DPI FOR APPLICATION PERFORMANCE MONITORING

ROHDE&SCHWARZ

Make ideas real

TABLE OF CONTENT

1. Introduction 3

2. Framework for application monitoring 4
2.1 Application level monitoring 4
2.2 Network level monitoring 4
2.3 Real-time alerts 4
2.4 Root cause identification 4
2.5 Predictive maintenance and issue prevention 4

3. Identifying application performance issues 5
3.1 End user devices and external connectivity 6
3.2 Code and databases 6
3.3 Application runtime and middleware 6
3.4 Network equipment 7
3.5 Network links or connectivity 7

4. Enhancing APM with DPI 8
4.1 What is DPI? 8
4.2 How does DPI enhance APM? 9

5. Use case: DPI in APM illustrated 11

6. OEM DPI software 14

7. Conclusion 15

1. INTRODUCTION

In recent years, the shift from on-premises deployment to the cloud has given rise to more applications being delivered via cloud-native infrastructures. They often feature highly distributed architectures boasting the use of microservices and containerization for faster and highly scalable deployments. These infrastructures are also becoming highly virtualized, with proprietary hardware being replaced by virtual machines (VMs) to enable dynamic allocation of computing, networking and storage resources across different application infrastructures.

These complex delivery architectures create a host of new challenges for enterprises when it comes to managing the performance and security of their suite of applications. Issues that impair the performance or security of any application may be embedded within any layer of the complex application delivery architecture, and can impact any of the billions of transactions being made on that application. This is especially so with the rise in the amount of data captured and retrieved on business critical applications.

Such data includes banking transaction logs, medical record updates and telecom customer data, which is generated by the millions every second. Issues such as delays in transaction response times, application timeouts and malware attacks thus become much harder to pinpoint and resolve, subjecting enterprises to higher risk exposures and potential financial and reputational losses.

This paper examines closely how the use of deep packet inspection (DPI) addresses these challenges with deeper analytics and real-time insights on IP traffic in the context of application performance monitoring (APM). It looks at how these analytics address the current gaps in monitoring, tracking and analyzing application performance. It highlights DPI's application awareness as a pre-requisite for managing bandwidth-, latency- and security-sensitive applications, especially with the advent of 5G. The paper also illustrates how the granular analysis and accuracy of IP traffic inspection enables enterprises to identify issues as they happen and even before they happen, paving the way for a better management of enterprise networks and resources.

ENTERPRISE APPLICATIONS

Enterprise applications cover the entire spectrum of mobile, desktop and web applications accessed by internal (staff) and external (suppliers, agents, customers) users. These applications typically support the business operations of an enterprise and are usually connected with one another. Examples of enterprise applications are enterprise resource planning (ERP), customer relationship management (CRM), enterprise customer service and business intelligence (BI) suites. Business operations of an enterprise are highly dependent on uninterrupted access to these applications.

With more enterprises digitalizing their business operations, these web and mobile applications are also becoming key touch points for suppliers and customers. This sees enterprise applications becoming increasingly critical not just for day-to-day operations but also for managing the supply chain and delivering marketing and sales objectives while driving stakeholder engagement, satisfaction and loyalty.

2. FRAMEWORK FOR APPLICATION MONITORING

In addition to the complexities discussed in the earlier chapter, enterprises are seeing an increasingly critical need to develop application monitoring and network management capabilities that cover monitoring, real-time alerting, root cause identification and predictive maintenance. The following points summarize these key capabilities, forming the framework upon which enterprises can assess their current application monitoring and network management methodologies, tools and solutions.

2.1 Application level monitoring

Continuously monitoring the performance of each:

- Application, for example the up-time rates of an enterprise email system
- Application attribute, for example the average connection time for voice calls made on an enterprise communication suite
- Application transaction, for example the response time for a user trying to retrieve the account details of a bank customer during peak banking hours

2.2 Network level monitoring

Continuously monitoring the performance of enterprise IP communications networks responsible for connecting a distributed application architecture and bridging it to third party API providers and ISPs. These networks can be the local area network (LAN), campus, wide area network (WAN) and data center networks. Monitoring at this level involves the performance of:

- Each network node: network nodes include customer premises equipment, wireless access points, routers, servers and POS terminals. The uptime of a web server and the forwarding capacity of a router are examples of performance types tracked at network node level.
- Each connectivity link: IP connectivity links can be wireless (4G, 5G, satellite) or wired (optical fiber, copper-based Ethernet). Bandwidth, throughput, latency and jitter are performance types tracked across individual links and network areas.

2.3 Real-time alerts

When enterprise applications register downtimes or delays in response times, experience sluggishness and transaction failures, run into errors, crash or perform intermittently, the enterprise needs to be alerted of these issues in real time. Service level agreements (SLAs), quality of experience (QoE) metrics and performance and security requirements must be built into the monitoring mechanism so that application performance can be measured against these metrics and any shortfalls can be reported immediately to the IT admins via built-in alerts.

2.4 Root cause identification

Upon learning about issues in real time, the monitoring mechanism must also be able to provide end-to-end granular visibility into the network, so that the underlying cause can be accurately ascertained as soon as the network is affected. Regardless of whether it is a particular link, a network node or an end user device, the problem area and underlying cause must be immediately visible to the IT admins down to the last machine, code line, service, link or device.

