Learn About Intrusion Detection and Prevention

This *Learn About* discusses the complex security threats businesses are facing and how the technology behind *intrusion detection and prevention* (IDP) can prevent attacks on business networks. Juniper Networks has offered IDP for years, and today it is implemented on thousands of business networks by the Juniper Networks SRX Series Services Gateways and Juniper Networks JSA Series Secure Analytics appliances.

Staying Open for Business

Security has long been important to network technology. But there is an increasing focus on it today because most business networks are designed to provide access to the Internet and other public networks in order to perform their core operational functions. As shown in Figure 1, a typical business network has several access points to other networks, both public and private. Thus, securing organizational networks as well as their multiple access points is now fundamentally important, if not critical, for businesses to survive.

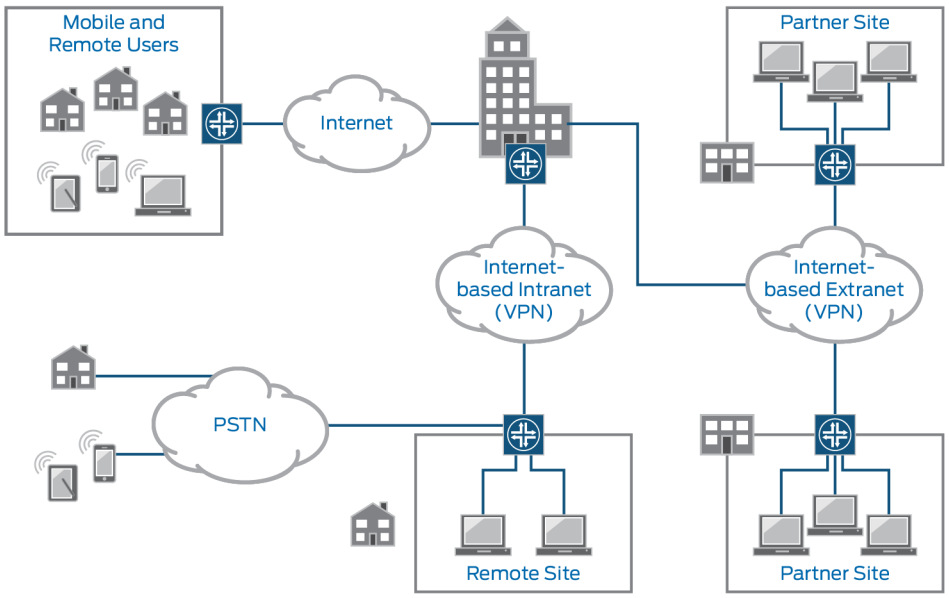

Figure 1 *The Network Today*

The challenge is maintaining the security of these networks while keeping them open to their customers. And of course, the threats are constantly changing; the network-based attacks that caught your attention only yesterday can and will continue to evolve. Currently, attacks are so sophisticated that they can thwart the best security systems, especially those that still operate under the assumption that networks can be secured by encryption or firewalls. Unfortunately, those technologies alone are not sufficient to counter today's attacks.

For example, Figure 2 illustrates the frightening frequency and sophistication of cyberattacks. Today's attackers have increased knowledge and understanding of the technology, infrastructure, and systems of their victims. In addition, the amount of knowledge an attacker needs to know about your network in order to launch a sophisticated attack is decreasing. This means sophisticated attacks are growing more severe each day. (*Source: Citi Online Academy, Digital Security – Cyber Security and Fraud Prevention*)

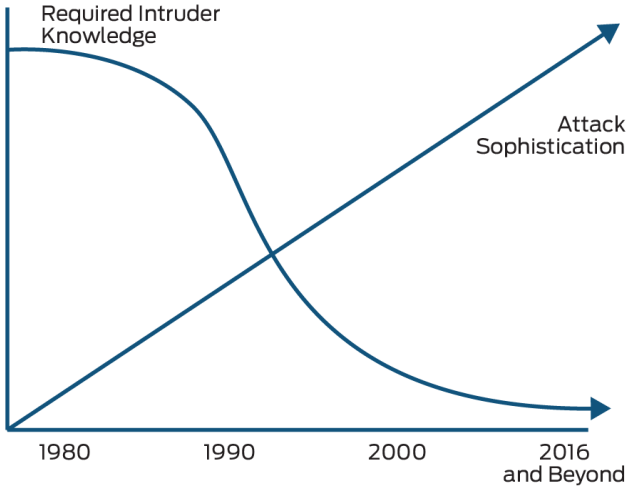

Figure 2 *Attack Sophistication vs. Intruder Technical Knowledge*

While your network still needs firewalls and encryption to improve security, you also need security systems that will watch your network and detect suspicious activity (such as attackers gathering intelligence about your network) 24 hours a day. When these tools observe any suspicious activity or event, they produce alerts for the network administrators. Often, they can detect attacker activity even before the attack begins.

