors surveyed admit that migration tools will have a significant effect on their decisions, while a third (33.3%) think that it will have a moderate effect. Another 31.1% of vendors see migration tools having minimal effect on their decisions, while the remaining 17.8% feel that their decisions will not be affected by these tools.

A lack of information on the availability of migration tools, and their role in facilitating migration can have a major impact on how vendors perceive their use and relevance. The results of the survey indicate very little emphasis on such tools during migration decisions, despite them being available at vendors' fingertips. With increased awareness, vendors are likely to reassess their decisions to explore alternative solutions more seriously. This is expected to push the adoption of commercial DPI, especially among vendors with immediate capacity and performance upgrade requirements.

More than a third of vendors are already migrating to commercial DPI

More than half (59.6%) of vendors have no plans at present to migrate to commercial DPI, and only 4.3% of vendors intend to do so in the next three years, according to the survey. Upfront licensing costs and migration complexities can discourage vendors from exploring commercial DPI solutions. Continuous improvements in open-source DPI and strong community support can also be a contributing factor. Additionally, inadequate information on commercial DPI solutions and a lack of an urgent need to scale up their DPI capabilities can lead to limited interest among vendors to explore alternatives. Despite this, more than a third (36.2%) of vendors are already migrating to a commercial DPI solution, confirming its potential among users of open-source DPI.

DIAGRAM 17 Influence of migration tools on vendors' decision to move to commercial DPI

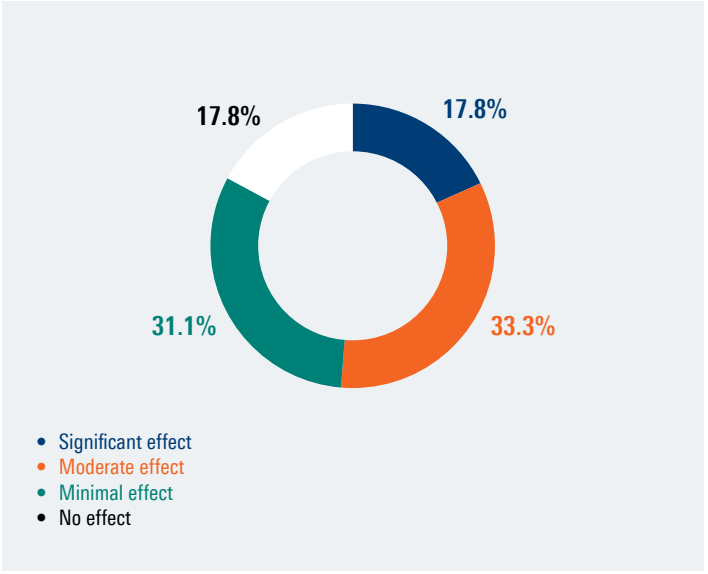

DIAGRAM 18 Vendors' plans to move to commercial DPI

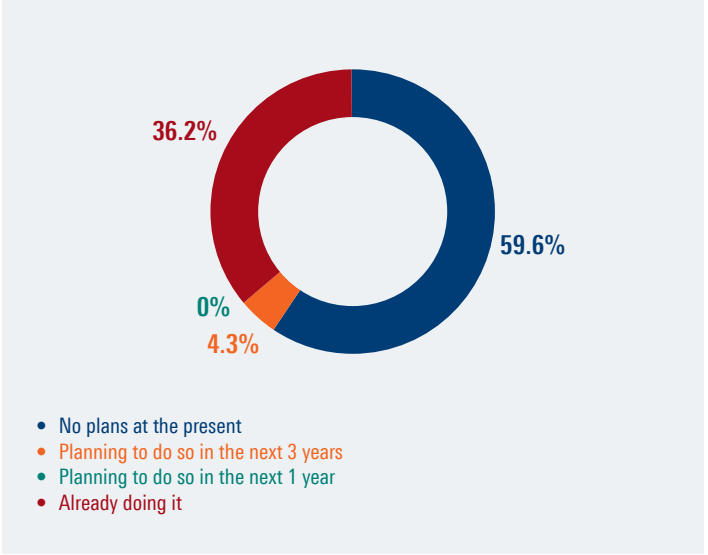

9. EXPLORING ADVANCED DPI CAPABILITIES

Introducing R&S®PACE 2 and R&S®vPACE

A leading player in the commercial DPI space is ipoque, a Rohde&Schwarz company, who has equipped leading networking, analytics and cybersecurity vendors, including startups, with powerful, real-time traffic intelligence. ipoque's next-gen OEM DPI engines, R&S®PACE 2 and R&S®vPACE, deliver advanced protocol and application classification capabilities based on DPI technology. These DPI engines can be easily integrated into any networking and cybersecurity solution to gain real-time IP network traffic visibility up to layer 7 and beyond, even when traffic is encrypted. They identify and classify thousands of applications and protocols, extract metadata in real-time and provide comprehensive application visibility to analyze, manage, and optimize IP network traffic. With weekly signature updates and continuous performance and reliability testing, ipoque's DPI technology offers the highest traffic detection rate in the market, enabling users to keep up with the rapid growth of IP traffic rates and concentrate on their core competencies.