2.5 Predictive maintenance and issue prevention

Advanced monitoring capabilities will also enable enterprises to identify hotspots and potential issues before they set off pre-defined triggers or alert thresholds. By using machine learning (ML), artificial intelligence (AI) and correlation of past data, parameters such as current traffic information, application response times or anomalies in server CPU consumption can be evaluated to predict potential issues before they happen. This enables predictive maintenance and preventive measures to be executed in a timely manner, avoiding potential financial and reputational loss.

3. IDENTIFYING APPLICATION PERFORMANCE ISSUES

The crux of managing an application's performance lies in an enterprise's ability to understand their application architectures. This is especially true for distributed application architectures, where performance and security issues can reside anywhere from the end user device to the bare metal used in any of the network nodes. The following graphic illustrates the architecture of a cloud-based application covering end user points, network devices, links, servers and clouds.

DISTRIBUTED APPLICATION ARCHITECTURE

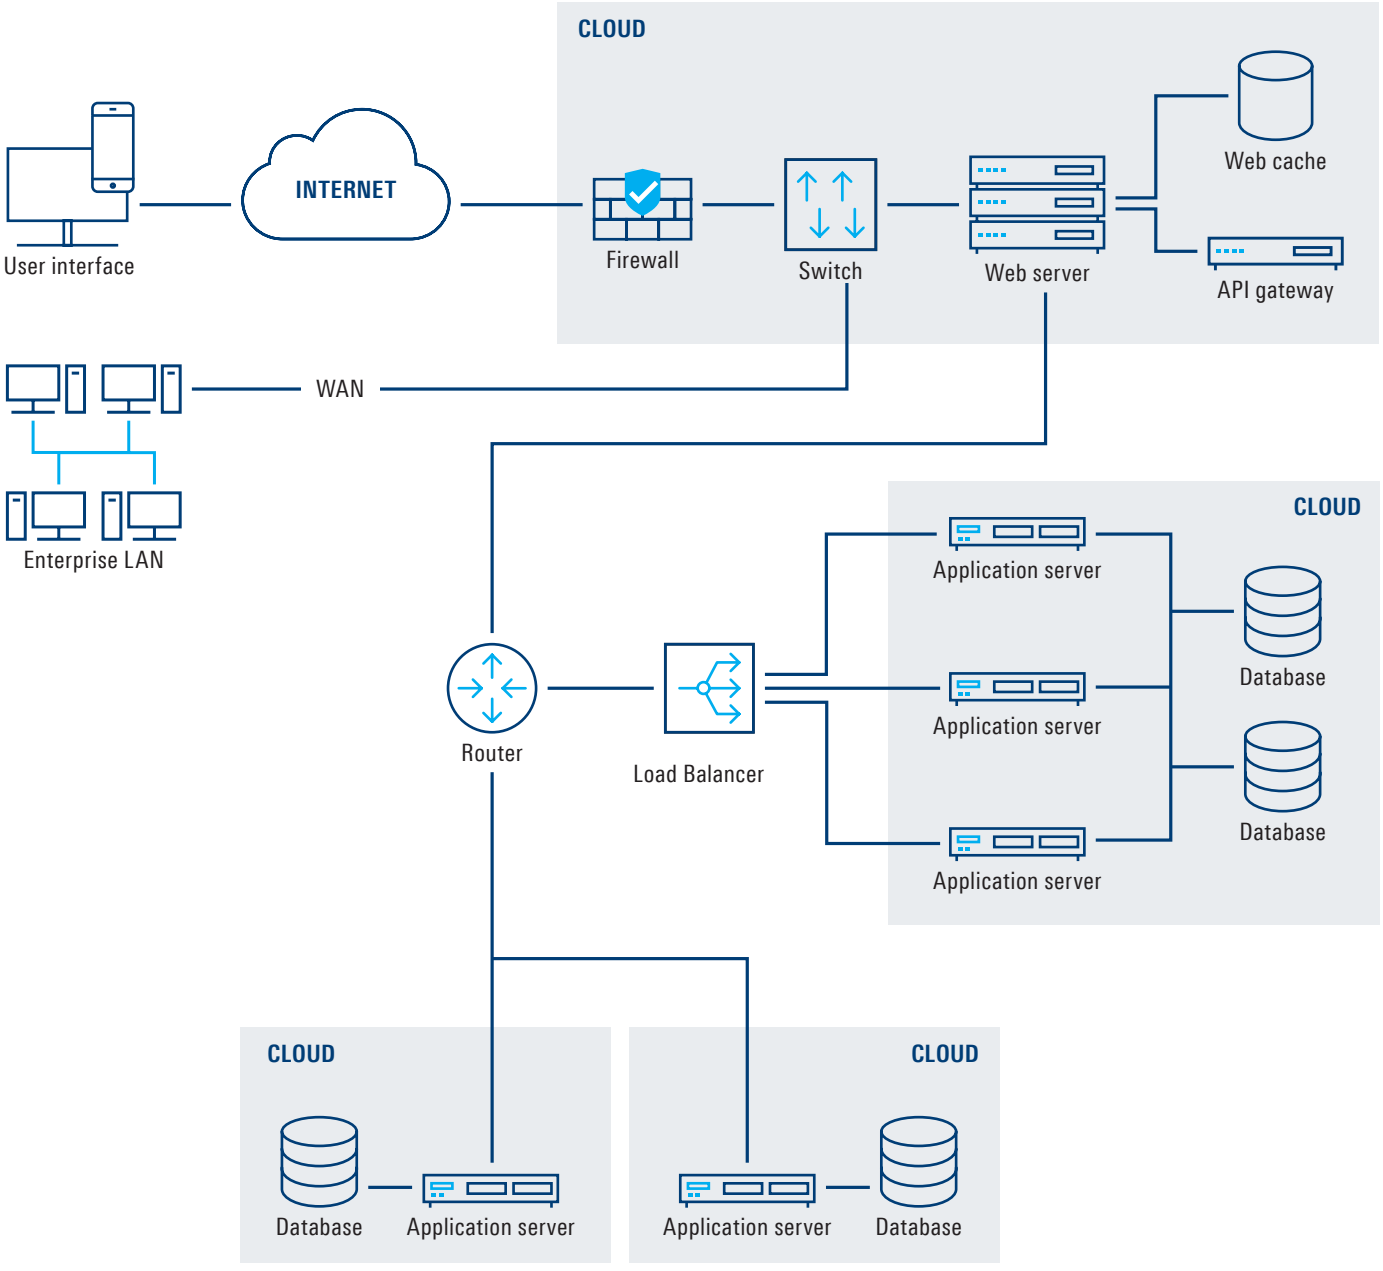

An application architecture is built in such a way that each transaction, such as a page request, runs on its own pathway. The pathway can be split into various physical and software layers starting from the point of request and ending at the server machines, before looping back to the requesting device. The diagram below lists these layers, grouped into two categories - the external domain and enterprise domain. The external domain covers end user devices and the wired/wireless links that are not part of an enterprise network while the enterprise network comprises of devices and nodes managed or controlled by the enterprise.

