Tivoli

Tivoli™ Cross-Site for Security User’s Guide

Version 1.1

October 20, 1999

Tivoli Cross-Site for Security User's Guide (March 1999)
Copyright Notice

Trademarks

The following product names are trademarks of Tivoli Systems or IBM Corporation: AIX, IBM, OS/2, RS/6000, NetView, Tivoli, and Cross-Site.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

ActionMedia, LANDesk, MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

Sun™ workstation, Solaris Operating Environment™ software, SPARCstation™ 20 system, Java™ technology, and Sun RPC are trademarks of Sun Microsystems, Inc.

VitalAgent™ is a trademark of INS.

JClass LiveTable, JClass BWT, JClass Field, and JClass Chart are trademarks of KL Group Inc. Copyright © 1997-1999 KL Group Inc. All rights reserved.

VeriSign is a trademark of VeriSign, Inc. Copyright © 1998 VeriSign, Inc. All rights reserved.

Contains RSA technology. Copyright © 1991 RSA Laboratories, a division of RSA Data Security, Inc. All rights reserved.

Copyright © 1990, 1993, 1994, 1995, 1996. The Regents of the University of California. All rights reserved.

Castanet and the Powered by Marimba logo are trademarks of Marimba, Inc. Copyright © 1996-1999 Marimba, Inc. All rights reserved.

Powered by marimba®

Other company, product, and service names mentioned in this document may be trademarks or servicemarks of others.

Notice

References in this publication to Tivoli Systems or IBM products, programs, or services do not imply that they will be available in all countries in which Tivoli Systems or IBM operates. Any reference to these products, programs, or services is not intended to imply that only Tivoli Systems or IBM products, programs, or services can be used. Subject to Tivoli Systems' or IBM's valid intellectual property or other legally protectable right, any functionally equivalent product, program, or service can be used instead of the referenced product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by Tivoli Systems or IBM, are the responsibility of the user.

Tivoli Systems or IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, New York 10504-1785, U.S.A.

Contents

Preface

The Cross-Site for Security application provides intrusion detection for systems in your environment that might be vulnerable to attack. The Cross-Site for Security agent that detects intrusion attempts should be deployed wherever the administrative domain connects to the Internet. The Cross-Site for Security agent can do the following:

- Detect scans and floods
- Monitor IP traffic
- Monitor port services
- Detect DNS, mount service, and network file system (NFS) requests and replies
- Detect portmapper service request and reply dumps
- Detect RStatd calls
- Detect requests for specific map names and file names
- Detect SMB-based attacks on PC file servers
- Detect Internet control message protocol (ICMP) attacks

All Cross-Site applications leverage a set of core services that is provided with the Cross-Site server and agents. These services enable all communications, tasks, and security for Cross-Site. The Cross-Site services include:

- Authorization and authentication
- The event service
- The task manager
- The policy manager
- The channel manager
- The report generator

This preface highlights new features of Cross-Site for Security version 1.1. It identifies the audience for this guide, the related and prerequisite documentation, the typeface conventions used in the Cross-Site documentation, the Cross-Site icons, and instructions for contacting customer support.

What’s New in This Guide

This guide documents the new features of Cross-Site for Security version 1.1. In addition to providing procedures related to Cross-Site’s AIX and DB2 support, the guide describes new features including flood squelching, extending the functionality of the IGNORE command, token ring support, and DNS address resolve alerts. Changes to the Cross-Site graphical user interface are described. A new appendix, “Signatures Detected by Cross-Site for Security”, describes many of the signatures that the Cross-Site for Security agent can detect while monitoring your network, and provides preventive measures to protect against intrusion attempts. Finally, the guide includes updated versions of the **ids.msg**, **ids.cfg**, and **ids.rules** configuration files.

Who Should Read This Guide

The target audience for this guide is system and web administrators who are familiar with maintaining operating systems, administering web servers, maintaining databases, and general information technology (IT) procedures. Specifically, users of the guide should have some knowledge of the following software:

- Solaris, NT, and AIX operating systems
- The Netscape Enterprise Server, including the administration server, and web protocols such as HTTP and SSL; for secure web servers, you must also be familiar with digital certificates
- Oracle or DB2
- The Java environment

Users of the guide should also have some knowledge of the Tivoli Enterprise Console and Tivoli NetView products if they intend to use the Tivoli Enterprise integration. Cross-Site provides integration with these products, for event forwarding.

Prerequisite and Related Documents

The following documents are related to the Cross-Site product suite and are considered prerequisite reading before using Cross-Site:

- *Tivoli Cross-Site for Security Release Notes*
  The release notes provide the most up-to-date information about known defects and their workarounds. Be sure to read this document before installing and using Cross-Site.
- *Tivoli Cross-Site for Security Installation Guide*
  This guide provides prerequisite information and instructions for installing Cross-Site. You might be interested in this guide after installing Cross-Site for

information such as connectivity commands and uninstallation procedures. This guide was shipped with the product. You can access the installation guide on the product CD-ROM or download it from the following URL:

http://www.cross-site.com/support/docs

- *Cross-Site Integration with Tivoli Enterprise Applications User's Guide*
  This guide describes how to use the features provided with Cross-Site for integration with Tivoli Enterprise applications. Specifically, Cross-Site provides integration with Tivoli Software Distribution, Tivoli User Administration, Tivoli Enterprise Console, and Tivoli NetView. You can download this user's guide from the following URL:

  http://www.cross-site.com/support/docs

Conventions Used in This Guide

This guide uses several typeface conventions for special terms and actions. These conventions have the following meaning:

Bold Command names and arguments, keywords, file names, World Wide Web addresses, or other information that you must use literally appear in bold. Names of windows, dialogs, and other controls also appear in bold.

Italics Variables and values that you must provide appear in italics. Italics are also used to emphasize words or phrases.

New terms appear in italics when defined in the text.

`Monospace` Code examples and commands appear in a monospace font.

Cross-Site Icons

The following icons represent various Cross-Site resources:

Identifies the executable (EXE) files on NT platforms for the Cross-Site agent and console.

Represents the Cross-Site suite; in particular, this icon represents a global view of all resources on the Cross-Site console.