Intrusion detection is the process of monitoring the events occurring in your network and analyzing them for signs of possible incidents, violations, or imminent threats to your security policies. And *intrusion prevention* is the process of performing intrusion detection and then stopping the detected incidents. These security measures are available as *intrusion detection systems* (IDS) and *intrusion prevention systems* (IPS), which become part of your network to detect and stop potential incidents.

How Does IDP Work?

IDP constantly watches your network, identifying possible incidents and logging information about them, stopping the incidents, and reporting them to security administrators. In addition, some networks use IDP systems for other purposes, such as identifying problems with security policies, documenting existing threats, and deterring individuals from violating security policies. IDP systems have become a necessary addition to the security infrastructure of most organizations, precisely because they can stop attackers while they are gathering information about your network.

Most IDPs typically record information and produce reports. But many IDPs can also respond to a detected threat by attempting to prevent it from succeeding. This process can use several different response techniques such as involving the IDP in stopping the attack itself, changing the security environment (for example, reconfiguring a firewall), or even changing the content of the attack.

Figure 3 illustrates the following general components of an IDP solution (note that specific network architecture will differ depending on the exact type of IDP):

- *Sensors* or *agents* monitor and analyze activity on the host, node, or network (the term *agent* is typically used for host-based IDP technologies).
- *Management servers* are available as either software or as an appliance that is the core of the IDP solution. A management server:
  - Manages the agents and sensors, collects data from them, analyzes the data received, and identifies intrusion attempts.
  - Compares events from multiple management servers to see if there are correlations between triggered events (multiple servers are common in larger networks, but not required in smaller deployments).
- *Database servers* act as a centralized repository for event information recorded by sensors, agents, and/or management servers. Many IDPs provide support for database servers.
- The *console* is the administrator program interface to the IDP system and is used to configure agents or sensors, run updates, and monitor and analyze events.

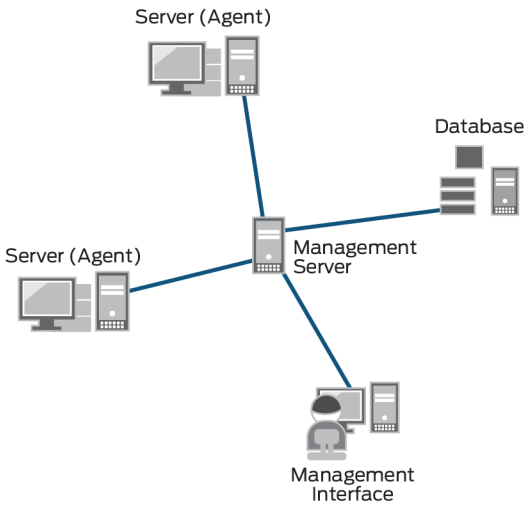

Figure 3 IDP Components

All of this observation generates lots of data. So in addition to monitoring and analyzing events in order to identify undesirable activity, most IDP technologies archive the recorded data locally, although it might also be sent to separate systems, such as centralized logging servers, security information and event management (SIEM) solutions, or enterprise management systems. Reports summarizing monitored events and details can be provided on any event of interest. And once data is flagged as suspicious, the IDP system notifies the security administrators through e-mail, text, messages on the IDP user interface, Simple Network Management Protocol (SNMP) traps, system log messages, and user-defined programs and scripts.