Each of these layers are defined by different hardware and software, with different risks and vulnerabilities that can impact an application's performance, security and QoE. The following sections discuss in detail the type of issues that enterprises can anticipate at each layer and how they impact application performance.

3.1 End user devices and external connectivity

End user devices impact an application's performance due to resource limitations. A video application requires sufficient memory and processing capability on the receiving devices. At the same time, if the device is infected with viruses and fails to respond as expected, user experience on an application will be severely impacted. Apart from these, applications accessed on networks managed and controlled by third parties such as ISPs, mobile network operators and WiFi hotspot providers can experience performance degradation. This occurs when network management is deficient, resulting in network congestion, poor signal strength, limited network coverage or insufficient data allocation on the user's data plan.

3.2 Code and databases

Errors in code and database queries will cause significant delays in application response times and can stall an application or return recurrent errors. Missing data from application databases results in the same problems, as does congestion originated during memory-intensive processes such as database searches. Similarly, configuration errors can give rise to various performance issues in an application.

3.3 Application runtime and middleware

Issues at this level range from incompatibility between programming language versions to gateway timeouts that arise when there are communication delays between a web server and a runtime (for distributed architectures). They can also be due to heavy traffic or distributed denial of service (DDoS) attacks. Memory leaks in runtime and middleware applications can cause the applications to stop running. Applications can also slow down significantly when a connection pool in an application server receives connection requests greater than the specified limit.