Represents the Cross-Site for Security application.

Represents Cross-Site's administrative services. For example, you can administer users, roles, events, and the event log.

? Represents the Cross-Site help system. The help system provides conceptual, procedural, and reference information about the Cross-Site services and applications.

Contacting Customer Support

The Tivoli Cross-Site support team is committed to providing you with the best possible service. As a leading provider of Internet-based solutions, Tivoli believes that the Web is the ideal venue for service. Therefore, you can obtain answers to critical questions on the Cross-Site web site (**http://www.cross-site.com/support**). This support site contains the following information:

- Hot issues, which are immediate issues that are the most recently identified and the most frequently requested.
- A list of frequently asked questions (FAQs), which provides answers to basic questions that address issues that are important, but not necessarily urgent.
- Updated versions of the Cross-Site documentation.

Use the **Ask Support** form to submit a request for assistance. The URL for the form is as follows:

http://www.cross-site.com/support/asksupp/submitpmr

Tivoli is very interested in hearing from you about your experience with Cross-Site products and documentation.

For telephone support inside the United States, contact Tivoli Cross-Site Support by calling 1-800-TIVOLI8. For support outside the United States, refer to your Customer Support Handbook for phone numbers in your country. The handbook is available on-line at the following URL:

http://www.support.tivoli.com

1

Cross-Site for Security Overview

This section of the guide begins by introducing the basic Cross-Site concepts, such as the Cross-Site server, the Cross-Site domain, and Cross-Site resources. It is important that you understand these concepts before you proceed. After you are familiar with them, you should have little trouble understanding Cross-Site for Security concepts. Understanding the basics enables you to build on them and easily accomplish Cross-Site tasks.

Cross-Site for Security enables you to monitor your network for possible intrusions. This is explained here in detail, and usage scenarios are provided as examples of how you might use Cross-Site for Security, depending on your network and firewall setup. Other Security concepts discussed here include the agents and the configuration files that you will use to accomplish your goal, and the ways to report on intrusion attempts and attacks.

Finally, several documents are provided to help you get started using Cross-Site for Security. These documents provide an overview of the steps you must perform to set up and use Cross-Site for Security. Refer to later sections for detailed information.

Cross-Site Concepts

This section introduces the basic terms and concepts that you need to know to perform Cross-Site operations. Cross-Site collectively provides functionality from four different areas:

- The Cross-Site core services
- The Cross-Site for Security application, or Security
- The Cross-Site for Availability application, or Availability
- The Cross-Site for Deployment application, or Deployment

The Cross-Site management server provides services upon installation. These *core services* include the event service, resource management, security (authorization and authentication), report generation, and administration.

Cross-Site applications use the core services to provide specific Internet management functions.

- *Cross-Site for Security* enables you to monitor network traffic and detect attacks and intrusion attempts. Cross-Site for Security monitors traffic in your “demilitarized zone” (DMZ), the area between two firewalls that insulates your intranet from the Internet, and on your internal network.
- *Cross-Site for Availability* enables you to gather information about the reliability of web sites that you access regularly. It provides a way for you to measure the effectiveness of the business you conduct over the Internet by alerting you when sites are performing below standard or servers are down.
- *Cross-Site for Deployment* provides a secure way to distribute applications and information to end users on an intranet or the Internet.

The elements of Cross-Site that enable the core services and applications to provide functionality are the following:

- Cross-Site management server, also referred to as the Cross-Site server or management server
- Cross-Site management repository
- Cross-Site console
- Cross-Site for Security agent, also called the Security agent
- Cross-Site application agents, often referred to by the name of the application that is installed on the agent. Currently Cross-Site features Deployment and Availability agents.

All of these elements comprise a *domain,* also referred to as an *administrative domain*. Because the *Cross-Site server* hosts all of the resources and provides services to all of the Cross-Site elements, it represents the domain. There is also the concept of a *foreign domain*, which is a domain associated with a Cross-Site management server other than the local server.

The Cross-Site server stores its information in the management repository. The *management repository* is based on a relational database that stores all Cross-Site data and information. You create the management repository in your database during Cross-Site installation.

The *Cross-Site console* acts as a window into the Cross-Site server and management repository. It is the user interface (UI) through which you can view all Cross-Site resources, tasks, events, and policy; you use the console to manage Cross-Site. More than one console can interface with a Cross-Site server, and you do not need to run the console on the server.

You can monitor all of the Cross-Site agents from the Cross-Site console.

Security agents reside on a network server, continuously monitoring network traffic for suspicious activity. When such activity is detected, the agents determine the severity and, if required, notify the Cross-Site server. Security agents have no UI.

Availability agents should be installed on client computers and provide the Cross-Site agent desktop for use by end users. The agents provide end users with a way to see Availability configurations and data.

Deployment agents should be installed on client computers and provide the Cross-Site agent desktop for use by end users. The agents provide end users with a way to see Deployment configurations and data.

The following is an illustration of a Cross-Site domain. This illustration highlights how each component communicates with the others.

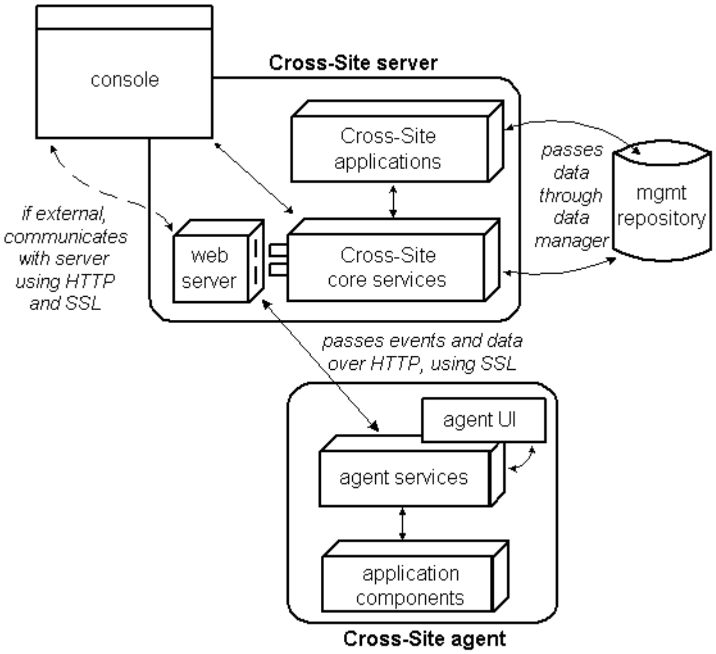

The console can run either on the Cross-Site server or on another computer external to the Cross-Site server. However, the computer on which you are running the console must be able to connect to the Cross-Site server using HTTP and Secure Sockets Layer (SSL). Agents are also connected to the server over the Internet and use HTTP and SSL to send information and download channels.

