

# Holistic IPS: The Convergence of Intrusion Prevention Technologies

By Avi Chesla

Over the last several years, the intrusion detection and prevention (ID&P) market has emerged from simple detection (and notification) capabilities to fully automated intrusion prevention solutions. In the “early days,” ID&P products were developed by small vendors whose focus was to address deficiencies in the emerging network security infrastructure. These early generation products were introduced to the market as complementary solutions to traditional firewalls and anti-virus products. Over time, vendors have introduced increasingly sophisticated products that often provide multiple security functions. Despite the sophistication of these new products, they still primarily rely upon one of the following two intrusion detection techniques:

- ▼ **Signature (or pattern) matching.** This method is sometimes called “content-based” and it incorporates what is often referred to as “deep packet inspection” capabilities.
- ▼ **Behavioral assessment.** This method is sometimes referred to as anomaly-based, statistical-based, or rate-based.

When implemented correctly, each of these technologies responds to a different type of threat, and therefore these technologies do not overlap, but rather complement each other.

Historically, Intrusion Prevention Systems (IPS) vendors provided only one detection/prevention method. Recently, in response to market demand, vendors have begun to offer products that incorporate elements of both detection methods, in what is now called the *hybrid* approach to IPS. However, these products apply the two disparate detection methods independent of each other and do not combine information obtained via each of the methods to improve the accuracy and performance of the overall IPS product.

As such, in order to fully leverage the capabilities of the hybrid approach, the two detection methods need to be integrated in an intelligent fashion. This integration is achieved by employing a decision technology embodied by a “*Correlation Engine*.” This engine harmonizes the two detection methods and creates a single, effective, accurate, efficient, and maintainable holistic IPS methodology.

This article introduces the concept of the Correlation Engine, and describes how it represents the next natural step in the evolution of intrusion detection and prevention. It begins with a brief overview of the main network attack threats. In this context, it outlines the characteristics and limitations of each of the two intrusion detection technologies. It then describes how appropriate correlation between these detection methods can provide a vastly improved level of preventive capabilities. It then suggests several correlation techniques and applicable system architectures. The article concludes with a description of the value that the correlation engine brings to the information security market.

## Network Attack Threats

In practice, there are three phases of operation:

1. **Intelligence (Information Gathering)**—A typical intrusion into computer networks involves pre-attack probe scanning activities, which help the attacker gain valuable knowledge about the target networks.
2. **Attack Planning**—Using the knowledge gained during the intelligence phase, a cyber-assailant can decide which attack type will be most effective in harming the target network.
3. **Attack Execution**—Most network attacks can be executed using readily-available attack tools.

## IPS Technologies

The two prevalent IPS technologies in use today are *signature-based* and *behavior-based*. Products employing either method are typically placed “in line” with network traffic, at key junction points, where they can evaluate traffic entering or leaving the “protected” networks.

The different types of intrusion detection methods are described below: **Signature-based technologies** were the first employed to detect network intrusions. A signature-based product matches data packet contents to a pre-defined set of known attack “fingerprints.” When a match is found, the system raises an appropriate alert and prevents the attack traffic from entering the network. To allow a signature-based system to perform properly (without excessive amounts of false positives and misdetections), the system must be continuously updated with the most recent attack signatures.

Signature-based detection methods are more effective against known, “single-bullet” attacks. Single bullet refers to attacks that can be perpetrated via a single packet (or using a very small number of packets). Due to the small amount of data involved in these attacks, it is very difficult to detect the attacks using behavioral techniques. Single-bullet attacks often exploit a known vulnerability in an application or computer operating system. Examples of single-bullet attacks include buffer overflows, protocol anomalies, viruses, and others system vulnerability exploitation types.

**Behavior-based technologies** represent a newer approach to detecting network intrusions. Products that use behavior-based methods construct a complete picture of communication patterns of entities that span the IPS device in order to establish baselines that best fit the protected network’s normal communication activities. Behavior-based products include a decision engine that is responsible for the following tasks:

- ▼ Detecting deviations from the network’s normal baselines,

- ▼ Deciding how a deviation from these baselines threaten the network, and
- ▼ Implementing effective countermeasures against the detected threats.