Why ipoque

For open-source DPI users, who are seeking to upgrade their DPI capabilities, R&S®PACE 2 or R&S®vPACE provide a complete solution that combines accurate traffic insights, exceptional performance, enriched features, high-quality assurance and world-class support.

In terms of traffic insights, ipoque's DPI technology offers the following capabilities that ensure real-time, highly accurate analytics across any type of traffic flow:

- **Advanced traffic classification techniques** comprising statistical, behavioral and heuristic analyses for instantaneous detection of protocols (e.g. HTTP/S, BitTorrent, SIP/RTP and MTProto), applications (e.g. Instagram, Microsoft365, Discord and WhatsApp) and service types (e.g. audio, video and chat). ipoque delivers the highest classification accuracy rate, with virtually zero false positives.
- **Superior metadata extraction** that captures the widest range of traffic attributes at the packet-, flow- and session-level, including performance and capacity metrics such as bandwidth, speeds, latency, jitter and packet loss.
- **AI-based encrypted traffic intelligence (ETI)** to tackle even the toughest encryption protocols (e.g. TLS 1.3, ESNI, DoX and QUIC). ETI combines ML techniques (e.g. k-NN and decision tree models), DL techniques (e.g. CNN, LSTM and RNN), high dimensional data analysis and advanced caching, to deliver real-time insights into encrypted, obfuscated and anonymized traffic flows.
- **A comprehensive, weekly-updated signature library**, boasting thousands of signatures that can be installed during runtime not requiring a reboot.

In the face of growing traffic loads and increasingly complex network architectures, ipoque's engines ensure scalability, performance and efficiency. The engines feature:

- **High-speed processing** for unlimited volumes of traffic, ensuring no added latency from continuous inspection.
- **The lowest memory footprint** in the industry, allowing a lightweight implementation that complements lean architectures.
- **A software-form factor** which allows both engines to be incorporated into all relevant architectures - traditional, virtualized or cloud-native.
- Optimization for different computing frameworks. R&S®PACE 2, which is based on **scalar-packet processing**, caters for traditional environments, while R&S®vPACE, which is based on **vector-packet processing (VPP)**, is built for cloud computing environments e.g. those using DPDK Graph and FD.io, enabling support for 5G UPFs, VNFs and CNFs.

Taking into account the diverse requirements of different solution verticals, R&S®PACE 2 or R&S®vPACE come with a host of enhanced features including:

- **Custom signatures** where vendors can include their own signatures to align traffic analysis to specific use cases.
- **First packet classification** which allows traffic flows to be identified from the first packet itself. This reduces filtering time and overheads, supporting use cases that involve ultra low-latency applications and networks with extreme traffic loads.
- **Tethering detection** which allows networking and cybersecurity vendors to identify unauthorized use of network resources without requiring additional tools.
- **IPFIX/NetFlow** interface to integrate seamlessly with IPFIX/NetFlow consuming network solutions.

Recognizing the need to keep abreast of changes in traffic trends, and the importance of customer support, ipoque ensures:

- **Continuous R&D** and a **stringent quality assurance** process that includes 24/7, globally-conducted, automated testing to capture regional traffic patterns and an automated and widely distributed system across the globe to check for updates and changes for all supported mobile protocols and applications.

For open-source DPI users who are seeking to upgrade their DPI capabilities, R&S®PACE 2 or R&S®vPACE provide a complete solution that combines accurate traffic insights, exceptional performance, enriched features, high-quality assurance and world-class support.

- **Extensive experience and expertise** spanning nearly two decades. This includes vertical knowledge that is necessary for optimizing DPI to various different industries and reporting requirements.
- **Superior customer service and support**
  - Flexible and adaptable SLAs
  - Quick integration, enabling faster time-to-market
  - 24/7 support from DPI experts before, during and after deployment
  - On-site system performance optimization
  - Hands-on training and application engineering
  - Opportunity to share feedback and influence the product roadmap
  - A global support team and availability

Using migration assistant for faster transitioning

More importantly, ipoque offers a DPI migration assistant that is built specifically for open-source DPI users. By embedding the assistant in their open-source DPI solution, vendors are able to extract existing information structures that are used to relay DPI outputs to various components in the host solution. Without a migration tool, replacing an existing DPI solution will require extensive code changes to ensure the new classification results are aligned to existing reporting formats. Where there are multiple components and an extensive reporting structure, vendors will require third-party tools to ease integration and speed up migration. An example use case is a firewall, where DPI's inputs must be forwarded in real-time to the filtering engine, the threat analytics component and the local cache.

ipoque's DPI migration assistant eliminates transitioning complexities by automatically translating the information structures used by an open-source DPI tool, enabling inputs from DPI engine, R&S®PACE 2, to be reproduced into formats that are readily usable by each DPI-dependent component within the host solution. For example, the UI component of a secure web gateway can have immediate access to advanced analytics from R&S®PACE 2, without requiring a new round of integration.