The functionality provided by the Cross-Site core services and applications in a Cross-Site domain enables you to manage your e-commerce environment. The Cross-Site console displays the following elements of your e-commerce environment as follows:

- Resources
- Policy
- Tasks
- Events

Resources are items that you can manage, including Cross-Site agents, users, sites, channels, and collections. *Policy* specifies how a resource or task (e.g., site scan policy) behaves. It enables you to create restrictions or guidelines for resources. For example, policy controls resource properties and upload intervals for Availability agents.

Users Enable you to manage who has access to Cross-Site. A user can represent an end user of an agent or an administrator who uses the Cross-Site console. By creating users you can limit, through a log-on ID and password, what Cross-Site information individuals may view. You can also specify which individuals can make changes to the management repository.

Availability sites
Represent web sites that you want Cross-Site to scan. When you create a site resource, you must enter the Web site’s URL and assign an Availability policy to it that defines parameters for the scan.

Deployment channels
Hold, or host, information that you want Deployment agents to download. When you create a channel, you must also apply Deployment policy. Policy for channels specifies how often agents check for updates, as well as other parameters.

Collections Containers that organize resources. A collection is similar to a folder in a Windows environment. However, you can apply policy to a Cross-Site collection. When you apply a policy to a collection, you apply policy to all the resources in that collection.

Tasks are actions that run on a resource. The parameters of the action are set through policy, but the action is carried out by a task. Tasks are used to retrieve information from Security and Availability agents and publish Deployment channels.

Events are alerts, error messages, or status messages that are displayed on the Cross-Site console. They inform you when services and applications are functioning, and notify you of policy violations. Status events are messages that give you statistics and updates regarding regular Cross-Site system functions. For example, a status event might indicate that a task has completed.

Alert events are error messages that warn you of possible problems. For example, if an Availability site violates policy by taking too long to load into a browser, an alert event is generated. Critical events are alert messages that indicate serious policy violations. For example, the Security agent sends a critical event to the console if it detects a threatening ping and the server is in danger of failing.

All of the concepts discussed in this section are described in greater detail in later sections. See each application’s Getting Started section for tips on how to use Cross-Site elements.

Cross-Site for Security Overview

Cross-Site for Security is an intrusion detection system that monitors networks to detect external and internal attacks. Cross-Site for Security works with the Cross-Site framework services to bring overall protection and management to on-going Internet commerce relationships.

Cross-Site for Security includes the following benefits:

- Real-time intrusion detection that alerts the Cross-Site administrator of potential attacks.
- Configurable policy that enables you to set different policy for agents in your DMZ, the area between two firewalls that insulates your internal network from the Internet, and agents on your intranet.
- Online modification of Security agent policy that enables you to respond to changing environments quickly.
- Integration with Tivoli's Enterprise applications that enables you to augment your Tivoli enterprise management system.
- Token-ring support for Windows NT and AIX systems

The following list outlines the types of intrusions that Cross-Site for Security detects:

- Signature, or pattern, detection
- Flood detection
- Network-based attacks
- Windows network attacks
- Remote procedure attacks (SUN/ONC RPC)

- Service exploitations
- Unauthorized network traffic
- Suspicious activity

Cross-Site for Security guards your network from unauthorized activities by using the Cross-Site for Security agent and Cross-Site management server. When an agent detects a critical attack, it sends an encrypted event using Secure Sockets Layer (SSL) to the Cross-Site server, which immediately logs the information and responds, according to how it is configured. You can configure the management server to send an alert to the Cross-Site console, post an e-mail message to the administrator, page an on-call administrator, or execute a predefined script.

This section lists the available protocols that Cross-Site for Security supports, introduces how Security intrusion detection works, and shows how Security fits into the suite of Cross-Site products. For information about configuring Cross-Site for Security agents, see “Security Configurations” on page 17.

Supported Protocols

This section explains the network protocols that Cross-Site for Security supports. The protocols are divided into the four standard protocol levels. Arrows show the possible information flow from one level to the next.

In the Open System Interconnection (OSI) model for data communication, Cross-Site monitors the flow of information from the Data Link layer through to the Application layer. The following diagram shows only a sample of protocols in the OSI. For more information about the protocols that Cross-Site supports, see the “Cross-Site for Security Features” on page 109.

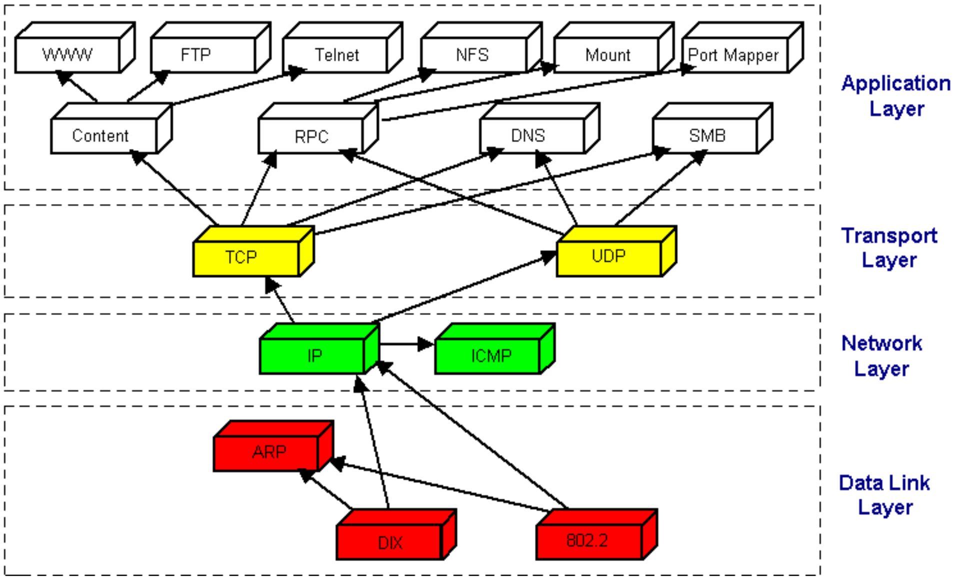

Information arrives through the Data Link layer, which is the interface to the hardware. At this layer, Cross-Site collects the packets of information, reads them, and rearranges them at the Network, Transport, and Application layers. The diagram above illustrates how packets are filtered across these layers.