Network-based IDP and *host-based* IDP are two different types of IDP technology characterized by the types of events they monitor and the ways in which they are deployed, as depicted in Figure 4:

- *Network-based* IDP monitors network traffic for particular network segments or devices, and analyzes the network and application protocol activity to identify suspicious activity.
  - Wireless components monitor wireless network traffic and analyze it to identify suspicious activity involving the wireless networking protocols.
  - Network behavior analysis (NBA) examines network traffic to identify threats that generate unusual traffic flows, such as distributed denial-of-service (DDoS) attacks, certain forms of malware, and policy violations (for example, a client system providing network services to other systems).
- *Host-based* IDP monitors the characteristics of a single host, and the events occurring within that host, for suspicious activity.

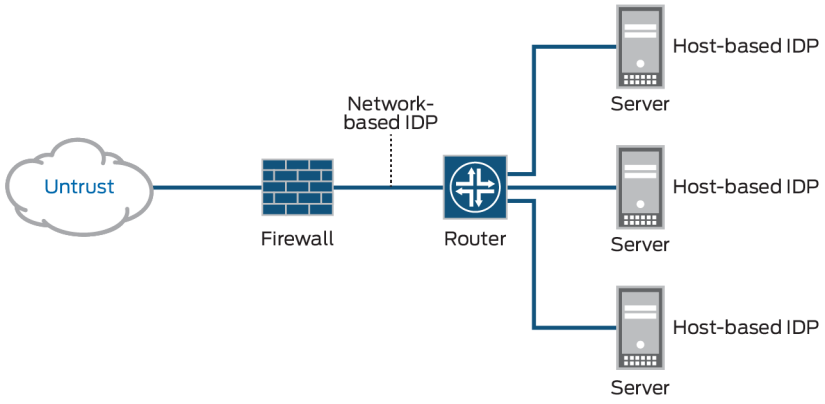

Figure 4 *Deploying Network-Based and Host-Based IDPs in a Network*

Most IDP technologies use multiple detection methodologies, either separately or integrated, to provide wider and more accurate detection. Table 1 lists three IDP detection methodologies (*signature-based*, *anomaly-based*, and *stateful protocol analysis*) that are typically used to detect incidents.

Table 1 Common IDP Detection Methodologies

Detection Method	Description	Example	Benefits
Signature-Based Detection	Signature-based detection compares signatures against observed events to identify possible incidents.	A telnet attempt with a root username, which is a violation of an organization's security policy or an e-mail, with *Free Pictures* in the subject line, and an attachment with the filename *freepics.exe*, all of which are characteristics of a known form of malware.	This is the simplest detection method because it compares only the current unit of activity (such as a packet or a log entry, to a list of signatures) using string comparison operations.
Anomaly-Based Detection	Anomaly-based detection compares definitions of what is considered normal activity with observed events in order to identify significant deviations. An IDP using anomaly-based detection has profiles that represent the normal behavior of such things as users, hosts, network connections, or applications. The profiles are developed by monitoring the characteristics of typical activity over a period of time.	A profile for a network might show that Web activity comprises an average of 13% of network bandwidth at the Internet border during typical workday hours. The IDP then uses statistical methods to compare the characteristics of current activity to thresholds related to the profile, such as detecting when Web activity consumes significantly more bandwidth than expected and alerting an administrator of the anomaly.	This detection method can be very effective at spotting previously unknown threats.
Stateful Protocol Analysis	Stateful protocol analysis compares predetermined profiles of generally accepted definitions for benign protocol activity for each protocol state against observed events in order to identify deviations. Unlike anomaly-based detection, which uses host or network-specific profiles, stateful protocol analysis relies on vendor-developed universal profiles that specify how particular protocols should and should not be used. The *stateful* in stateful protocol analysis means that the IDP is capable of understanding and tracking the state of network, transport, and application protocols that have a notion of state.	When a user starts a File Transfer Protocol (FTP) session, the session is initially in the unauthenticated state. Unauthenticated users should only perform a few commands in this state, such as viewing help information or providing usernames and passwords. An important part of understanding state is pairing requests with responses, so when an FTP authentication attempt occurs, the IDP can determine if it was successful by finding the status code in the corresponding response. Once the user has authenticated successfully, the session is in the authenticated state, and users are expected to perform any of several dozen commands. Performing most of these commands while in the unauthenticated state would be considered suspicious, but in the authenticated state performing most of them is considered benign.	This analysis identifies unexpected sequences of commands, adds stateful characteristics to regular protocol analysis, and adds reasonableness checks for individual commands (for example, min/max lengths).

