



Intrusion Prevention

What is it and what do businesses need to do about it?

What is Intrusion Prevention?

ID, IDS, IDP, and IPS – all of these terms are circulating in today's security environment. They are often used interchangeably, but don't necessarily mean the same thing. Also, the term 'firewalling' is used in a slightly different way when intrusion detection and prevention are discussed. In this context, "firewall" is generally understood to be "stateful packet filtering" or "stateful inspection" which does not extend to application layer security.

Following are a few definitions of the most commonly used terms for segment of the security market.

ID – Intrusion Detection -- the process of monitoring network activity and reporting on suspicious behaviour.

IDS – Intrusion Detection System -- developed as a **passive** detection, recording, and forensic tool. An IDS usually comprises two components, a management centre where data is stored and configuration parameters are set up with one or more sensors. These sensors are typically connected to a mirroring or "span" port of a switch so that they can "see" all the network traffic on that switch. The sensors transmit the data back to the management centre for storage and analysis. IDSs are an expensive and complex component of network security, usually available only to larger enterprises.

IDP – Intrusion Detection and Prevention, the process of monitoring network activity, reporting on suspicious behaviour, and then actively blocking (preventing) attacks.

IPS – Intrusion Prevention System is an **active, in-line blocking** system that looks for intrusion attempts at both the network layer and application layer and blocks these attempts in real-time. It incorporates many of the technologies of IDS, but has more stringent performance requirements since it must evaluate threats in real-time without introducing unacceptable performance bottlenecks or latency.

Intrusion Prevention Technologies

Intrusion prevention systems commonly combine the following methods of attack recognition:

- **Protocol Anomaly**
- **Signature-Based**
- **Behaviour-Based**
- **DOS and DDOS**

Protocol Anomaly

Protocol anomaly attacks work because most devices expect a protocol to be used properly and so they don't do range checking or any enforcement of the proper use of the protocol. This occurs at all layers, for example, at the TCP layer, fragmentation attacks such as Teardrop rely on improperly overlapping packet fragments to create dangerous content from innocuous fragments. At the application layer are even more opportunities for attack - FTP bounce is a good example. It relies on the FTP server not checking to see if the connecting client IP address is the same one the client specifies for a file to be sent to or taken from. Buffer overflow attacks are another example of exploiting inadequate range checking in many types of server.

Protocol Anomaly Detection works by modelling the protocol used and comparing the traffic with known correct behaviour. If the behaviour is not correct, then the traffic is discarded.

Signature-Based

Signature-based detection compares the traffic with known patterns from specific attacks. Currently there are between 1,000 and 5,000 known patterns and this is growing daily. This form of detection when all signatures are compared with all packets leads to the most occurrences of false positives, as a pattern might match even though it is not an attack. This presents a major issue for network administrators.

A recent improvement in this field has been to look for known attack signatures only during the particular phase of a protocol when the attack could succeed. For instance, the Wiz attack can only work when an SMTP server is in command mode – so this is the only case when we need to compare the traffic to the Wiz signature. This does require the IPS to model the protocol in a very similar way to that required by Protocol Anomaly Detection.

The drawback of this type of detection is that attacks without identified signatures can still get through. As attacks are designed that propagate increasingly quickly, e.g. the Slammer worm propagated worldwide in 8 to 10 minutes, the effectiveness of signature-based detection (unless combined with other detection forms) will decline.

Behaviour-Based

Behaviour-based detection looks for certain behaviours that are almost always hostile. This is a relatively new technology, and there are many forms of behaviour that can be identified and ways to use that information. For example:

- Recognising precursors to attacks such as port scans
- Recognising attack behaviours such as spoofing and the illicit use of IP options

- Host Profiling where traffic patterns for network hosts are identified and irregularities such as traffic on a new port, or uncharacteristically large volumes of traffic are flagged. Once identified, adaptive countermeasures are taken to block the intruder.

Denial of Service (DOS) and Distributed Denial of Service (DDOS)

Denial of service attacks flood the network with traffic in order to block access to, or tie up mission critical resources. Forms of DOS attack include Ping Flood, SYN flood, and Smurf attack. Detecting and preventing these attacks typically relies on algorithms that look for unusual or excessive traffic and discard the traffic on that basis. For example, large numbers of incomplete TCP handshakes (SYN requests) from a particular IP address.

Distributed Denial of Service uses an array of compromised systems across the Internet to flood a site. This is much more difficult to prevent since any one attacking system may not be generating abnormally large volumes of traffic so the total traffic flow, although large enough to result in denial of service, may not look to be abnormal. Examples of this are Trin00, TribeFlood Network (TFN) and Stacheldraht. Detecting and blocking distributed denial of service typically requires even more sophisticated modelling of traffic behaviour.

Intrusion Prevention or Intrusion Detection?

Intrusion Detection

- Only detects and doesn't prevent attacks. Because they sit in parallel to the network, IDS sensors are generally unable to stop an attack in progress
- Requires manual intervention to prevent an attack, by which time the damage may have already been done
- Typically generate more false alarms as they look at a lot more (not necessarily relevant) traffic
- Generate huge amounts of log data – which needs to be interpreted to be useful
- Sees all parts of the network (multiple sensors)

Intrusion Prevention

- Is inline and prevents attacks automatically
- Generates few false alarms
- Generates low log data volumes as only allowed traffic or recognised traffic need be logged
- Is generally used at the network perimeter so doesn't necessarily see all parts of the network

In summary, for small and midsize companies intrusion prevention represents a more practical and useful solution to the problem. Intrusion detection is still useful in larger environments where it is important to know much more about internal network traffic.

What is the impact for companies that don't have intrusion prevention?

A company that is relying on firewalling (in this context Stateful Packet Filtering) as their primary defense is going to be vulnerable to attacks that utilise legitimate protocols that must be allowed through the firewall. At a minimum, these protocols are likely to be HTTP, FTP, SMTP and DNS, which probably constitute around 98% of all traffic. These companies may also be vulnerable to DOS and DDOS attacks if the firewall is incapable of recognizing the attack.

What do businesses need to do about it?

Businesses need to add intrusion prevention technology to their defences. This can be done in two ways -- by placing an IPS system inline with their firewall, or by replacing their firewall with a firewall (or security gateway) that combines Stateful Packet Filtering with intrusion prevention technologies.

Having separate firewall and IPS systems raises some issues:

- If the firewall is outside and is not able to handle DOS or DDOS, these attacks might succeed as they overload the firewall before the IPS system is able to mitigate the attack
- Separate management consoles increase complexity
- Configuration complexity may lead to vulnerabilities
- Separate logs make auditing more difficult

Combining firewalling and IPS in a single appliance has inherent advantages:

- Better DOS/DDOS protection
- Simplified management and less likelihood of configuration error
- Unified logging
- Lower cost
- Lower latency (not processing everything twice)