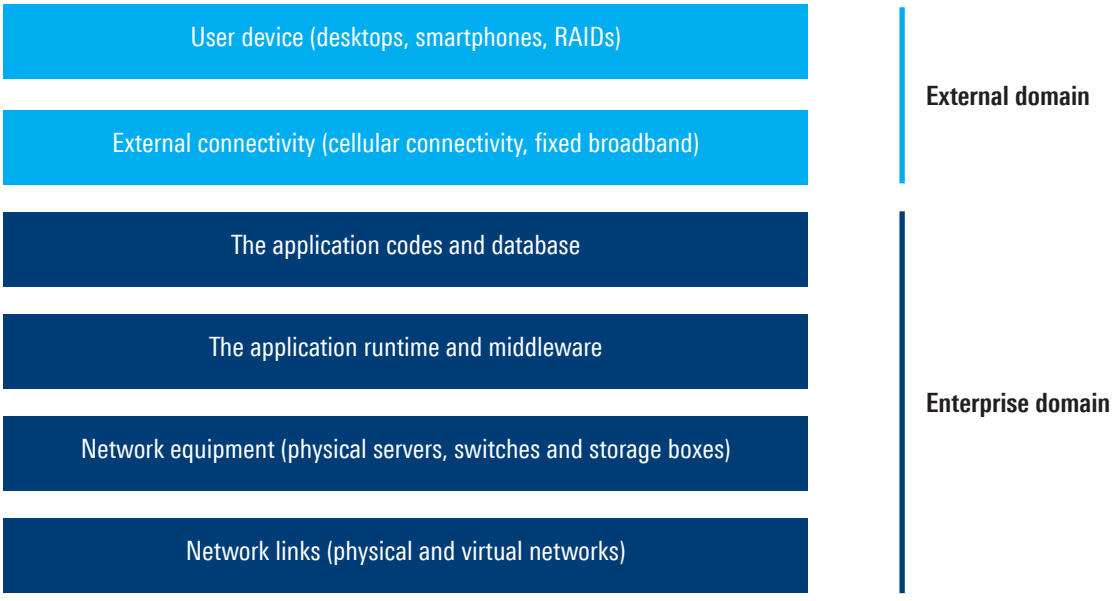

Multi-layer application delivery architecture

3.4 Network equipment

Power issues, insufficient processing capacity due to CPU load, lack of memory, shortage of storage space and issues with server-network interfaces will all lead to deteriorated application performance. In virtualized environments, processing, memory and storage bottlenecks can reduce the capacity of a VM and impair its ability to carry out certain tasks. Issues at this level often point to insufficient resources, poor resource distribution, unexpected usage peaks or resource exhaustion due to cyberattacks such as DDoS attacks on the application layer.

3.5 Network links or connectivity

Issues with network links or connectivity are split into three categories: network performance, network security and suboptimal traffic management policies.

Network performance issues
Network level issues often relate to bandwidth availability and the ensuing speed. However, bandwidth provisioning alone does not guarantee the performance of an application. Networks can suffer from high latencies caused by issues in connection links and propagation delays. Packet loss happens when a link is congested, network devices are overutilized or the hardware is faulty. Another potential issue is network jitter, which can be caused by outdated hardware, router misconfiguration, queue contention and serialization.

With the arrival of 5G and the rise in ultra-reliable low latency (URLLC) applications, such as gaming and remote surgery applications, minute degradation in network latency will greatly impair application performance and the end user experience. The same applies to ultra-high definition (UHD) content consumption over the internet and mobile video applications. According to Fast Company[1], Amazon's calculations showed that a 1 second slowdown in page load time costs the company USD 1.6 billion in annual sales, while Google's calculations showed that it risks losing 8 million searches a day with 0.4 seconds slower search results.

Network security issues
Another key connectivity issue relates to security. Malware, including viruses, trojans, worms, adware and spyware, does not just infect network devices. By using compromised devices to launch intensive botnet attacks into the network or by sending large amounts of stolen data from computers on the network to the attacker, it also consumes network bandwidth. Similarly, DDoS attacks take up the network bandwidth by sending hundreds of concurrent requests every second, thus exhausting resources on network devices such as firewalls and web servers, causing an application to stall. Other cyberattacks that impair the network include advertisement of false routes which drive traffic to compromised hosts, causing higher latencies and longer response times for an application.

Suboptimal traffic management policies and enforcement
Enterprise applications performance is also highly susceptible to network management policies. Lack of differentiated traffic management policies based on different traffic types will result in all traffic being routed through the same routes and network services. A business critical application or a URLLC application that requires edge processing will be routed all the way to the core network and may be filtered through numerous firewalls, resulting in low response times and high latencies. This results in major application performance issues, as well as overconsumption of network resources. At the same time, low risk traffic undergoes a higher degree of security filtering typically assigned to suspicious activity, while a video application that requires caching in the web server may be delivered newly from its source every time it is requested by the user. Poor traffic management policies and weak enforcement impact application performance as much as any other factor.

THE COST OF DOWNTIME

According to an IDC Study, the annual average cost of unplanned application downtime for the Fortune 1000 is between $1.25 and $2.5 billion. The average hourly cost of an infrastructure failure is $100 000 per hour. To gain the required visibility to effectively identify and remediate application issues quickly and accurately, application development teams need the appropriate tools and dashboards.

1) Source: Fast Company
https://www.fastcompany.com/1825005/how-one-second-could-cost-amazon-16-billion-sales

4. ENHANCING APM WITH DPI

When it comes to managing application performance, most enterprises today use APM. APM is a unified monitoring technique that uses analytics to deliver visibility into the end-to-end performance of an enterprise application measured against pre-determined performance, security and QoE targets.

With software agents deployed across the distributed application architecture (for example, across web and application servers), APM monitors the performance of each network node by tracking application transaction flows (for example, by instrumenting bytecode during application runtime). Using either real-user monitoring or digital agents, it collates application performance metrics and matches them with these results for identification of issues in the application layers. This enables APM to conduct root cause analysis, a process that is now supported by ML and AI for higher diagnostic accuracies. This information is then relayed to enterprise IT teams for further investigation and resolution. In addition to tracking transaction flows, capturing performance analytics, performing diagnosis on issues and identifying problem areas, APM tools collect continuous data of an application (for example, application usage patterns by location, device types, peak usage times and resource consumption) for correlation analysis and reporting purposes.

4.1 What is DPI?

DPI is a traffic recognition method using behavioral, heuristic and statistical analysis to classify IP traffic and extract metadata in real time. With granular information on protocols, applications and application attributes, it provides information on the following levels:

- Network level traffic – Speed, latency, jitter, bandwidth consumption, type of devices, user locations, type of applications, overall state of security and type of threats
- Application level traffic – application-specific performance metrics such as average page download times and average transaction completion times, speed, latency, bandwidth consumption, user locations and security threats
- Application attribute level traffic – attribute-specific performance metrics such as average content download times, average search results display times, average call connection times and detection of security threats

DPI engines leverage continuously updated traffic libraries established through analysis of current traffic patterns, behavior and trends. Traffic identification by DPI also enables checking anomalies and suspicious traffic patterns against updated libraries of the latest cyberthreats, such as DDoS attacks, malware and ransomware. By doing this, it can identity and block threats in the network and application layers in real time.