Choosing IDP Systems

IDP technologies can provide a wide array of security capabilities for your network. Look for these common, but necessary, security capabilities:

- **Information gathering**: Systems that identify hosts and the operating systems and applications being used, as well as identifying general characteristics of the network.
- **Logging**: Your IDP should perform extensive logging of data related to detected events. This data can be used to confirm the validity of alerts, investigate incidents, and correlate events between the IDP and other logging sources. You should know that specific types of IDPs log additional data fields, such as network-based IDPs that perform packet captures, and host-based IDPs recording user IDs. Your IDP should permit administrators to store logs locally and send copies of logs to centralized logging servers (for example, syslog, security information and event management software). Also, your IDP should ideally synchronize its clocks using the Network Time Protocol (NTP), or through frequent manual adjustments, so that log entries have accurate timestamps.
- **Detection**: IDP technologies should typically offer extensive detection capabilities. The types of events detected and the accuracy of detection can vary greatly depending on the type of IDP technology being used. Most IDPs require at least some fine-tuning and customization, such as setting prevention actions to be enabled for particular alerts, to improve their detection and effectiveness.
- **Prevention**: Finally, most IDPs should offer multiple prevention capabilities. While the specific capabilities vary by IDP technology type, your IDP should allow administrators to specify the prevention capability configuration for each type of alert, including enabling or disabling prevention, as well as specifying which type of prevention capability should be used.

Juniper Networks IDP

Juniper Networks uses its SRX Series Services Gateways for Intrusion Detection and Prevention services and its IDP policy configuration lets you selectively enforce various attack detection and prevention techniques on the network traffic passing through your chosen SRX Series device (the IDP-enabled device). You can define policy rules to match a section of traffic based on a zone, a network, or an application, and then take active or passive preventative actions on that traffic. The SRX Series device contains a full set of IDP signatures to secure networks against attacks.

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks website. This database includes attack object groups that you can use in IDP policies to match traffic against known attacks. Although you cannot create, edit, or delete predefined attack objects, you can use the CLI to update the list of attack objects that you can use in IDP policies.

The SRX Series device can forward packet capture (PCAP) data from its traffic to a Juniper Secure Analytics (JSA) appliance using the PCAP Syslog Combination Protocol. With the PCAP Syslog Combination Protocol, the JSA appliance is capable of receiving both syslog and the additional PCAP data once configured with the SRX Series.

This Juniper Networks IDP system is shown in Figure 5, in a very small site deployment that larger networks can scale. The SRX Series device displays the visibility of incoming or outgoing traffic and the JSA appliance collects events, allowing real-time streaming of events and monitoring of events through a common dashboard.

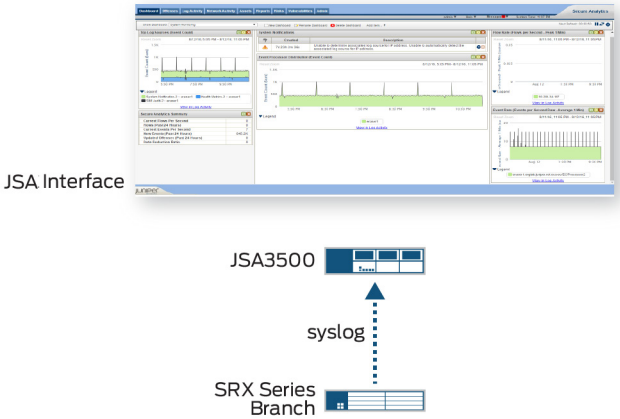

Figure 5 *Small Site Deployment – JSA Appliance*

Packet capture data is forwarded to the JSA appliance on a specified port, which is separate from the port that receives forwarded syslog data. The data contained in the packet capture and the outgoing port from the SRX Series is all configured using the SRX Series user interface.

Use Case: Protect Server and Application Vulnerabilities

Let's employ a simple use case to examine how the Juniper IDP system works. Assume that *Company X* is hosting its own commercial website as shown in Figure 6. Traffic is sent to the SRX Series services gateway for monitoring.

