Global Information Assurance Certification Paper

Intercepting Intrusions With Entercept

Andy Garcia

March 26, 2001

Introduction

When speaking of host-based Intrusion Detection Systems, IDS, one name that should be mentioned is ClickNet Software's entercept 2.x. Entercept's unique design allows it to analyze software calls before the OS executes them. It is able to accomplish this layer of OS protection by working between the kernel and the API layer. If it becomes suspicious of a call, it can block the call from executing and alert the system administrator of the event that triggered it. While many IDSs can detect that the system has been compromised, this product can be configured to prevent it from happening. Also worth noting is the ease of installation. Entercept is comprised of two main components – the console and the agent(s). A single console can accommodate up to 1000 agents. The agent must be installed on each server that entercept is to protect.

System Requirements

It is recommended that the console be installed on a Windows NT 4.0 server (sp 4 or greater), Windows 2000 server or Windows 2000 advanced server. The agent can be run from three platforms: Windows NT (workstation, server, and enterprise), Windows 2000 (professional, server, and advanced server), and Solaris (both MicroSparc and UltraSparc). Note: All testing outlined in this document was performed on a Windows NT 4.0 workstation with sp6a.

Key Features

While most host-based IDS on the market can prevent the system they are monitoring from being compromised from known attacks, entercept can prevent many unknown attacks from compromising the server as well. It is able to accomplish this by not just looking for individual strings, but by also analyzing generic signatures before passing the call on to the kernel. It is also designed to determine how the server should be interfacing with it's applications and alert the console (or block the action) if the application behaves outside of its "norm." This type of application "shielding" allows the agent to protect the server from unknown attacks.

How does it work?

As previously mentioned, there are two primary components – the console and the agent(s). The agents are configured to communicate with the console by pointing to the console's IP address. Traffic between the agent(s) and console is encrypted using asymmetric-key cryptosystems to prevent it from being tampered with while in route. Remember, asymmetric-key cryptosystems requires two keys – a private key and a public key that is created from the private key. When the console is installed, it automatically creates both keys. The installation wizard will prompt for the public key when installing each agent. The most secure way to install it is to put it on a floppy disk and install it on the agent manually.

Once configured, the agent will analyze all calls made to the kernel. If it deems the call to be "normal" it will let it pass through to the kernel and it will execute as always. If, however, the agent is suspicious about the call it can perform one of four actions (depending on how it is configured): ignore, log, prevent, or terminate the process. More detail on these actions will be explored in the next section.

Getting started

So far, this document has discussed how the agent(s) and console communicate in a secure manor and touched on a few key features of the product, so now it's time to jump in and explore how it works.

Installing the console involves little more than kicking it off and answering a few straightforward installation questions. Installing the agent is just as easy – but the static IP address of the console must be known and a copy of it's public key must be on hand. Once the agent is installed, all monitoring and maintenance can be performed from the console. There is an agent configuration utility found in C:\Program Files\Clicknet\enterceptagent\Config.exe that can be run from the agent to perform minor configuration changes, but for the most part, all work can be done remotely from the console.

The console has seven major components that can be accessed from their corresponding icons in the left vertical pane. These components are: security event monitor, system event monitor, agent management, policy management, exception management, security level modifiers, and agent version management. Once configured, a vast majority of work and monitoring will be done from the security event monitor and the system event monitor.

The security event monitor lists security events reported by the various agents that it is responsible for. It describes the following characteristics of each event captured:

Event Characteristic	Example
Source:	Server001
Event Name:	SAM Permission Modification
Recording Time	3/20/01 11:48:16 PM
User:	NT Authority\System
Process:	Winlogon.exe
Reaction:	Prevent
Note:	

Additional details can be found by looking at the event's properties. Each event has three tabs associated with it's properties: details, description, and advanced details.

