How Context-Aware Data Makes Security Threat Detection Better

WHITE PAPER

INTRODUCTION

Detecting security threats is tough, and network security analysts' jobs are getting tougher as threats and attacks become increasingly more sophisticated. To fight back, security architects integrate more security and monitoring tools into the enterprise. Chief information security officers (CISOs) not only want more tools, but they also want tools that work cooperatively, even across vendors. Discovery, forensics, and remediation all require correlation analysis among different tools that perform different functions. Correlation analysis becomes easier when network tools all get reliable access to relevant traffic at the same time. Enhance that traffic with context, and it can quickly make security analysts and the tools they use better.

Enhance monitored network traffic with context and it can make security analysts and the tools they use better.

WHAT IS CONTEXT-AWARE DATA PROCESSING?

Context-awareness is the ability to extract knowledge from or apply knowledge to information. That definition may sound arbitrary and oversimplified, but at its core, this is the foundation of what context-awareness is. The extent of knowledge is what separates the best context-awareness capabilities.

CONTEXT-AWARENESS AT THE FOUR SEASONS

The following is a real-life comparison of this digital concept. The Four Seasons hotel chain is well-known for its legendary customer service, which includes staff addressing each guest by name. How do they do this? They use context-awareness. After a guest makes a reservation, hotel staff may Google him or

her, print his or her photo, and place it in staff-only areas. The driver who picks up the guest at the airport may get a name from luggage tags. Or, he or she may use flight data to estimate arrival time. And sometimes, the doorman may just ask, "Under what name is the room?" Once they know a guest by name, the information spreads throughout the hotel like wildfire, and employees are expected to address the guest formally.

CONTEXT-AWARE DATA PROCESSING ON GOOGLE MAPS

In the digital world, context-aware data processing takes data, applies intelligence, and produces greater insights. Open Google maps to get directions, and it picks up your location and indicates it with a blue dot on a map. It is amazingly reliable and accurate. Google uses global positioning system (GPS) coordinates, Wi-Fi location services, and your web browser's location information. Google then applies location knowledge to information it receives from users accessing the application. The combination of information, along with correlated knowledge, is what makes its context-aware data processing so powerful.

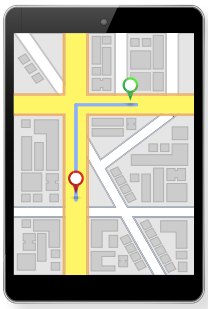

WHY CONTEXT-AWARE DATA PROCESSING IS IMPORTANT

Imagine if Google maps could not identify your location. You would have to know precisely where you are to determine the route to your destination. And getting to your destination would require following pre-defined instructions along the route rather than the voice-guided turn-by-turn directions. Any route deviation could result in your getting lost. This probably sounds similar to using a paper map where the data is provided without context and you are expected to deliver the intelligence.

A Context Exercise

You are a network analyst at a large advertising firm. You work in the corporate office in New York City, but there are offices around the world in Paris, Sydney, and London. Additionally, the firm works with numerous clients, partners, and sub-contractors globally, and there are big files moving across the network constantly. Today, you are analyzing traffic on your network to see if there is something to investigate. Below are three traffic patterns shown in different ways—one without context and the other with context. Which ones would you investigate further?

Without context	With context
An application is connected to remote port 21	An unknown application in Sydney is using FTP and connected to port 21 in North Korea
An application is connected to internal port 22	The IT-approved FTP client in Paris is using FTP and connected to an internal FTP server in London
An application is connected to remote port 5000	An unknown application in New York is using FTP and connected to port 5000 in Iran

Without context, you may decide to investigate all three, or none, or two. But that same traffic with context should make clearer which traffic needs your time and attention.

Google maps has intelligence, which makes it a superior choice for navigation over paper maps. It almost always knows where you are, which means it knows when you have diverted and can quickly reroute you. It can even advise you to change course, because it found an alternate route based on real-time traffic conditions.

