

# An Introduction To Distributed Intrusion Detection Systems

*Nathan Einwechter* 2002-01-08

## An Introduction To Distributed Intrusion Detection Systems

by *Nathan Einwechter*, Senior Research Scientist

last updated January 8, 2001

---

### What is a dIDS?

A distributed IDS (dIDS) consists of multiple Intrusion Detection Systems (IDS) over a large network, all of which communicate with each other, or with a central server that facilitates advanced network monitoring, incident analysis, and instant attack data. By having these co-operative agents distributed across a network, incident analysts, network operations, and security personnel are able to get a broader view of what is occurring on their network as a whole.

A dIDS also allows a company to efficiently manage its incident analysis resources by centralizing its attack records and by giving the analyst a quick and easy way to spot new trends and patterns and to identify threats to the network across multiple network segments. This article will discuss distributed intrusion detection systems, including the general setup of a dIDS and a fictional case study to demonstrate the distributed analysis abilities. It will also try to give the reader some insight into the benefits of running a dIDS system, from both incident analyst and corporate views.

### Overview

#### The Central Analysis Server

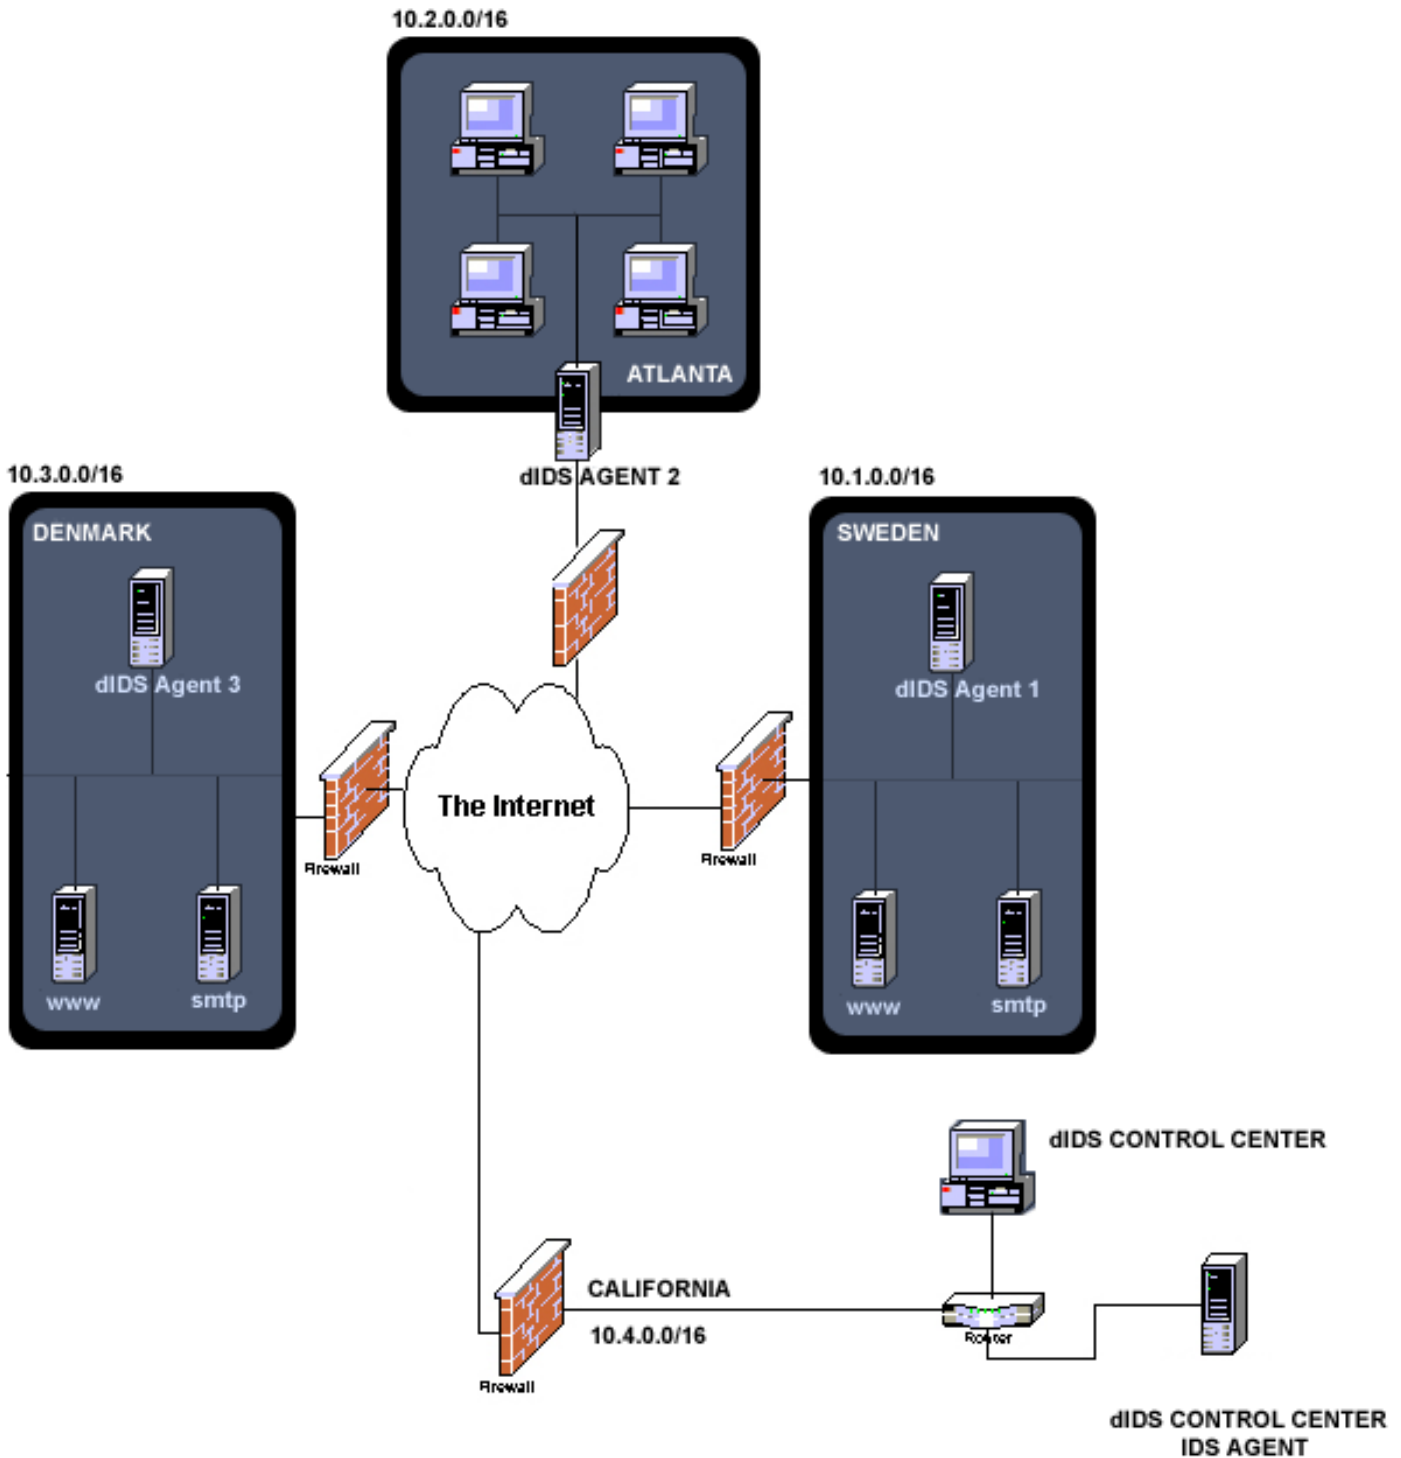
The central analysis server is really the heart and soul of the operation. This server would ideally consist of a database and Web server. This allows the interactive querying of attack data for analysis as well as a useful Web interface to allow the corporate guys upstairs to see the current attack status of your network. It also allows analysts to perform pre-programmed queries, such as attack aggregation, statistics gathering, to identify attack patterns and to perform rudimentary incident analysis, all from a Web interface.