Cross-Site for Security monitors the protocols in each layer. Intrusion patterns, or signatures, are detected within these modules for the particular services they support.

The next section describes how Cross-Site for Security works and gives details about how Security behaves when it detects suspicious activities.

Monitoring Traffic

The Cross-Site for Security agent is an intelligent network sniffer. It continually monitors the packets coming into and out of each computer on the network. The Security agent filters these packets, looking for various signatures that represent suspicious activity. These signatures can indicate attacks on the network, referred to as *incidents*. The way the Security agent filters packets is detailed in “Security Policy” on page 57.

The Security agent runs as a daemon on UNIX and AIX, and as an NT service on Windows NT. Cross-Site for Security is configured to start automatically when the system boots. It remains resident and runs on the system in the background, regardless of whether or not a user is logged in.

When an incident is detected, the Security agent determines the severity of the incident and decides whether to immediately notify the management server or log an alert to a local file. This log file is periodically uploaded to the management server. For more information, see “Security Events” on page 75.

In addition to detecting incidents and generating alerts, the Security agent regularly contacts the management server to let it know that the agent is “alive” and running. This method of communication is called sending a *heartbeat*. If the Security agent does not send a heartbeat to the management server, the server knows that something is wrong and sends an “Agent not responding” event to the Cross-Site console. You can configure the interval by which each Security agent sends heartbeats in the **ids.cfg** file. For information on how to configure the **ids.cfg** file, see “Configuring Scan and Flood Parameters” on page 63.

When the Security agent sends a heartbeat to the management server, the server responds to the communication. Additionally, the Cross-Site server notifies the Security agent of updated configuration information, new signatures, and upload schedules. The Security agent automatically downloads and installs these updates.

For information about where to install the Cross-Site server and Security agents, see “Determining the Location of Servers and Agents” on page 17.

The server component of Cross-Site for Security acts as a central collection point for the data gathered by each Security agent. It processes this data and saves it in the *management repository*. You can use the Cross-Site console to view a list of events created from critical agent alerts. To view other alerts, you must generate reports from the **Resources** explorer.

How Cross-Site for Security Fits into the Cross-Site Suite

Cross-Site for Security leverages the Cross-Site services and console, just as other Cross-Site applications do. To view Security-related resources, policy, tasks, and events, select the Global or Security view icon on the Cross-Site console.

Resources are configurable objects that Cross-Site manages. Security provides the agent resource. A Security agent is a network sniffer that you install either on a production server or on a dedicated UNIX system. The Security agent monitors all the packets that pass through the network and sends alerts if it detects any critical activity. For more information about agent resources, see “The Security Resource” on page 51.

Edit the Security agent’s properties from the **Resources** explorer. Configure the Security agent by creating and applying *policy* in the following ways:

- Specify the signatures to which you want the agent to respond.
- Choose the priority level of each type of intrusion.
- Create or edit alert messages.

You can use policy to configure agents in different ways. For example, the environments on your intranet and Internet are different. You might want agents outside of your intranet to respond differently to certain signatures than agents inside your intranet. You configure the Security agent resource by modifying the policy files. For more information about configuring Security policy, see “Security Policy” on page 57.

Security Configurations

Cross-Site for Security is a flexible system that can be configured according to your business requirements. Cross-Site for Security configuration options support a wide variety of firewall and access policies.

This section describes the main decisions you have to make before installing Cross-Site for Security:

- Where should I install the Cross-Site management server: in my DMZ, the network space that buffers the internal network from the Internet, or on my internal network?
- How many Security agents do I need: one on every server or one for an entire network?
- Where should I install the Security agents: do I install them on a UNIX system connected to my router or directly on my server boxes?

This section discusses three general scenarios that depict how you might configure your Cross-Site for Security system. However, these scenarios are by no means exhaustive. The flexibility of Cross-Site for Security enables you to create the custom configuration that works best for you.

Determining the Location of Servers and Agents

Positioning management servers and Security agents depends on the following basic factors:

- Whether you want to monitor your DMZ, internal network, or both
- Where your firewalls are located

- What kind of network you have: switched or unswitched ethernet or token ring
- Whether you want to install the Security agents on existing servers or on dedicated UNIX machines

  Note: *If you use a dedicated computer for a heterogeneous network, you will want to use a UNIX system because an NT system cannot monitor UNIX-specific RPC traffic.*

The next section describes the factors that you must consider when deciding where to install your Cross-Site management server.

Where Should You Position the Cross-Site Management Server?

The Cross-Site management server can be installed either in the DMZ or behind the internal firewall. The most important item to remember when selecting a location is that the Security agents have to be able to connect to the management server.

Security agents communicate with the management server using secure HTTP. For communication to take place between the agents and the management server, your firewall settings must be configured to allow the following communication:

- The management server must be configured to accept HTTPS traffic on the destination port. Typically HTTPS is set to port 443.
- Currently, the Security agent machine SRC port is chosen dynamically by the operating system to a port number greater than 1023 and less than 65536.

Typically, browsers inside the internal firewall may access management servers in the DMZ, as described above. However, if your network does not permit internal machines to access the Internet, you might need to make alterations to your firewall configuration.

Because the management server contains potentially sensitive security information, you might feel more comfortable placing it behind the internal firewall. Positioning the management server on your intranet provides added security and enables you to monitor internal machines.