With techniques such as behavioral, heuristic and statistical analysis, DPI is able to extend packet inspection to both encrypted and non-encrypted traffic. Protocols identified by DPI can be classified into the following categories:

Metadata category	Example metadata
Traffic volume	Per user, per protocol, per application, per flow, per direction
Service detection	Differentiation between for example Skype audio and video calls
Quality of service (QoS)	Jitter, throughput, latency, roundtrip time, ramp up time, packet loss, retransmissions
Security and data leakage	File up- and downloads, entropy-based DNS tunneling detection
Client information	HTTP/QUIC user agents, operating system

DPI is typically deployed as a network service or a functionality, for example within a network service such as a firewall. DPI's ability to identify, classify and manage IP traffic in real time thus delivers both full network visibility and real-time application awareness.

4.2 How does DPI enhance APM?

While APM is able to deliver all of the discussed metadata categories at a high level, enterprises are demanding new capabilities that require packet-level intelligence. DPI provides this intelligence, which enables APM to go beyond the usual simple network management protocol (SNMP) and flow-based traffic monitoring to detect performance issues. This is necessary for a converged, end-to-end view of enterprise application performance, which requires real-time analytics on, among others, enterprise network links and connectivity such as LAN and WAN, as well as on user devices and external connectivity. It is also necessary for a deeper dive into application layers, protocols and attributes. Without packet-level analytics, there will be delays in diagnosing application performance issues and a significant proportion of them will remain undetected and unresolved.

Towards this end, DPI does not just fill in these gaps, but also greatly enhances the functionality of both APM and network performance management (NPM). The following five points discuss in detail the mechanisms by which DPI delivers these capabilities.

Uncovering user device and external connectivity issues

Enterprises can configure their devices in such a way that the device health data is continuously logged into the LAN/WAN network reports. This can be executed via APM software agents installed into user devices, including mobile devices.

With DPI, this information is enhanced in many ways. On the performance front, it is possible to identify the specific mobile operator, ISP or WiFi hotspot provider behind mobile devices accessing an enterprise application whose last mile connectivity is poor or whose data allocation is insufficient. DPI achieves this by matching packet metadata with the application metadata and using the IP addresses on those packets to determine the originator networks where packet speeds are slowing down. Matched against user login session, enterprise IT teams can identify users who are experiencing slow performance in real time, and matched against user ratings, can confirm the external connectivity issues as the underlying cause for poor QoE.

At the same time, DPI enables APM to pinpoint the exact users whose devices are compromised and which might be the source of malicious attacks experienced on the network. DPI uses application metadata to detect traffic anomalies and then matches these anomalous packets with user login sessions on a given application.

Adding depth to issue diagnosis at the application layer

DPI functionality at network nodes enables inspection of each IP packet flowing in and out of each node. Speeds, latencies and security threats are identified in real time at the entry and exit point of these nodes, which helps explain the performance metrics captured by APM. For example, detection of malware flowing into the web server can be used to explain the surge in CPU usage at that node.

In another example, at various network links in an application architecture, DPI detects the speed of traffic belonging to an application attribute, e.g. the messaging attribute on an enterprise communication application. This enables the enterprise to identify the underlying cause of slow message delivery on this application, which could be congestion at the messaging server.

Diagnosing enterprise connectivity issues

In distributed architectures, connectivity issues and routing inefficiencies between various network nodes can result in massive delays in application response times.

In this respect, DPI reveals, in real time, latency, packet loss and jitter on each link, each LAN and the overall enterprise network, organized by each transaction, protocol, application attribute and application. This enables the enterprise to identify issues such as intermittent or slow WAN connectivity and link this to the performance of an application or transaction.

CORRELATION ANALYSIS AND APM

Performance metrics provided by APM are analyzed in relation to each other. Response time delay, bandwidth and throughput for example are closely related. Response time delay is influenced by bandwidth, while response time delay and bandwidth collectively influence throughput. Issues such as slow propagation, poor routing/switching or delays caused by the medium of transportation impact response time delay, bandwidth and throughput at varying degrees. Correlation analysis is critical to link various metricsfor the purposes of accurately diagnosing a network event and correctly pinpointing the underlying causes.

Providing analytics for application development and testing

Developing enterprise applications on distributed architectures requires detailed insight on the application traffic flow by transaction and application attributes. DPI-embedded APM tools provide application developers, testers, system administrators and DevOps teams with application- and network-level IP traffic data. With said data, they can design application infrastructures and transaction routings and decide on the deployment of network services such as firewalls, content delivery networks (CDNs) and APIs. For testers, load testing and troubleshooting at code and application infrastructure level become more effective with real-time IP traffic information provided by DPI. This enables enterprises to optimize network resources and arrange for lower latencies, higher speeds and enhanced security for their applications while reducing the total cost of ownership (TCO).