#### The Co-operative Agent Network

The co-operative agent network is one of the most important components of the dIDS. An agent is a piece of software that reports attack information to the central analysis server. The use of multiple agents across a network allows the incident analysis team a broader view of the network than can be achieved with single IDS systems.

Ideally these agents will be located on separate network segments, and geographical locations (See

diagram below.) The agents can also be distributed across multiple physical locations, allowing for a single incident analysis team to view attack data across multiple corporate locations.



Although any IDS could be used on the agent machines, it is highly suggested that [Snort](#) be used. It has been demonstrated, however, that any attack logging system can be incorporated into this agent network. This can range from router attack logs, to ipfw, firewalls, and even Windows personal firewall systems.

## **Attack Aggregation**

Attack aggregation is another core part of the dIDS system. This part of the system is programming logic based on the central server. Aggregation simply refers to the method in which users group or order the information gathered from the agent network. One example of this would be to aggregate information according to attacker IP, putting all attacks from an attacking IP together with other attacks from the same IP. Another example is the aggregation of attack data according to destination (attacked) port, or even by date and time. Uses for aggregation will be explained later in this paper.

## **Advantages of a dIDS**

### **Why a dIDS?**

Due to the greater view the agent allows the analyst to achieve, the dIDS offers the incident analyst many advantages over other single mode IDS systems. One of these advantages is the ability to detect attack patterns across an entire corporate network, with geographic locations separating segments by time zones or even continents. This could allow for the early detection of a well-planned and coordinated attack against the organization in question, which would allow the security people to ensure that targeted systems are secured and offending IPs are disallowed any access. Another proven advantage is to allow early detection of an Internet worm making its way through a corporate network. This information could then be used to identify and clean systems that have been infected by the worm, and prevent further spread of the worm into the network, therefore lowering any financial losses that would otherwise have been incurred.

The second major advantage is that a single analysis team can now do what previously required several incident analysis teams due to physical distance. This obviates the need to pay for distinct incident analysis teams for each separate geographic location of the organization's offices. Another issue that it addresses is attacks from within the corporations network by angry, upset, or bored employees. By tying the central analysis server in with the companies DHCP or RADIUS servers, the incident analysts can track down people launching attacks from within the company, and track what they have attempted to do, as well as provide evidence against the perpetrators.

## **Incident Analysis With dIDS**

Incident analysis using the dIDS system is really what it is all about. This is where all the power, potential, flexibility, and strength of the system as a whole lies. It is the reason why the dIDS was first conceptualized, to allow for advanced analysis of attacks occurring over multiple network segments, and at an advanced level.

### **Analysis Using Aggregation**

Aggregation is the main component used to facilitate this advanced method of analysis across a networks multiple segments. By aggregating similar or related data, the analyst is able to easily see how an attack progressed through the different stages: from active network reconnaissance, to the final attack. It is possible for the incident analyst to see what kind of time frame the attacker was working within and to correlate other attack attempts against the networks to determine if there were multiple co-operative attackers. The most common methods of aggregation are according to attacker IP, destination port, agent ID, date, time, protocol, or attack type.

- Aggregating by attacker IP allows the analyst to view the steps of an attacker's attempt from start to finish across the multiple network segments.
- Aggregating by destination port allows an analyst to view new trends in attack types, and to be able to identify new attack methods, or exploits being used.
- Aggregating by agent ID allows an analyst to see what variety of attacks and attackers have made attempts on the specific network segment the agent is on. Consequently, the analyst can determine if there are multiple attackers working in conjunction, or if there are network segments that are of more interest to attackers than others, thereby giving the security team a list of common targets to work on.
- Aggregating by date and time allows the analyst to view new attack patterns, and to potentially identify new worms or viruses that are only triggered at certain times.
- Aggregating by protocol helps in a purely statistical manner, which could allow an analyst to identify new attacks in particular protocols, or identify protocols on a network segment that should, under no circumstances, be there anyhow.
- Aggregating by attack type also allows for attack pattern matching and to correlate coordinated attacks against multiple network segments.

By utilizing all of these aggregation methods, the analyst is given an unlimited number of different sets of data to correlate against other attacks, detect coordinated distributed attacks, attacks from within their own network, and to detect new exploits and vulnerabilities being deployed by the underground hacking community.

The broad view given by the dIDS system also allows the analyst to ensure a minimum of false positives and false negatives by being able to see beyond a single network segment, into the network as a whole. For example, if the analyst saw that one out of five network segments got seven unrequested ICMP Echo packets, it could be a simple issue of false addressing or improper routing somewhere. However, if the analyst were to see that three separate network segments were reporting seven unrequested ICMP Echo packets, it is much more likely that these packets would be malicious in nature. This would cause the analyst to take note of the activity and perhaps check into the incident further or flag it for review at a later date.

## Analysis Case Study

You come into the office early one morning, boot up your PC, and surf to your Central Analysis server to see what has been going on throughout the night. First thing you do is check incident reports aggregated by the attackers IP. You notice that slews of probes were sent to two internal use IIS Web Servers, located at 172.16.2.106, and 172.16.1.98. These segments' agent Ids are "Main Office." and "Production" The following shows up on the incident report:

**Source IP: 206.219.23.16**

Attack Time	Agent Alias	Target IP	# of Machines Targeted	IP Protocol	Target Port	# of Probes
25 Sep 2001 16:23:45	Main Office	172.16.2.106	1	6	80	12
25 Sep 2001 16:24:01	Production	172.16.1.98	1	6	80	12
25 Sep 2001 16:24:35	Main Office	172.16.2.106	1	6	27374	3
25 Sep 2001 16:24:01	Production	172.16.1.98	1	6	27374	3

Now we'll want to see if any other attacks were attempted on either of these machines. So, we'll aggregate the attack data by the target IP addresses, which are included in the previous report. Now we get two reports.