You may also monitor the DMZ with a management server positioned inside the internal firewall. As with the other configuration, a standard web browser must be able to access the management server from the Security agent machine residing in the DMZ.

To test whether your agents can access the management server. Type the following URL into a browser on the agent machine:

```
protocol://server_name:port/servlet/CrossSiteServlet
```

If the Cross-Site server is a secure server (running over SSL), specify **https** for *protocol*. The *server_name* and *port* variables are the values specified when the Cross-Site server was installed.

If the server is running, the URL displays a page containing the following message:

Cross-Site Management Server Status: ALIVE
Version: *x*

If you can access this URL, then the agents can access the management server.

After you have decided where to install your management server, you must think about Security agents. The next section describes where you should install agents and the factors to consider as you decide how many agents you need.

Where Should You Position Security Agents?

You can install Security agents on any server to monitor traffic to that machine, or on a dedicated UNIX machine to monitor large networks. If you are using a dedicated UNIX machine to host the Security agent, you can place the UNIX machine on the network in the following areas:

- Attached to the network you want to monitor (unswitched ethernet)
- Connected to a network router with the packet-forwarding feature turned on (switched ethernet)

If you have an unswitched ethernet, you also have the option of installing a Security agent on any server connected to the network. That agent can scan all of the traffic going to each server on its network. By installing an agent on a server in the DMZ and one on the intranet, your entire system is monitored for intrusion attempts.

If you have a switched ethernet, the preceding scenario will not work because a Security agent on one server cannot detect traffic going to another server. For a switched ethernet, you can install a Security agent on all servers in your network. This approach has the advantage of fault tolerance. However, you are unable to use wide-scan detection because of the nature of switched ethernet.

Another way to monitor a switched ethernet is to connect a dedicated UNIX machine to the router serving the networks you want to monitor. The router should be set to VLAN or “mirroring” mode. This setting enables monitoring of all network traffic.

If you have a token-ring network, a single agent can see all packets coming into the network.

Basic Security Configurations

The following figures illustrate sample configurations for Cross-Site for Security. Your final configuration is likely to be slightly different than these samples.

Management Server in the DMZ

The following graphic illustrates a configuration in which the Cross-Site server resides inside the DMZ. From this location, the management server can monitor traffic moving in the DMZ and the internal traffic on the intranet.

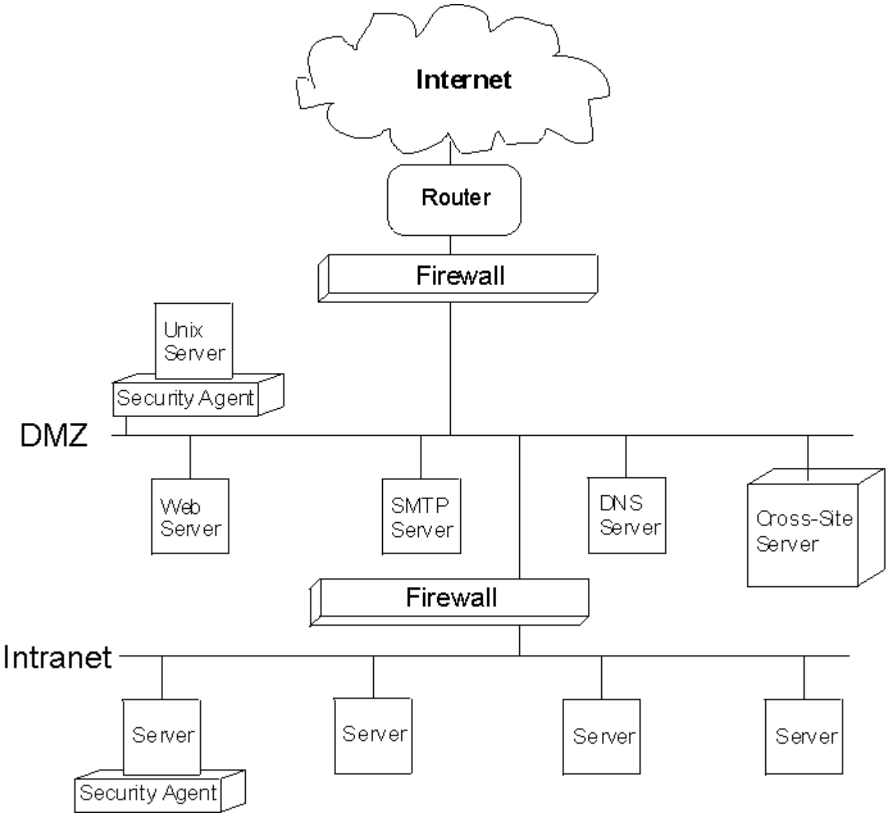

This example assumes a configuration supporting an unswitched ethernet and has a Security agent on a dedicated UNIX machine attached to the network segment. This computer is installed inside the DMZ to monitor Internet traffic. Additionally, a Security agent is installed on a production server inside the internal firewall to monitor intranet traffic.

Management Server on Intranet

To provide increased security for the management server, you can position it inside the internal firewall. To check the connection in the DMZ, make sure a browser in the DMZ can access the management server inside the intranet by accessing the following URL:

```
protocol://server_name:port/servlet/CrossSiteServlet
```

If the Cross-Site server is a secure server (running over SSL), specify **https** for *protocol*. The *server_name* and *port* variables are the values specified when the Cross-Site server was installed. If the server is running, the URL displays a page containing the following message:

Cross-Site Management Server Status: ALIVE
Version: ***x***

If a web browser cannot access the management server, you might need to alter the firewall configuration. See “Where Should You Position the Cross-Site Management Server?” on page 18 for more information about firewall settings. The following graphic shows a scenario in which the Cross-Site server is located on the intranet.