Complementing network performance management

In addition to providing a converged view of an application's performance across networks, enterprise IT infrastructure and software stacks, DPI's detection capabilities go a step further to support enterprise NPM. NPM is essentially a combination of APM (end-to-end application performance monitoring and timely analytics) and the following:

- Timely resolution of network issues and institution of automated remedial actions
- Optimization of the network through improvements in network architecture, network services, network devices and network resource allocation
- Enhancement of the network using virtualization, containerization, cloud-based deployments and other technologies

DPI ADDS THE FOLLOWING ADVANCEMENTS TO NPM:

Enforcement of traffic management rules in real-time

With advanced forensics provided by DPI on IP traffic, network devices such as routers/switches are able to enforce traffic management rules built into the network in real time. These network devices are able to execute for example, routing of traffic through firewalls, load balancers and CDNs based on the traffic classification information of each IP packet provided by DPI.

This avoids network level performance issues such as bottlenecks and congestion. It also prevents overconsumption of network resources, for example in cases where all IP traffic, regardless of the risk classification, is routed through uniform security filtering via multiple firewalls and proxy servers, resulting in unnecessary processing of otherwise low-risk traffic.

Development of effective traffic management policies

DPI enables the development of effective traffic management policies. Network management is only as good as the network policies that govern it. Granular analysis provided by DPI with accurate packet-level classification combined with big data analysis, ML and AI enables network administrators to determine at protocol, application attribute and application level, rules that best deliver the expected performance and QoE while addressing possible security threats.

For example, based on past data on application latency and bandwidth consumption, enterprises may institute a policy rule that routes all traffic from voice calls on an enterprise IP-voice network through high priority routes that offer higher speeds and lower latencies instead of the default routes.

Development and optimization of enterprise networks

Additionally, the DPI functionality in combination with the monitoring information provided by APM enables NPM tools to assist Enterprises in the development and optimization of their enterprise networks. DPI provides not just the performance and security information for the entire application stack, it also provides detailed breakdown on resource consumption throughout the network.

This data helps enterprises to improve and optimise the design and architecture of their networks, for example their enterprise WAN, by provisioning additional bandwidth, implementing distributed architectures for selected applications and deploying virtualized network functions (VNFs) where scalability is required.

5. USE CASE: DPI IN APM ILLUSTRATED

The following is a use case that illustrates how granular, real-time IP traffic analytics provided by DPI create greater visibility into application delivery and performance, and how that helps in diagnosing and resolving performance issues.

Monitoring the performance of an enterprise supply chain management software

Enterprise X supplies fresh food produce to hoteliers, restaurants and large fast food chains. The produce is sourced via regionally distributed farms in 12 countries. Internet of things (IoT) sensors are deployed at farms, across warehouses and on freight for end-to-end tracking. Information from the sensors is relayed in real time to the enterprise supply chain management software (SCMS) for tight control on harvesting cycles, as the produce is highly perishable. The SCMS provides order and inventory information, as well as added features such as live messaging and an issue logging system. Enterprise X maintains an SD-WAN that connects its regional hub with country branches and hosts its various application servers and databases in the regional hub.

USERS OF ENTERPRISE SUPPLY CHAIN MANAGEMENT SOFTWARE (SCMS)