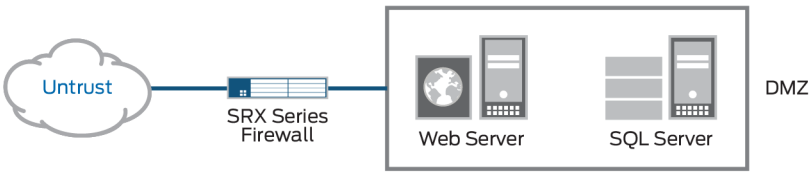

Figure 6 *Company X Network Overview*

When the traffic is sent to the SRX Services gateway, it is discovered that the company's website is vulnerable to a specific SQL injection attack as shown in Figure 7. Packet capture provides the following details of the attack:

- The external connections are coming from the UNTRUSTED zone.
- The webserver (10.10.10.80) is located in our DMZ zone.
- The attack happens over TCP/80 or HTTP.

- The attack uses the GET command.
- The attack uses the following pattern:

```
form.php?q=1+UNION+SELECT+VERSION
```

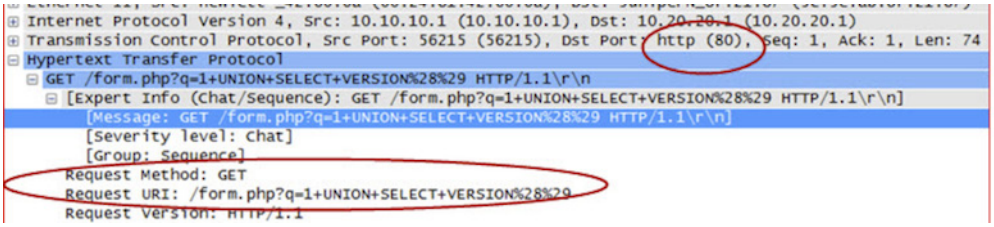

Figure 7 *Packet Capture*

Once the SRX Series services gateway has been spotted, and alerts are sent, the administrator can create custom attack objects to detect SQL injection as shown in Figure 8.

```
root@SRX# set security idp custom-attack HTTP:CUST_SQL_INJECT

root@SRX# edit security idp custom-attack HTTP:CUST_SQL_INJECT

[edit security idp custom-attack HTTP:CUST_SQL_INJECT]

root@SRX# set severity major

root@SRX# set attack-type signature protocol tcp destination-port match equal value 80

root@SRX# set attack-type signature direction client-to-server

root@SRX# set attack-type signature context http-get-url

root@SRX# set attack-type signature pattern "/form\.php\?q=1/+UNION/+SELECT/+VERSION"
```

Figure 8 *Creating a Custom Attack Object*

Each rule is composed of match conditions, objects, actions, and notifications. When you define an IDP rule, you must specify the type of network traffic you want IDP to monitor for attacks by using the following characteristics – source zone, destination zone, source IP address, destination IP address, and the Application Layer protocol supported by the destination IP address. The rules are defined in rulebases, and rulebases are associated with policies.

Figure 9 shows the result – the network drops the attack when the SQL injection attack is attempted.

```
name="SERVICE_IDP" application-name="NONE" rule-name="1" rulebase-
name="IPS" policy-name="COMPANY_X" repeat-count="0" action="DROP"
threat-severity="HIGH" attack-name="HTTP:CUST_SQL_INJECT" nat-source-
address="0.0.0.0" nat-source-port="0" nat-destination-address="0.0.0.0" nat-
destination-port="0" elapsed-time="0" inbound-bytes="0" outbound-bytes="0"
inbound-packets="0" outbound-packets="0" source-zone-
name="UNTRUSTED" source-interface-name="fe-0/0/7.0" destination-zone-
name="DMZ" destination-interface-name="fe-0/0/6.0" packet-log-id="0"
message="-"]
```