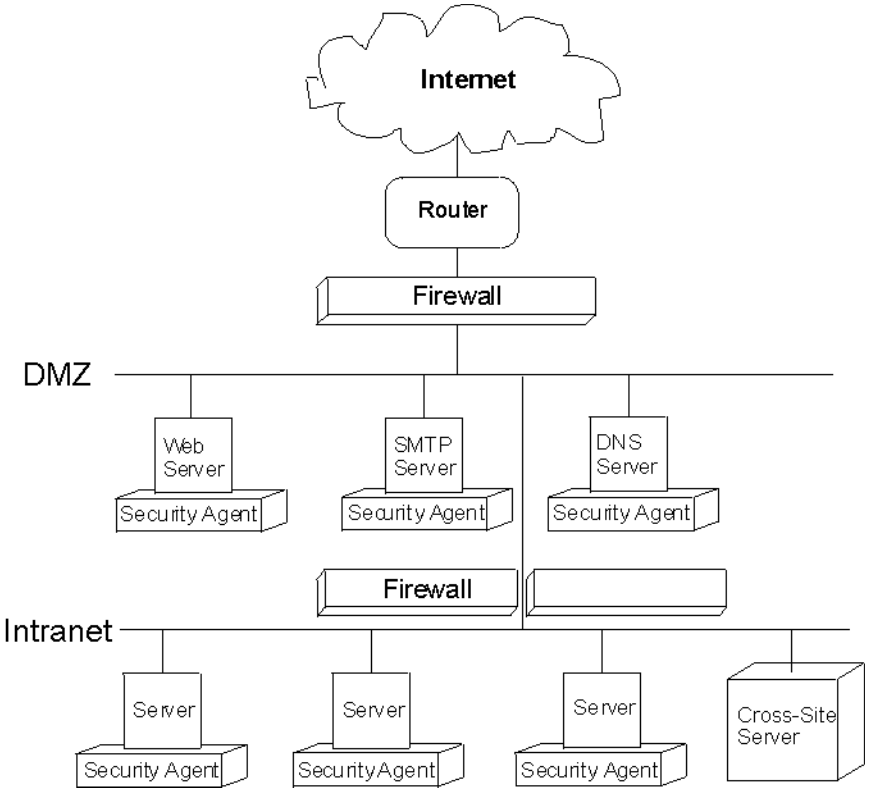

This example assumes a configuration supporting switched ethernet, with the Cross-Site management server inside the internal firewall and Security agents on each server. Installing Security agents on every server is one way to configure Cross-Site for Security on a switched ethernet. You can also connect Security agents to network routers in order to monitor all incoming traffic.

An advantage to placing the Cross-Site management server in the DMZ is that, from there, the management server can support Security agents for resources located outside corporate boundaries. For example, you might have a server located at your Internet service provider (ISP). This server could be connected to the Cross-Site server over the Internet. The ability to support sites connected over the Internet enables Cross-Site for Security to scale to the size of your business and support geographically separate locations.

Management Server in DMZ Supporting Internet Connected Server

The following example shows a configuration that uses a switched ethernet with Security agents on each server. This configuration has a Cross-Site server in the DMZ. The management server shown supports a server at your ISP that is built on an unswitched ethernet. A Security agent resides on one of the four production servers.

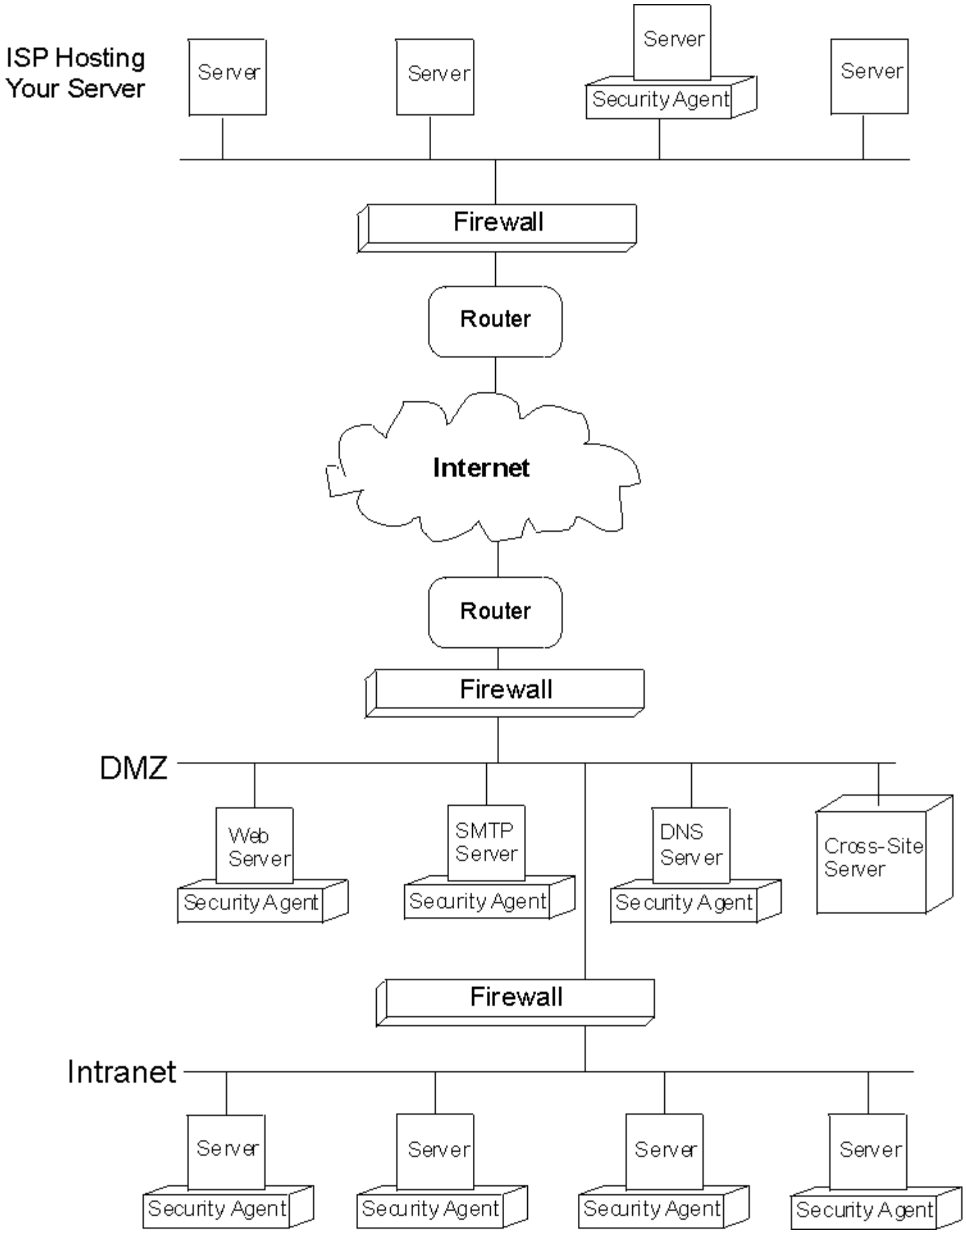

This section outlines the basics of positioning your Cross-Site management server and agents. The most important rule to follow is that you must ensure that all Security agents can connect to the management server. Remember that Security agents communicate with the management server using HTTPS. Also, take into consideration whether you have a switched or unswitched ethernet. Finally, if you use a dedicated computer to a router, always use a UNIX system.