Just as context-aware data processing matters for navigation, context-aware data processing matters in network security, too. Security and monitoring tools use network traffic to perform inspection, analysis, and recording. Some security tools like an intrusion detection system (IDS) look at session and application layer data, trying to find pattern matches against a database of threat signatures. But, not every tool is designed for every traffic flow. Before distributing network traffic to your tools, context-aware data processing applies its intelligence to network traffic flows to intelligently distribute only relevant data to security and monitoring tools. For example, this means email monitoring tools will get email traffic only.

Context-aware data processing involves more than simply identifying the type of application traffic. It is about understanding the context of users, devices, and locations, as well as applications. It is about filtering traffic by geography and removing duplicate packets before it reaches a monitoring tool. Performing additional functions like this can be critical to network security monitoring, but it all starts with application intelligence.

PSEUDO APPLICATION IDENTIFICATION AND ASSOCIATED PROBLEMS

Using port numbers can be helpful to identify applications. For instance, File Transfer Protocol (FTP) traffic generally uses Transmission Control Protocol (TCP)/Internet Protocol (IP) port 20/21, Simple Mail Transfer Protocol (SMTP) uses port 25, and Hypertext Transfer Protocol (HTTP) uses port 80. Encrypted web traffic (HTTPS) uses port 443. Any requests to those ports are generally used by FTP, SMTP, and HTTP/S application protocol traffic, respectively. Other applications sometimes use specific ports, like 1433 for Microsoft Structured Query Language (SQL) Server. While using port numbers to identify applications can be helpful, this approach also has problems.

PROBLEM #1: AN UNCOMMON PORT NUMBER IS USED

The first problem with identifying application traffic by port number is that it could be inaccurate. Port numbers can be modified. For instance, a web server administrator may want to change the port from 80 to 8080 or maybe even run two webservers—one on port 80 and the other on port 8080. The same can be done with FTP. FTP could be setup to listen on port 21 or port 5000. While administrators can use just about any available port for any application protocol, they generally do follow conventions. But what if you are monitoring for port 20/21 traffic and a hacker has opened a backdoor in your network to his FTP server listening on port 5000, or worse, he has setup his FTP server to listen on port 80 and transmit on port 443.

PROBLEM #2: PORT NUMBERS ARE SHARED

The second problem with identifying application traffic by port number is that it could be inadequate. For instance, many email services use the same ports for their service. If you are monitoring network traffic looking for threats within emails, it will be difficult to distinguish email traffic from one provider to another. Your company may use Office 365 for its email; however, company policy does not prohibit personal email use. But port-based application identification does not allow you to differentiate among email providers, making it very difficult to monitor only one or a few provider's traffic. The following is a list of some popular email services and the ports they use.

	POP	IMAP	SMTP (SSL)	SMTP	Web-based
Gmail	995	993	465	587	80
Yahoo! Mail	995	993	465	587	80
AOL Mail	995	993		587	80
Office 365	995	993		587	80

In addition to email ports, many web-based applications use port 80 for plain text traffic or port 443 for encrypted traffic, regardless of the functions they perform. With the massive amount of cloud-based applications using port 80 and 443, it is virtually impossible to identify the individual applications using these ports. You may want to monitor back-office applications like Concur and Workday differently than Evernote or Skype, but with port-based filtering, that is not possible. All of these popular cloud applications cannot be monitored separately using port number filtering.

Back Office Apps	Collaboration Apps	Cloud Storage Apps
Concur	Evernote	Box
Workday	Skype	Dropbox
Salesforce	WebEx	Hightail

PROBLEM #3: SIMILAR FUNCTIONALITY BY OTHER APPLICATIONS

The third problem with identifying application traffic by port number is that you could miss alternatives. FTP used to be extremely popular to move files from one server to another. Today, there are numerous file transfer applications that perform services similar to FTP, such as Box, Dropbox, and Hightail. And, they all perform functions using web traffic ports 80 and 443. If you are monitoring network traffic looking for internal exfiltration, you may want to monitor these services specifically. But, filtering all destination port 80/443 web traffic and sending it to a data loss prevention (DLP) device is impractical, as it could overload the DLP with traffic that you do not want to monitor or inspect.