With the migration tool in place, vendors can transition to ipoque's DPI technology anytime. This allows them to access latest application / threat awareness and high-performant processing without the hassle of replicating, reconfiguring and testing every single data stream. A seamless transition will also enable vendors to focus their efforts on meeting the needs of more complex and challenging networking use cases, thus enhancing their value proposition.

10. CONCLUSION

With traffic demands increasing year by year, networking, analytics and cybersecurity vendors will leverage traffic intelligence more than ever. Fine-grained analytics that provide real-time application and threat awareness will become necessary in supporting instantaneous policy responses and effective steering of traffic flows. This, and market competition, will inevitably push vendors to seek high-performant DPI technologies that are extremely reliable and efficient.

This report highlights the value proposition of open-source DPI as an affordable and accessible means to deliver traffic intelligence. The report also assesses some of its inherent limitations. The following summarizes its key findings:

- Vendor neutrality, low initial costs and support from a wider community are the biggest factors driving the uptake of open-source DPI.
- Open-source DPI offers comprehensive classification for protocols but has restricted visibility into applications, service types and threats. While it utilizes advanced statistical, behavioral and heuristic analysis, the use of AI is quite limited. Accuracy issues are also noted, with vendors reporting widely varying accuracy rates.
- Encryption protocols such as TLS extensions and DoX present a huge challenge for open-source DPI in terms of visibility. Open-source DPI also has limited visibility into obfuscated and anonymized traffic.
- In terms of technical capabilities such as memory efficiency, throughputs / speeds, dynamic process upgrade and support for VNFs and CNFs, there is an obvious gap between vendors' expectations and what is offered by open-source DPI. Likewise, the majority of vendors think that open-source DPI libraries lack comprehensiveness and are not updated frequently enough.
- Custom classification, first packet classification, IPFIX integration and full customization are some of the advanced features offered in open-source DPI. Other merits include ease of integration, API stability and long-term continuity. Open-source DPI however, lacks global presence and has poor support and customer service.
- Low upfront costs provide open-source DPI a competitive edge, however, this may be offset by many other long-term cost implications. These include losses from issue resolution delays and costs relating to customization, operations and upgrades.
- Open-source DPI can increase network susceptibility to security risks via loopholes created during internal customizations and hidden bugs and errors. Other concerns are the adoption of unverified or untested versions and loopholes created during the integration process.
- Growing traffic volume and complexities and the need for comprehensive libraries are some of the major influencing factors in vendors' decision to switch to commercial DPI; the availability of custom migration tools are expected to promote and accelerate these decisions.

These findings confirm the ability of open-source DPI in meeting baseline requirements for a network visibility tool. It also confirms the importance of commercial DPI solutions in providing vendors an effective pathway for scaling up their traffic filtering capabilities and capacity, and improving efficiency and cost-effectiveness.

In this regard, the availability of custom migration tools are expected to not only drive vendors to transition from open-source DPI to commercial DPI solutions, but also accelerate existing plans to do so. Ultimately, the right commercial DPI solution, implemented at the right time with the right migration tools can power networking and cybersecurity solutions with the most accurate and comprehensive insights into traffic, greatly enhancing their ability to monitor, manage and secure today's IP networks.

ipoque
ipoque, a Rohde&Schwarz company, is a global leader in the world of network analytics software. We leverage our deep domain expertise to create customized software solutions that empower our customers to transform data into intelligence. As a subsidiary of Rohde&Schwarz, we take advantage of potential synergies.

Rohde&Schwarz
The Rohde&Schwarz technology group develops, produces and markets innovative information and communications technology products for professional users. Rohde&Schwarz focuses on test and measurement, broadcast and media, cybersecurity, secure communications and monitoring and network testing, areas that address many different industry and government-sector market segments. Founded more than 80 years ago, the independent company has an extensive sales and service network in more than 70 countries.

The Fast Mode
The Fast Mode is a leading independent research and media brand, delivering breaking news, analysis and insights for the global IT/telecommunications sector. With a global reach spanning millions of readers annually, The Fast Mode partners with global technology companies to publish breakthrough ideas, critical analysis and latest updates on initiatives in the IT and telecoms space, focusing on IP/optical connectivity, network intelligence, security, cloud, internet of everything, CX and digital services.

Rohde&Schwarz GmbH&Co. KG
www.rohde-schwarz.com

ipoque GmbH
Augustusplatz 9 | 04109 Leipzig, Germany
Info: + 49 (0)341 59403 0
Email: info.ipoque@rohde-schwarz.com
www.ipoque.com

The Fast Mode
Info: +60 12 2016 186
Email: admin@thefastmode.com
www.thefastmode.com

Version 01.00 | May 2024
State of open-source DPI: Challenges, opportunities and alternatives