The installation directions for Cross-Site management servers and agents are in the *Tivoli Cross-Site for Security Installation Guide*. Configuration of Security agents is described in “The Security Resource” on page 51.

Getting Started with the Cross-Site Suite

Cross-Site enables you to manage your Internet business-to-business relationships. Cross-Site for Availability monitors the performance of your commerce web site and other sites that you use regularly. Cross-Site for Deployment distributes and updates applications and data for end users. Cross-Site for Security monitors your networks, alerting you to potential network attacks. The Cross-Site server and core services support these applications by providing basic services, such as the event service, task manager, policy manager, data manager, channel manager framework security, and auto updating.

The following steps describe what you need to do to get your Cross-Site management server and agents up and running.

1. If you have not done so, install the management server and Cross-Site agents. For installation information, see the appropriate installation guide.
2. Tivoli recommends that you change the **admin** and **install** users' passwords, if you have not already done so. When you install Cross-Site, these users are created. Changing their passwords provides more security for your system. For information about how to change a user's password, see "Creating and Configuring Users" on page 159.
3. Create users for administrators who need to use the Cross-Site console. Assign the **admin** role to each user; administrators who intend to log in to the console must be assigned the **admin** role. For information about creating and configuring new users, see "Creating and Configuring Users" on page 159.
4. Create one or more users to which you can assign the **install** role. This role enables end users to install the Security agent. For more information about creating and assigning roles to new and existing users, see "Creating and Configuring Users" on page 159.

5. Configure the event service to define the polling interval, the number of events in the event log, and event forwarding directions. For information about setting the polling interval and configuring the event log see, “Viewing Real-time Events.” For information about forwarding events, see “Managing the Event Service” on page 163.

Before using Cross-Site, familiarize yourself with the online help and tutorials. This documentation will acquaint you with the Cross-Site console and Cross-Site concepts.

Getting Started with Cross-Site for Security

Cross-Site for Security enables you to monitor your network for signs of intrusion. It works with the Cross-Site framework services to protect and manage information flowing into your network from the Internet and your intranet. Cross-Site for Security uses real-time intrusion detection, configurable policy, and integration with Tivoli's Enterprise applications.

Cross-Site for Security uses two components: the Cross-Site for Security agent and the Cross-Site management server. These pieces work together to recognize the signature of security threats such as denial of service (floods), known service vulnerabilities, unauthorized traffic, and other suspicious activity.

Note that some of the steps involve performing operations on the Cross-Site console. If you are unfamiliar with the console, you may wish to read "Areas of the Console" on page 31 and "Elements of the Explorer" on page 35.

Before you can use Cross-Site for Security to monitor your network for intrusions, you must set up your Cross-Site server and Security agents. The following steps describe how to get started with Cross-Site for Security:

1. Decide where on your network to install the Cross-Site management server and Security agents. Where you position these components depends on your network topology. For example, you'll want to consider where your Internet service provider connection is, where firewalls isolate networks, and your remote access points. For more information about positioning the Cross-Site server and Security agents on your network, see "Security Configurations" on page 17.

2. Install the management server and Security agents. Installation instructions for Cross-Site management servers and agents are documented in the installation guide.

3. After installation, the Security agents automatically start monitoring the network. Using the default policy, the agents generate incidents and send critical alerts to the Cross-Site console. You can view the alerts as events on the console or forward them to another administrator or developer for further processing.

4. Create a task to upload alert data from the Security agents. (By default, critical alerts are sent to the console and are displayed as events in the event log.) This task uploads all the alerts generated by Security incidents to the management server. To view the uploaded alert data, you must generate one of several types of Security reports.

5. Evaluate the events generated by the Security agents and decide if you need to alter their policy. You alter policy by editing the Security configuration and rules files. These files are called **ids.cfg** and **ids.rules**, and must be edited manually in a text editor. Use the **Policy** explorer in the Security view to edit alert messages. Here you determine the priority and text of Security alerts. For more information about creating and editing Security policy see, “Security Policy” on page 57.

6. If you changed Security agent policy, reevaluate the current settings based on the generated events.

7. Repeat steps 3 and 4 until the management server receives only pertinent events. You might want to create different policies for agents installed in the DMZ, or the area between firewalls, and agents installed on your intranet.

8. Configure Security agent resources. You can edit the agent name and intrusion detection policy assigned to the agent. You make these changes by selecting the agent you want to configure from the **Resources** explorer. For more information about configuring agent resources, see “The Security Resource” on page 51.

After you complete these steps, Cross-Site for Security logs incidents based on the intrusion signatures that concern you.

2

Console Overview and Navigation

This section introduces the Cross-Site console, which is the user interface you will use to manage your Cross-Site domain. The console is the main window into the Cross-Site services and applications.

The console is different than a Microsoft Windows application, which has menus and many windows and dialogs that enable you to perform tasks. Cross-Site's management information and resources are displayed in one window. Drop-down lists and icon buttons, instead of menus, enable you to accomplish tasks. The console may appear to be revolutionary and difficult to use at first. However, once you become familiar with the console, using it becomes very simple. There's a common paradigm for accomplishing tasks used throughout the console. Once you learn the steps for accomplishing one task, other tasks are easier to learn and perform.