There are several built-in filters in entercept that allow for filtering of events displayed on the monitor. The first group of objects that can be filtered is the security level associated with that event - high, medium, low, information, and system. The second group of objects that can be filtered aid in tracking down a particular event or attack. This group

can filter by security level, source, event name, day, user, process, or by any combination of those listed above. To prevent false-positives from saturating the security event monitor, entercept has an exceptions feature. Exceptions are discussed in more details later, but they are created from within the security event monitor. When entercept is first setup, count on having a tremendous amount of false-positives reported – this is characteristic of any new IDS or monitoring device that is added to a network.

The system event monitor looks and feels much the same as the security event monitor. It's job, however, is to monitor entercept's system activity – excluding security events. It displays events such as services starting/stopping, console administrators logging on/off, agents starting/stopping, etc. It can be filtered much the same as the security event monitor.

The agent management module allows the administrator to create and manage agent groups. By default it ships with two built-in groups – all agents and new agents. The agent groups are listed vertically in the left pane and the agents (servers being monitored) are listed horizontally in the right pane. This module allows the agents to be segmented into smaller, more manageable, groups. These groups are created and defined by the administrator, so they can be functional, geographical, or whatever makes since for that environment. At a quick glance, the key characteristics of each agent can be seen from the console - it's name, current state (active, no connection, etc.), and requested state (warning or protection mode).

The policy management module allows the administrator to create and manage policies that define how entercept reacts to an attack or suspicious event and who the policy pertains to. Additionally, it defines who gets contacted; by what means they are contacted, and what reaction entercept should take when an event is triggered.

There are four event levels that are configured in the properties of each policy. These event levels are high, medium, low, and info. For each event level there is a reaction that can be taken – ignore, log, prevent, or terminate process. (Note: later in this document, under the section entitled '*Going for a test drive*' more details on these four reactions will be explored). For each of the four event levels, characteristics such as who gets contacted and how they are contacted are defined. The four notification methods are: email, pager, kicking off a spawn process, or kicking off an SNMP trap.

The exceptions management console contains a list of events that exceptions have been made for. As mentioned earlier, exceptions are defined in the security event monitor and\or the system event monitor by right-clicking on an event and selecting *'Create Exception*.' When first installed, the console and agents should be configured to just *log* all events. They should remain this way until that administrator has a chance to create exceptions for all the events that shouldn't get reported to the security and system management consoles.

Each exception created has three tabs in it's properties that allow the administrator to configure and manage the exception with ease. The tabs are: details, event description,

and advanced details. The details tab is where the rule is named and most importantly where it is defined. Exceptions can be defined by agent (either an individual server or all servers), user (a particular user or all users), or the process that triggers the event. The event description tab describes the event that caused the flag and gives a brief description of why this may be of concern. The advanced details tab identifies the server or workstation that triggered the event and describes the exact registry key or file(s) that caused the event to be triggered. Filtering in this module is much the same as in other modules; it can be accomplished by the event name, the agent, by the user, process, or any combination of those listed above.

The security level modifiers module is important because it gives the administrator the ability to change the trigger levels from the default settings that ship with the agent to levels that are more appropriate. For instance, the administrator might list a certain event as a medium concern while the agent may list it as high by default. This module is where these modifications can be made. Before moving the agents to a production environment, (i.e. switching from warning mode to protection mode), it is highly recommend that the administrator get a list of every agent and verifies that their default settings are in alignment with the company security policy. For those that aren't in alignment, change their trigger levels manually and document the changes. To change them, select the *new level modifier* icon in the toolbar (or right-click in the right pane and select *new…*) to bring up a list of all events that could trigger an alert. Each event is displayed with their signature name, default security setting, the platform the agent pertains to, and the application that is associated with the event. Highlight the event that needs to be changed and select the *change level* button. Entercept will then give the opportunity to document the change and assign a new security level to it – disabled, info, low, medium, or high. Additionally, it can be configured to apply to all the agent groups or to specific groups from here.