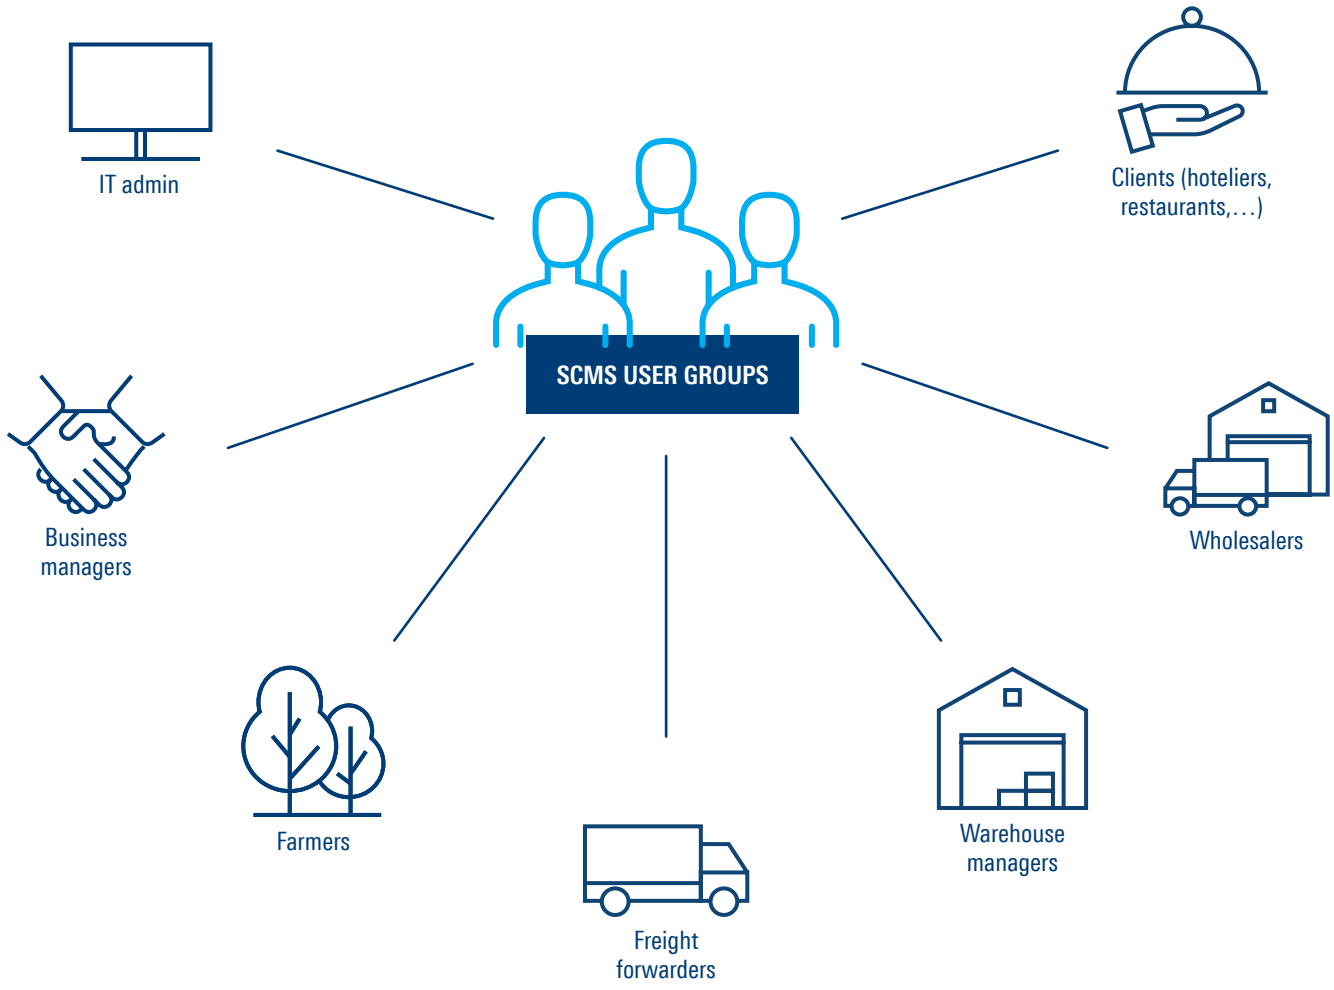

SAMPLES OF MEASURES AND METRICS THAT TRACK QOE, PERFORMANCE AND SECURITY OF THE SCMS:

Categories	Sample measures	Metrics
QoE measures	User reviews	End user ratings on application
	Service abandonment	Cart abandonment rate
Security measures	**Overall application**	
	Malicious activity	Volume of unusual usage pattern incidents
	Unauthorized access 1	Volume of failed login attempts
	Unauthorized access 2	Volume of unusual login patterns
	Processes	
	Malicious activity	Volume of unusual usage pattern incidents by: 1. Page requests 2. Request errors
Performance measures	**Overall application**	
	Application speed 1	Average page response time
	Application speed 2	Average page load time
	Application speed 3	Average transaction processing time
	Application load	Total transaction volume
	Transaction completion	Transaction status: successful, failed or cancelled
	Application attribute	
	Messaging speed	Message delivery time
	Transactions	
	Transaction load	Total transaction volume for [retrieval of inventory summaries, retrieval of daily order summaries, stock updates, order creation, order tracking, order payments]
	Transaction completion	Transaction status - successful, failed or cancelled for [retrieval of inventory summaries, retrieval of daily order summaries, stock updates, order creation, order tracking, order payments]
	Transaction speed	Average transaction processing time for [retrieval of inventory summaries, retrieval of daily order summaries, stock updates, order creation, order tracking, order payments]

APM tools used by enterprise X highlighted the following red flags for the month of May. The following are the figures extracted from the monthly dashboard:

- High total transaction volume from May 22 – May 24, 201X
- High average transaction processing times for retrieval of daily order summaries in the month of May 201X
- Higher average message delivery time for May 27, 201X

Enterprise X has deployed DPI within its APM tool. Analytics from DPI are merged with the data collected by APM to provide insight that helps enterprise X diagnose the red flags accurately in real time and restore the performance of SCMS. The following illustrate the traffic identification and issue resolution by DPI for each of the red flags above.

RED FLAG 1

High total transaction volume from 22 May – 24 May 201X

DPI's advanced packet filtering had identified traffic anomalies during the period of May 22 – May 24. Matched against DPI's advanced libraries, malware had been identified and network firewalls had been notified to block the identified packets in real time. This investigation confirmed that the high total transaction volume captured on the APM dashboard was due to the identified malware activity. Further analysis based on the IP addresses and the application login sessions revealed the branch office where compromised computers were being used on the enterprise WAN. Computers were cleaned and network security was restored.

RED FLAG 2

High average transaction processing times for retrieval of daily order summaries in the month of May 201X

DPI detected insufficient bandwidth usage by the SCMS application and discovered that bandwidth allocation for SCMS in the SD-WAN network became too low due to priority traffic routing implemented for another new enterprise application launched in May. The database server holding order summary data was the most impacted due to having the smallest bandwidth link. Bandwidth upgrade is implemented immediately. Alerts are then built into the enterprise's APM tool to trigger if future delays in transaction processing times across all transactions occur.