Behavior-based methods are primarily effective against “continuous behavior attacks.” These types of attacks take effect through a collection of events (not through a single event). These attacks are often cloaked within legitimate packet structures and apparently “legitimate” data packet contents. This means that pre-defined signatures are not effective in detecting the attack. Examples of these types of network attacks include the following:

- ▼ Denial of service flood attacks
- ▼ New/unknown worm
- ▼ Network and application pre-attack probes
- ▼ Brute-force and dictionary attacks
- ▼ Other attack types that can take effect by generating an abnormal sequence of network events

Because the two main attack types are detectable by either signature-based or behavior-based methods (but not both), they represent complementary techniques for preventing network attacks.

### The Challenge: Holistic IPS

Although most of today’s IPS’s comprise both signature and behavioral technologies, in order to maximize the effectiveness of the system, it is not enough to put one and one together.

The challenge is to be able to merge these two distinct technologies into a one, more powerful technology. In order to do this, an additional Intelligence needs to wrap the existing technologies. This intelligence, defined in this article as the correlation engine, analyzes alerts coming from both detection technologies, decides about the level of correlation between these alerts and generates countermeasures accordingly. Thus, more accurate detection and prevention is achieved, pro-active countermeasures can be implemented, and the level of forensics analysis information is a lot higher. The following section introduces the concept of IPS correlation.

### Correlation: An Additional Layer of Intrusion Prevention Intelligence

There are two objectives that an intrusion prevention system achieves by implementing a signature/behavior correlation engine:

- ▼ Automated threat level analysis for each detected attack
- ▼ Automated generation of accurate countermeasures

The first step in establishing an effective correlation is to choose what type of security events will enable the system to perform a streamlined decision process.

#### Pre-Attack Probe Events

Just as in military situations, a network attack begins with probing for the “enemy’s” vulnerabilities. A typical cyber-attack starts with pre-attack probes, where the assailant scans for network or application vulnerabilities. In order not to raise suspicion, current pre-attack probe techniques are capable of continuously changing their scanning rate, conducting slow-rate scans, and sending decoy information during the scan—thereby making these techniques hard to detect. Although these events themselves do not

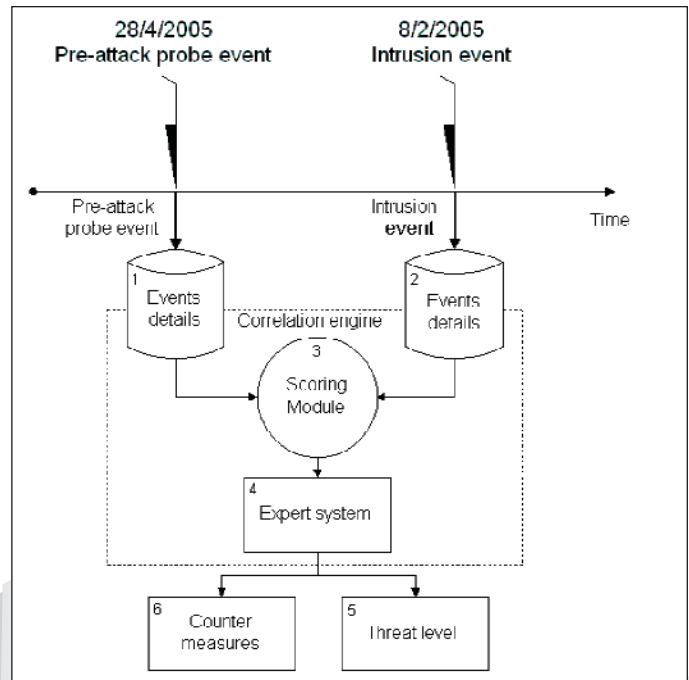


Figure 1: Events Correlation Process

impart a threat, identifying them is an important step in the process of threat analysis.

Information gleaned from pre-attack probes, such as the sender’s origin, the hosts targeted, the targeted ports and applications, and the protocol type used all help in understanding which network resources attract the most interest. A high-interest level necessitates an increased state of alertness for the resource “under attack.”

Because pre-attack probes can be characterized as suspicious or abnormal behavior, which doesn’t necessarily violate deterministic protocol rules or match pre-defined attack signature, behavior-based intrusion prevention technologies will probably be most effective in detecting them.