Figure 9 *The Attack Is Now Known to the Network*

References and Resources

- The Juniper TechLibrary documentation includes everything you need to understand Juniper's IDP system. See http://www.juniper.net/techpubs/en_US/junos15.1x49-d40/information-products/pathway-pages/security/security-idp-index.html.
- A tech note on Juniper SRX Series device forwarding of packet capture (PCAP) and syslog data to the JSA appliances. See http://www.juniper.net/techpubs/en_US/jsa2014.7/information-products/topic-collections/jsa-managing-juniper-pcap-data.pdf.
- The SANS Reading Room maintains, and makes available at no cost, a wide collection of research documents about various aspects of information security. It features over 2,460 original computer security white papers in 96 different categories. See http://www.sans.org/reading-room; http://www.sans.org/reading-room/whitepapers/detection/network-ids-ips-deployment-strategies-2143.
- The TechTarget network of technology-specific websites gives you access to industry experts, independent content, and analysis. See:
  - http://searchsecurity.techtarget.com/Do-you-need-an-IDS-or-IPS-or-both
  - http://searchsecurity.techtarget.com/feature/Enterprise-benefits-of-network-intrusion-prevention-systems
  - http://searchsecurity.techtarget.com/feature/The-basics-of-network-intrusion-prevention-systems
- Webopedia is an online tech dictionary for IT professionals, educators, and students. It also provides in-depth articles, study guides, and links to sources of further information on the topic, where applicable. See http://www.webopedia.com/DidYouKnow/Computer_Science/intrusion_detection_prevention.asp.
- The Computer Security Resource Center (CSRC) is the primary gateway for gaining access to NIST computer security publications, standards, and guidelines plus other useful security-related information. See http://csrc.nist.gov/publications/nistpubs/800-94/SP800-94.pdf.
- This O'Reilly book is a complete field guide, authorized by Juniper Networks, and is the perfect hands-on reference for deploying, configuring, and operating Juniper's SRX Series networking device. See http://chimera.labs.oreilly.com/books/1234000001633/ch13.html.
- World Academy of Science, Engineering and Technology (WASET) is a scholarly open science, peer reviewed, interdisciplinary, monthly and fully referred international research journal focusing on theories, methods, and applications in Science, Engineering, and Technology. WASET serves as a forum for scholarly intellectual exchange and as a platform to present cutting-edge research. See http://waset.org/publications/14713/network-based-intrusion-detection-and-prevention-systems-in-ip-level-security-protocols.
- CST (Computer Security Technology) provides consultancy services and managed security services for IT departments that may lack the time, resources, or expertise to handle security themselves. CST complements your own resources and helps fill any resourcing and skill gaps within your own security posture. See http://www.cstl.com/Products/Juniper/Juniper-IDP-Solution/WhitePaper/Juniper-IDPWhitePaper.pdf.

Learn About Intrusion Detection and Prevention

by Keerthi Latha M R

Intrusion detection *is the process of monitoring the events occurring in your network and analyzing them for signs of imminent threats; intrusion prevention is the process of stopping the incidents detected through* intrusion detection*. Together they are a formidable team blocking unwanted access to your network and reinforcing its security capabilities. Learn about these technologies and how to integrate Juniper Networks IDP Services to be part of your network.*

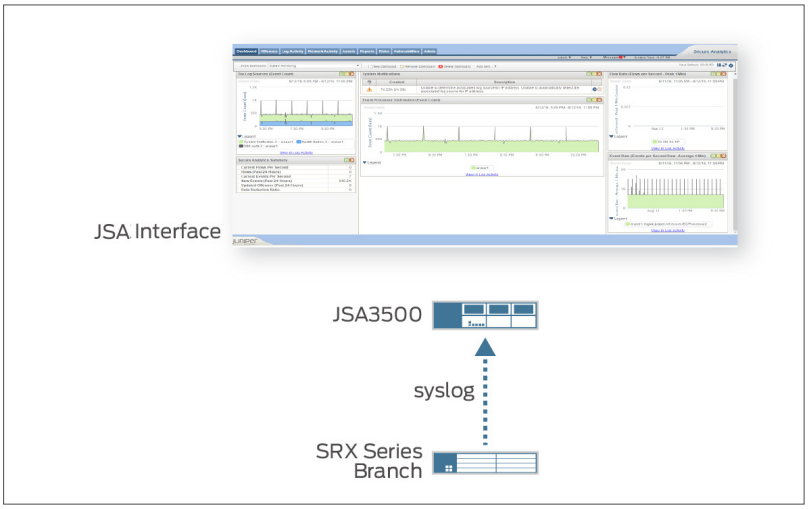

About the Author:
Keerthi Latha M R is an Information Development Engineer at Juniper Networks with over 10 years of experience in writing and developing documentation for networking and telecommunications.

Version History: First Edition, September 2016 2 3 4 5 6 7 8 9

For more information go to the TechLibrary at: www.juniper.net/documentation