**Target IP: 172.16.106**

Attack Time	Agent Alias	Target IP	# of Machines Targeted	IP Protocol	Target Port	# of Probes
25 Sep 2001 16:23:02	Main Office	24.26.198.98	1	6	21	3
25 Sep 2001 16:23:21	Main Office	24.26.198.98	1	6	137	5
25 Sep 2001 16:23:45	Main Office	206.219.23.16	1	6	80	12
25 Sep 2001 16:24:35	Main Office	206.219.23.16	1	6	27374	3

**Target IP: 172.16.1.98**

Attack Time	Agent Alias	Target IP	# of Machines Targeted	IP Protocol	Target Port	# of Probes
25 Sep 2001 16:23:14	Production	24.26.198.98	1	6	21	3
25 Sep 2001 16:23:29	Production	24.26.198.98	1	6	137	5

25 Sep 2001 16:24:01	Production	206.219.23.16	1	6	80	12
25 Sep 2001 16:24:46	Production	206.219.23.16	1	6	27374	3

Next we'll combine these two reports by asking the database to give us all attack data with an attacker IP of 24.26.198.98, and 209.219.23.16, against "Production" and "Main Office." We'll also have it sort by date and time:

**Source IP: 24.26.198.98 OR 206.219.23.16**

<b>Attack Time</b>	<b>Agent Alias</b>	<b>Target IP</b>	<b># of Machines Targeted</b>	<b>IP Protocol</b>	<b>Target Port</b>	<b># of Probes</b>
25 Sep 2001 16:23:02	Main Office	24.26.198.98	1	6	21	3
25 Sep 2001 16:23:14	Production	24.26.198.98	1	6	21	3
25 Sep 2001 16:23:21	Main Office	24.26.198.98	1	6	137	5
25 Sep 2001 16:23:29	Production	24.26.198.98	1	6	137	5
25 Sep 2001 16:23:45	Main Office	206.219.23.16	1	6	80	12
25 Sep 2001 16:24:01	Production	206.219.23.16	1	6	80	12
25 Sep 2001 16:24:35	Main Office	206.219.23.16	1	6	27374	3
25 Sep 2001 16:24:46	Production	206.219.23.16	1	6	27374	3

This basically, gives us a step-by-step view of how the attack was carried out by the two attackers. This example is very simplistic, but there have been several demonstrated highly complex attacks that have been identified using this analysis method.

Using these reports, it would be apparent to the analyst that a coordinated attack was attempted against both of these IIS servers. Further analysis can be achieved by viewing the actual IDS logs submitted to the central analysis server, and a decision on network vulnerability to attempted attacks can be performed, as per the GIAC standard for incident analysis reports. It would also be advisable to see the aggregated report on the second attackers IP, to see if any other systems had attack attempts from this system, that were not included in this coordinated attack. It has been demonstrated in the past that this system helps the analyst identify large scale, coordinated attacks against large corporate networks, attacks that might otherwise have gone unnoticed due to communications breakdowns between teams and the inability to

correlate all the data involved across multiple network segments without a system like this.

## Conclusion

The dIDS system gives the analyst a quicker, easier, more efficient method to identify coordinated attacks across multiple network segments, and to trace back the activities of the attackers. The system also, ultimately, saves the corporation whose networks it is deployed on money by reducing the number of Incident Analysts needed, as well as the amount of time required to gather logs from the various IDS systems setup in a large corporate network. By having all of these attack records stored in a single place, it allows the analyst much more flexibility in discovering attack patterns, and other attack issues which may have otherwise gone unnoticed.

As attackers, and attack methods become increasingly complex, the need for a dIDS system in large corporate, and military networks increases drastically. With the increased complexity of these attacks, analysts are leaving themselves open to the problems of communications breakdowns, where one analyst sees a single attack on his segment, and dismisses it as nothing. While several other segments receiving the same attacks in a coordinated manner, their analysts may be dismissing the seriousness of the attack. However, when all the attack data is viewed together, a dramatically different perspective the attack may emerge.

dIDS systems are the next logical level for IDS systems to move to. They are able to be setup with pre-existing architectures and IDS systems, making them even more cost-efficient. It should also be noted that there are currently a few systems in place that fall along the lines of a dIDS. Instead of being based in the corporate environment, focused on in this paper, they are deployed across the entire Internet, which thousands of sensors submitting data to them every day. One such service is [SecurityFocus's ARIS Predictor](#) service, which should be noted due to its large scale demonstration of the use of a dIDS in the reporting of attacks to ISPs, server owners, and their proven ability to identify new worms, and attacks, therefore making the Internet as a whole safer for all.

*Nathan Einwechter is currently a Senior Research Scientist with Fate Research Labs, aswell as a System Developer/ Incident Analyst with myNetWatchman.Com. While working with myNetWatchman, Nathan assisted in the discovery and analysis of the "W32.Leave.Worm" along with SANS, the FBI, and NIPC. The discovery of which was made possible by analysis of data collected by a dIDS system.*

[Privacy Statement](#)

Copyright 2006, SecurityFocus