#### Intrusion Events

An intrusion event can be characterized as an attempt to exploit vulnerable applications in order to take unauthorized control over them. Intrusion events are detected both by signature-based and behavioral-based intrusion prevention technologies in a complementary manner.

Each intrusion will necessarily have a unique (and dynamic) threat level. The threat level is a function of intrusion type and the target type. For example, an intrusion event designed to produce a denial of service (DoS) condition can be a high-level or low-level threat. If the target (e.g. server) has a known (related) vulnerability, then the threat will be considered high. However, if the target is not vulnerable to this type of DoS, then the threat level will be low. Most intrusion prevention systems already include some internal threat analysis mechanisms that are responsible for deciding about the level of threat of each event. However, this decision is taken without any relationship to other previous events.

### Correlation Engine

For today’s hybrid systems, the two detection methods operate independently.

The correlation engine proposed in this article strives to improve detection/prevention capabilities by correlating different types of events generated at different times by different types of detection methods.

The following diagram proposes a correlation process between an intrusion detected by a signature-based product, and a pre-attack probe as detected by a behavior-based product.

As shown in Figure 1, the proposed correlation process can be divided amongst two inference modules, which are the Scoring Module and the Expert System. Each of these is described in the following section.

### Scoring Process

The Scoring Module tracks attack event details that are maintained in the event databases of each of the detection engines and weighs their severity. This module compares details that are captured as signatures and the pre-attack probe events that are captured by the behavior-based engine. The Scoring Module assigns a correlation value for the similarity of information that is detected by the two detection engines. The different score levels are shown in Figure 2.

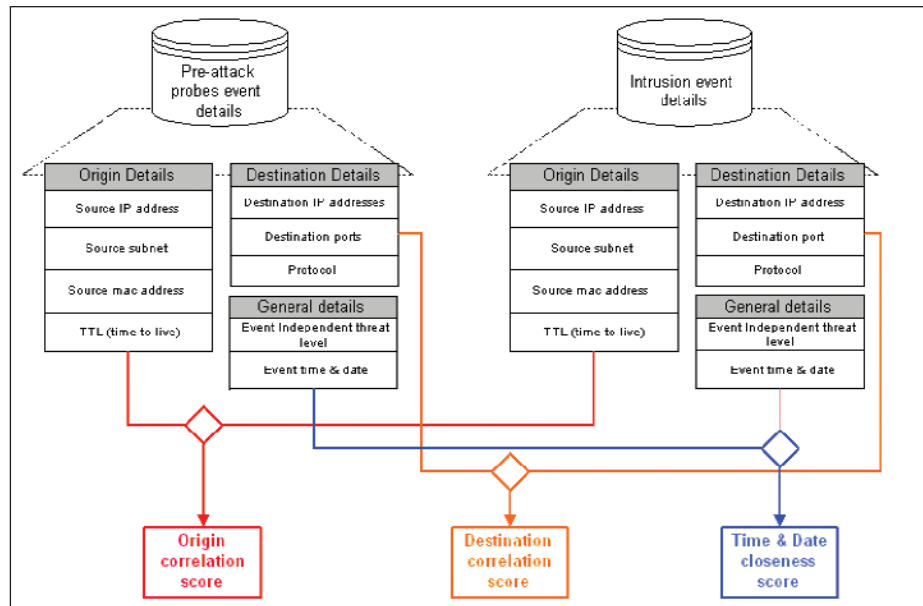


Figure 2: Scoring Process

Each of the three correlation score types shown in Figure 2 has a different meaning, as follows:

**Origin Correlation Score**—This score indicates the level of similarity between potentially similar events' "source" information. Origin details include the following data:

- ▼ Source IP address or source subnet
- ▼ TTL values (that reflect the number of hops that suspicious packets passed before they were detected)
- ▼ source MAC address (which reflects the last hop's physical interface address from which the packet was received)

A high Origin Correlation Score means that there is a history of suspicious pre-attack probe events that were generated from the same network region (i.e., subnet mask and TTL) or even from the same host.

**Destination Correlation Score**—This score indicates the level of similarity between potentially similar events' "destination" information. Destination details include the following data:

- ▼ Destination IP address, or addresses range
- ▼ Destination port
- ▼ Protocol type

A high Destination Correlation Score means that there is a history of suspicious pre-attack probe events which had similar destination characteristics.