Finally, there's the agent version management module. When first installed, there will only be the default agent displayed. Periodically, new agents are be released by the vendor (which can be configured to be pulled down automatically or manually). A list of all agents available to the system will be displayed here. This module is important because it allows the administrator to manage the status and distribution of the agent(s). For instance, before releasing a new agent into the production environment, the agent can be configured to a testing status and assigned it to a test group. This enables the events to just get logged, thus giving the administrator the opportunity to determine which events need exceptions and which should have their trigger levels reclassified before moving to a production environment.

There are two other modules to be aware of. They are the reports module and the options module. The reports module is found in the toolbar and is used for running various reports defined by the administrator. It is similar to most reporting tools on the market so a detailed explanation is not necessary. The options module is found in the toolbar from the tools pull-down. This is where key characteristics of the program are configured. These characteristics include: the size of the database, number of security events retained, number of system events retained, status of agents after deployment, domain settings,

email configuration, scheduling of agent updates, password policy, and communication between the console and agent(s).

Going for a test drive

This document has demonstrated that entercept is easy to install and configure, but does it really work as advertised? To test, I ran multiple scenarios against entercept to test it's capabilities and limitations. The testing included everything from creating a new admin account, creating files with multiple extensions, password tests, and running Netcat and L0phtCrack against it. I ran each scenario multiple times, each time changing the configuration of the policy to see how it reacted. In almost all cases, entercept worked as I anticipated it would and performed the way it is documented to work. Below is a snippet of some of the testing I performed and the resulting actions. All testing documented below was performed on a Windows NT workstation, sp6a.

Scenario 1: Creating a backup of the registry and SAM using rdisk /s.

Configuration 1: Current state: On-warning\ Requested state: On-warning\ Reaction state: Log
Result 1:
- I was able to complete the process successfully.
- The security console generated three events in security monitor.
    - Low, Repair Directory Access, [Log (warning)].

Configuration 2: On-warning\On-Protecting**Prevent**
Result 2:
- Rdisk launched, but dropped out without running the request.
- The security console logged the event.
    - Low, Repair Directory Access, [**Prevent**].

Configuration 3: Left configurations defined above and created an exception.
Result 3:
- I was able to complete the process successfully.
- No events were displayed in the security console.

Scenario 2: Creating a new user account with admin privileges.

Configuration 1: On-warning\On-warning\Log
Result 1:
- I was able to create the account successfully.
- The security console generated three events linked to the permissions defined on the account.

Configuration 2: On-warning\On-Protecting**Prevent**
Result 2:
- User Manager displayed an 'Access denied' error when attempting to create the account.
- All actions were logged in the security monitor and listed as [**Prevent**].

Conclusion

When considering a new host-based IDS entercept 2.x should be looked into. It is easy to install and configure and doesn't require a high-end security analyst to monitor and maintain. It also allows the administrator to remotely monitor several agents simultaneously. The agent analyzes each call before it allowing it to execute, thus enabling entercept to prevent the compromise from occurring - rather than just reporting to the administrator that an attack has occurred.

Works Cited

Andress, Mandy "entercept essential to corporate security plans" InfoWorld.com URL: http://www.infoworld.com/articles/pi/xml/00/03/13/000313pientercept.xml (03/10/2000)

"entercept" SC Magazine.com URL: http://www.scmagazine.com/scmagazine/standalone/clicknet/clicknet.htm

entercept Security Technologies web site URL: http://www.entercept.com/products/entercept/ (2001)

entercept – User Guide, Copyright 2000, ClickNet Software Corporation, Printed 07/00 (CNCS-UG1-10001)

Hulme, George, "ClickNet's entercept 2.0 Stops Attacks" Information Week.com URL: http://www.informationweek.com/story/IWK20000814S0004, (08/14/2000)

Jackson, William, "New software protects Microsoft Web server app" Government Computer News URL: http://www.gcn.com/vol19_no24/com/2731-1.html (08\21\2000)

Stevens, Alan "First Looks: Network Edition – entercept 2.0" PC Magazine URL: http://www.zdnet.co.uk/pcmag/ne/2001/02/05.html (02/2001)