REAL APPLICATION INTELLIGENCE AND THE BENEFITS OF IT

Real application intelligence goes beyond port numbers. Real application intelligence uses a variety of contextual clues to identify known and unknown applications. While the contextual clues may be different, the outcome is similar to Google maps pinpointing your location on a map or the Four Seasons greeting you by name. Real application intelligence has benefits, listed below.

BENEFIT #1 – IT STARTS SMART AND GETS SMARTER

Real application intelligence comes with hundreds of application signatures built-in, as well as a service feed to add new ones and keep everything up to date. This feed is similar to a threat intelligence feed from a security vendor, but the updates are application signatures, rather than threat signatures. And, it automatically starts building a signature for traffic it does not recognize.

BENEFIT #2 – YOU SEE THE APPLICATIONS ON YOUR NETWORK

Real application intelligence shows you what applications are on your network. A dashboard displays statistics, so you can see what applications are generating traffic, the amount of traffic, and the number of sessions contributing to the traffic volume. It also separates known and unknown applications.

BENEFIT #3 – YOU CAN FILTER APPLICATION TRAFFIC AND FORWARD TO SPECIFIC TOOLS

Real application intelligence must be actionable. Identifying specific application traffic and making forwarding decisions based on that traffic gives you control over the security and monitoring tools used to perform analysis. Whether you have a specific type of application traffic you want to monitor constantly, or you want to analyze application traffic that looks suspicious, it is easy to get specific traffic to a tool.

APPLICATION FILTERING VS APPLICATION INTELLIGENCE

Some vendors offer monitoring networks with a feature called “Application Filtering.” This is not the same feature as Ixia’s Application Intelligence. Their Application Filtering offers only capability, not intelligence. Ixia’s Application Intelligence involves both the capability and the built-in intelligence.

Ixia’s Application Intelligence

- Includes more than 200 signatures of the most popular applications on enterprise networks
- Provides ongoing updates of built-in application signatures
- Starts building signatures for applications it does not recognize
- Allows you to quickly build custom application signatures without deconstructing a packet capture

You may be wondering who provides the intelligence needed to get Application Filtering to work. You do. Application Filtering requires you to train (i.e. program) your monitoring network, so it can recognize specific application traffic types. This process is difficult, from analyzing packet captures to writing search expressions. It is also problematic when traffic changes, and the search expressions you wrote no longer recognize specific traffic. Imagine if you had to train Google maps to recognize your location.

Moreover, Application Filtering requires a network administrator to find packets, analyze them, write scripts to identify similar packets, program them into the monitoring network, and maintain them going forward. A more technical guide to this process is listed below.

How To Build Application Signatures for Application Filtering

1. Determine the type of traffic you want to filter
2. Capture the entire user session of that type of traffic, including TCP handshake
3. Isolate the TCP connections in the packet capture
4. Look for different characteristics that will help identify the application, like content-type and domain name
5. Follow each TCP stream, and find unencrypted text that is pertinent, like SSL certificate fields
6. Test different text strings across multiple packets to find the fewest false positives
7. Write a regular expression and/or use the text string to identify the application
8. Program the filter into the monitoring network following a separate set of steps

Ixia’s Application Intelligence has a graphical user interface (GUI) with point-and-click filtering capabilities. No packet captures are necessary. Just click “New Filter” and select the application or applications you want to filter. No technical guide to this process is listed, because no cookbooks or scripting is necessary.

SECURITY FABRIC FROM IXIA

The Ixia Security Fabric™ is powerful network visibility that ensures resilient delivery of relevant traffic to security and monitoring tools. Powered by dual data processing engines, Security Fabric ensures the right data gets to the right tools, even at high speeds.