**Time & Proximity Score**—A high level score means that two events occurred in close time proximity.

### Expert System Processes

The correlation scores described in the previous section feed the Expert System (shown in Figure 3). The Expert System contains embedded logic (i.e. expert rules) that mirrors the thought processes of a security expert. These rules are responsible for transforming sets of correlation scores and event details into operative actions. Actions include setting the event's threat level and choosing the most appropriate countermeasures (i.e., blocking rules).

Figure 3 demonstrates how the Expert System sets the threat level, which is the first process conducted by the Expert System:

As shown in Figure 3, each correlation score is mapped into the expert system's high, medium and low membership function space. This association method fosters the appropriate synthesis of correlation scores with the linguistic if/then expert rules into decisions, i.e., a single threat level. Moreover, using this type of inference system allows adding, removing or modifying the expert rules without applying any changes to the system itself.

The following is an example of an application of these rules and the consequent threat level:

- a. If OCS is HIGH AND DCS is HIGH THEN Threat Level is **HIGH**
- b. If OCS is LOW AND DCS is LOW THEN Threat Level is **LOW**
- c. ...

Adding more expert rules in a consistent manner will improve the level of intelligence of the system.

After applying the rules to a set of events, the system aggregates the results of each one and averages the results in order to set a single threat level.

### Constructing Pro-Active Blocking Rules

After determining a threat level, the Correlation Engine then produces automated countermeasures that will effectively mitigate the detected intrusion event. To do this, the Expert System collects the origin event details and destination event details that were found to have high correlation scores and organizes them according to pre-defined rules in order to construct pro-active blocking criteria.

The following scenario describes the progression of a representative attack and the IPS response. This describes the correlation process; in particular the process in which pro-active blocking rules are constructed. Each step in the progression is explained by way of examples, below:

**1ST Event**—The IPS has detected suspected activities:

- ▼ generated from a particular host (190.1.2.5)
- ▼ aimed to probe port 25 (i.e., mail application) on a number of local hosts (IP addresses 200.0.0.2–200.0.0.20) in the protected network

The suspected probing activities were detected by the behavioral-based detection engine. Event details of these “abnormal” mail application activities were recorded.

**2nd Event**—An attempt to exploit a mail application’s known vulnerability was detected by the IPS. The exploit was detected by the signature-based detection engine. The attempt was:

- ▼ Generated from a particular host (190.1.2.5)
- ▼ Aimed at port 25 on the host with the IP Address - 200.0.0.5, in the protected network.

**Event Correlation**—The Correlation Engine collects the details of the second event in real-time and analyzes the level of correlation between these, to previous events.

In this scenario the following event details were discovered to have a high level of correlation:

1. The source IP address of the first and second events is the same (190.1.2.5).
2. The destination port (25) of the first event and second event is the same.
3. The IP address of the attacked host (200.0.0.5), associated with the second event, is one of the probed hosts involved in the first event.

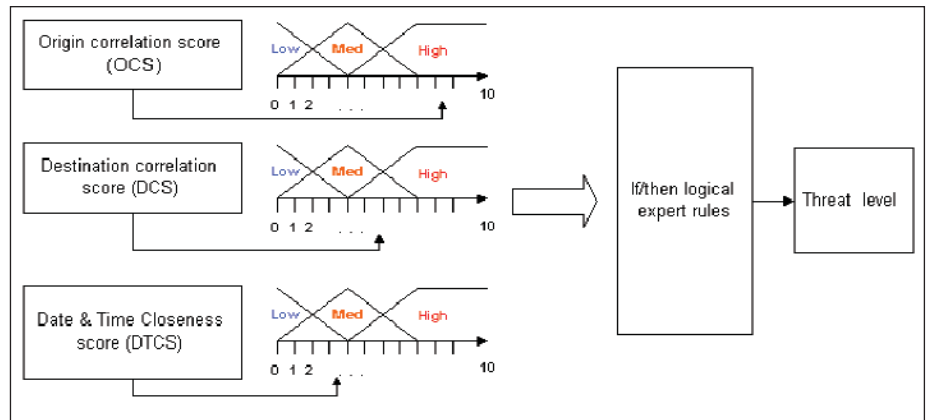
The Correlation Engine performs the following operations:

1. **Alert**—Because both the origin and destination details have high correlation scores, a **High**-level threat alert is generated.
2. **Block**—Pro-active preventive measures are implemented immediately. All further activities from the same source IP address to port 25 (i.e., mail application) on all probed hosts in the protected network are blocked.