RED FLAG 3

Higher average message delivery time for May 27, 201X

DPI detected no latency degradation for the messaging traffic. DPI however identified a huge volume of video traffic on the messaging service. Feedback from sales representatives showed that there was a UHD video file that was shared heavily by a large number of clients seeking to place new orders. The video explains a new product introduced by enterprise X. There was also a surge in users logged into the messaging session for chats with enterprise X sales representatives. No action was deemed necessary.

6. OEM DPI SOFTWARE

The use of DPI spans everything from application monitoring to network management. APM and NPM vendors, covering both equipment and software, require DPI as an essential capability. This capability is today available as a service from Rohde&Schwarz.

Rohde&Schwarz is a global leader in the world of IP network analytics software. Rohde&Schwarz leverages deep domain expertise to create customized software solutions that empower the communications industry to transform network data into intelligence. They enhance enterprise and service provider network solutions with their market-leading DPI software R&S®PACE 2 and the technically advanced IP probe R&S®Net Sensor OEM.

R&S®PACE 2

R&S®PACE 2 is a protocol and application classification engine that combines DPI, behavioral, heuristic and statistical analysis to detect network protocols and applications and extract metadata in real time. It classifies thousands of applications and protocols and provides content and metadata extraction regardless of whether the protocols use advanced obfuscation, port hopping techniques or encryption.

Continuous evolution of web protocols and applications requires software libraries that are frequently updated to include the latest changes in IP traffic trends. R&S®PACE 2 is essentially a ready-to-use DPI software library to which new application signatures are added weekly.

In addition to this, R&S®PACE 2 features a dynamic upgrade of its libraries, allowing near real-time incorporation of the latest protocols and application signatures to ensure that there are no gaps between the central signature registry and the software libraries currently used for traffic filtering on the client's end. These upgrades, combined with continuous performance and reliability testing significantly increase traffic detection rates as well as traffic classification accuracy.

At the same time, R&S®PACE 2 is delivered with a software development kit (SDK), giving APM and NPM vendors the flexibility of deploying powerful analytics as a functionality or 'on-demand' as a resource via APIs or other interfaces across their application monitoring software.

R&S®PACE 2 also boasts the most efficient memory and CPU utilization in the industry, featuring the smallest processing footprint. The software requires only around 400 bytes per flow and no memory allocation during runtime.

By providing DPI as a service, R&S®PACE 2 enables network software and equipment vendors to reduce costs and risks associated with developing and maintaining this highly complex technology.

R&S®Net Sensor OEM

For companies looking for a ready-to-use DPI network probe, R&S®Net Sensor OEM is an ideal choice. R&S®Net Sensor OEM is a flexible passive probing software based on the R&S®PACE 2 DPI engine. It provides fast packet processing to reveal a clear picture of the entire network and its subscribers. The IP probe classifies both plain and encrypted IP traffic to offer detailed visibility. Additionally, it can correlate control and user plane data to keep track of specific subscriber sessions and experiences in a network.

R&S®Net Sensor OEM also provides aggregated information of the collected data. This includes information on the type of applications and protocols by user, time, duration, frequency and usage. Depending on individual demands, additional modules can further enhance the solution: aggregation and correlation functions, a database and graphical user interface or integration of third-party products and solutions offer additional network insight and flexibility. APIs enable customers to extend R&S®Net Sensor OEM with their own modules.

The following are key benefits of an OEM DPI deployment:

- Customization based on individual requirement
- Focus on core competencies for higher efficiency and profitability
- Faster time-to-market thanks to optimization of the development schedule
- Reduction and optimization of development costs
- Maximization of return on investment

7. CONCLUSION

In summary, as the number of cloud-based applications managed by enterprises continues to rise, enterprises must have enhanced capabilities in place for application- and network-level monitoring. These have to be accompanied by real-time alerts and timely identification of the root causes, and complemented by predictive maintenance and preventive mechanisms. To achieve this, APM tools have to go beyond SNMP and flow-based traffic monitoring to packet-level identification and classification, in order to deliver a converged end-to-end visibility on application performance. This enables enterprises to extend application monitoring from simple performance tracking to a whole array of advanced diagnoses and analytics. This includes uncovering user device and external connectivity issues and adding depth to existing diagnosis of issues at the application and enterprise connectivity layers, as well as providing analytics for application development and testing and, more importantly, complementing NPM in the:

- Enforcement of real-time traffic management rules
- Development of effective traffic management policies
- Development and optimization of enterprise networks

With the granular real-time analysis it provides, DPI enables APM and NPM tools to correlate both application and network metrics to provide a detailed view of the end-to-end performance of their applications and networks. Without this capability, it will be impossible for enterprises to manage application delivery on a multi-cloud, distributed architecture. This and the shift towards highly virtualized architectures makes fine-grained, packet-level intelligence delivered in real time a critical 'must have' functionality. DPI delivers this capability seamlessly, providing APM and NPM vendors, as well as other network analytics providers, with both real-time insights and aggregated values that enable them to deliver the highest QoE, performance and security on enterprise applications and networks.

ipoque

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

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.

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

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

PD 3608.3078.52 | Version 01.01 | December 2019
White paper | DPI for Application Performance Monitoring
Data without tolerance limits is not binding | Subject to change

3608307852