The Security Fabric's context-aware data processing engine provides administrators with a dashboard of graphs, charts, maps, and statistics of known and unknown applications, geographies, and devices on their network. The engine can also act, by forwarding traffic to specific tools based on contextual details and generating enhanced NetFlow data, such as country, region, browser, operating system, and application name. Ixia's Application Intelligence is unique and is part of what makes its Security Fabric so powerful.

Anatomy of Ixia's Application Intelligence

1. A network tap or Switch Port Analyzer (SPAN) is used to copy network traffic
2. The network traffic is forwarded to the Security Fabric
3. The context-aware data processing engine processes each packet
4. The processing engine checks the five tuples of the packet to see if a session already exists
5. If a session exists, the application is identified
6. If a session does not exist, the engine checks its database of standard and dynamic application signatures
7. If an application signature does not exist in the database, the engine builds a dynamic application signature based on protocol, Secure Sockets Layer (SSL) certificate details, and other factors for future packet identification

Security Fabric's context-aware data processing engine also provides you with control of data conditioning functions to achieve reliability and efficiency in your network of security and monitoring tools.

- **Deduplication** - If you are tapping into multiple network segments for security and monitoring tool analysis, duplicate packets will exist. Sending duplicate packets to the same tool wastes resources. The Ixia Security Fabric removes duplicate packets at line-rate speed to deliver one copy of the packet to the tool(s) you have designated to receive it.
- **Timestamping** - Some forensic and data analysis tools, like security information and event management (SIEM), perform better when they can easily correlate events among device logs. Ixia's Security Fabric can insert a high-accuracy timestamp into every packet at ingress, so logs have more accurate timestamps across devices.
- **Burst protection** - Sometimes your traffic will microburst—a condition where the total bandwidth is temporarily more than your tool port's capacity. When this happens, Security Fabric's deep buffering safeguards against your tool dropping packets by providing your tool a window to catch up.

In order to perform data conditioning functions at line-rate speed, a dedicated hardware accelerator chip is necessary. Generic central processing units (CPUs) simply cannot perform data conditioning functions and process packets fast enough. As a result, packets are dropped when the network gets busy. Ixia's Security Fabric avoids this problem by using a separate, dedicated hardware accelerator chip to perform many packet processing functions to ensure no packets are lost, regardless of network activity.

CONCLUSION

Directing network traffic to the right security and monitoring tool requires seeing beyond port numbers to understand the application used to send the traffic. Otherwise, you may overload security and monitoring tools with traffic they do not need or completely miss the traffic they do. Application intelligence might also be critical in order to comply with your security monitoring strategy. For instance, your company may have subscribed to Box, but not Dropbox. In this case, recording or analyzing Dropbox traffic may be important.

Ixia's context-aware data processing provides you and your monitoring network with additional capabilities and knowledge beyond simple packet delivery. It includes the intelligence to show you what applications and devices are on your network and the power to deliver specific traffic to your tools, even in heavy traffic conditions.

Find out more about context-aware data processing and the Ixia Security Fabric by visiting http://www.ixiacom.com/securityfabric.

ABOUT IXIA

Ixia provides testing, visibility, and security solutions, strengthening applications across physical and virtual networks for enterprises and governments, service providers and network equipment manufacturers (NEMs).

Ixia helps customers manage the unpredictable world of IT and protects them against security threats through actionable insight into the performance, stability and security of their applications and networks. Whether it is testing a product, validating the integrity of a security infrastructure or monitoring a real-time operation, Ixia can help.

WORLDWIDE HEADQUARTERS

26601 W. Agoura Road
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800

(FAX) 1.818.871.1805

www.ixiacom.com

EUROPEAN HEADQUARTERS

Ixia Technologies Europe LTD
Clarion House, Norreys Drive
Maidenhead SL64FL
United Kingdom

Sales +44.1628.408750
(Fax) +44.1628.639916

ASIA PACIFIC HEADQUARTERS

101 Thomson Road,
#29-04/05 United Square,
Singapore 307591

Sales +65.6332.0125
(Fax) +65.6332.0127