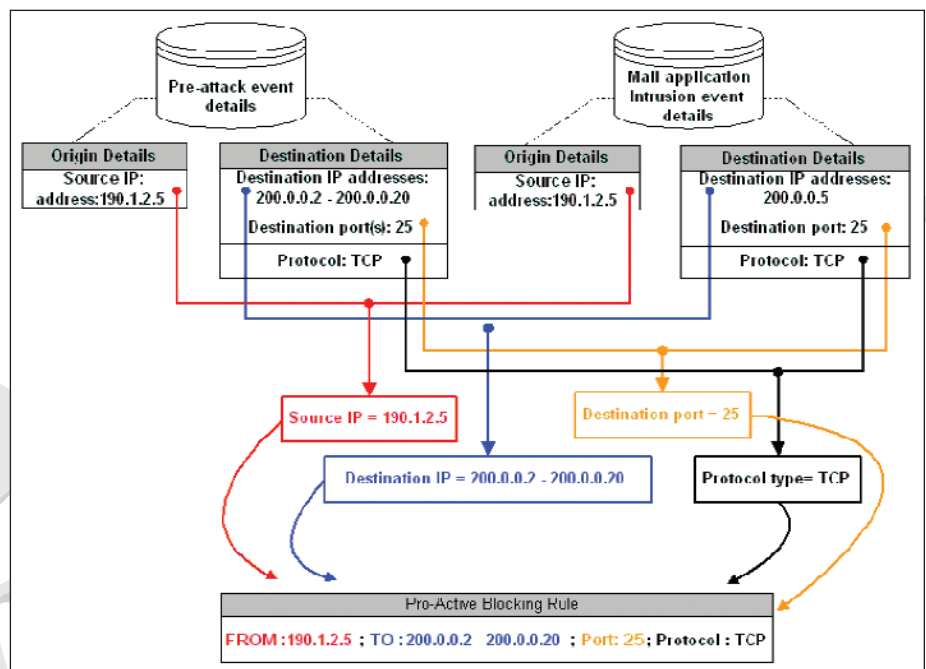
Figure 4 demonstrates a correlation process that results in a blocking rule that accurately addresses the threat. The blocking rule shown is pro-active because it includes blocking criteria that address future possible intrusion targets. In the case shown, future intrusion targets are determined according to the destination details of the pre-attack probe event, and which were detected earlier. In other words, the range of IP addresses that were probed prior to the intrusion attempt are now at high risk of an additional intrusion. Therefore the blocking rules pro-actively aim to protect them from the relevant threats.

The following emphasizes the major benefits that pro-active blocking rules provide:

- ▼ The rules effectively prevent further intrusion (exploits) attempts, including for those for which exploits do not reside in the known attack signature database.
- ▼ The preventive countermeasures accurately reflect the tangible



**Figure 3: Expert System Threat Level Decision**



**Figure 4: Expert System Blocking Rule Generation**

threat level that each detected attack imposes on the protected network, thus improving the relevancy and accuracy of the prevention rules.


- ▼ Precludes the need for deep-packet inspection (in order to detect exploits). This operation consumes significant CPU resources. The IPS can safely ignore traffic that is associated with the implemented blocking rules and therefore improve IPS performance.

Note that constructing blocking rules is a process that needs to be conducted with great caution, since imprecise blocking rules result in the blocking of legitimate packets. As such, when correlation scores are low, the best strategy is to block only the ongoing detected intrusion and not to implement proactive blocking rules until the correlation process results in higher scores.

## Conclusions

The next stage in the evolution of Intrusion Detection and Prevention technology is the introduction of correlation factors that will take advantage of information “hidden” within signature-based and

behavior-based IPS engines, in order to create a more powerful, more accurate IPS solution.

These next-generation products will apply attack information detected by each of the existing detection engines and will produce a synthesized, unified intrusion prevention instrument, thus increasing the accuracy and performance of tomorrow's intrusion prevention systems. At the same time, these products will produce more detailed reports, enabling increasingly effective attack forensics. 

---

*Avi Chesla is Chief Technology Officer at V-Secure Technologies.*