This section describes the console so that you may become comfortable with it. Each area of the console is described, so that you can identify the kind of information displayed in each section of the console. Each drop-down list and button is described so that you can quickly identify and use them. Finally, each of the basic Cross-Site services is described and correlated to an area of the console.

Areas of the Console

The console is Cross-Site's desktop; it is the user interface on which you perform basic operations to manage your e-commerce environment. The console displays and organizes the resources, policies, tasks, events, and reports of the different Cross-Site applications. Every user, from the administrator to someone who only views data, uses the same interface.

The Cross-Site console is divided into three areas: the view icons, view, and status bar.

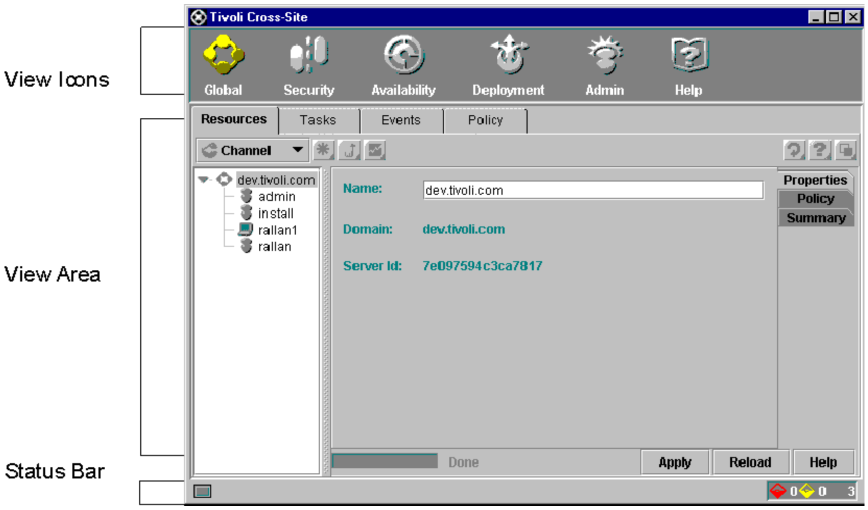

Name: dev.tivoli.com

Domain: dev.tivoli.com

Server Id: 7e097594c3ca7817

The view icons represent each *view*, or area of Cross-Site management. Select the discipline to manage by selecting a view icon. For example, by selecting the Admin view, you can perform administrative operations for Cross-Site.

When you select a view, tabbed panels are displayed below the view icons. These tabbed panels are called *explorers*, and they further divide the areas of Cross-Site management so that resources, policies, tasks, and events are displayed on their own tabbed panels.

The *status bar* is at the bottom of the console and provides information about Cross-Site system progress and the events generated by the Cross-Site server and applications.

The View Icons

The view icons enable you to select an area of Cross-Site management. For example, the Admin view is for Cross-Site administration. It enables you to configure settings for the Cross-Site server and its core services. Likewise, the Security view provides a way for you to view and manage Security resources, policies, tasks, and events.

When you select an icon, a view, in which you can perform functions, is displayed. The views and their icons are as follows.

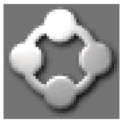

The Global view enables you to view and configure all Cross-Site resources, tasks, events, and policy. This view is an aggregate of the Cross-Site resources, including those available only by installing the Cross-Site applications. This view is provided by Cross-Site core services.

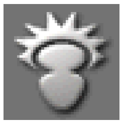

The Admin view enables you to administer the Cross-Site server by configuring settings for the server and its services. Cross-Site core services provides the Admin view. This view provides access to administrative functions, such as creating roles, assigning permissions, managing certificates, and viewing the event log.

The Help view provides a central location for viewing help for the Cross-Site server and applications. Again, this view is part of the Cross-Site core services. You can navigate through the help by using the table of contents tree and searchable index. Also, tutorials are provided to help you get started using Cross-Site. A support page provides information on how to contact customer support.

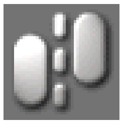

The Security view enables you to control the Security agents. From this view, you can schedule to upload Security information from agents and view Security reports and events.

To change views, click on another view icon. Views replace each other. You cannot display multiple views on the console at one time.

The View

The view is the area below the view icons. When you select a view icon, the Cross-Site console displays the view’s explorers in this view area. Cross-Site explorers are much like tabs provided in the Windows environment. An explorer is a panel that enables you to review and configure resources, policy, tasks, and events in a view. Explorers separate the elements that are managed. For example, if you are working in the Admin view, the **Roles**, **Certificates**, and **Permissions** explorers are displayed. (The **Event Log** is not an explorer, it has no invitation tree.) Explorers differ from view to view. You can see each view's basic set of explorers by selecting a view icon.

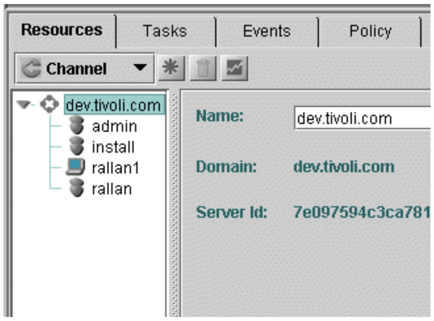

dev.tivoli.com
admin
install
rallan1
rallan

Name: dev.tivoli.com
Domain: dev.tivoli.com
Server Id: 7e097594c3ca781

Each explorer provides a tree for navigation on the left and tables or property panels for configuration on the right. Cross-Site explorers are much like tabs provided in the Windows environment. Explorers are stacked and you can bring an explorer to the top of the stack by clicking on its tab.

Unlike Windows tabs, you can *float* and *dock* Cross-Site explorers. The Float and Dock buttons are in the upper right of the view or panel.

The Float button displays a panel in its own frame, or dialog, outside of the console. By floating explorers from different views at the same time, you can compare information. For example, if you are using the Help explorer, you might float it so that you can read the directions as you work in another view.

The Dock button displays the panel on the console in its respective view. This is the opposite of floating a panel. Note that if you dock a floating panel that resides in a view different than the one you are currently using, you will not see the docked explorer until you display the panel's view.

Explorers, by default, are displayed as tabbed panels in the view. However, you can dock other panels in the Cross-Site console, such as reports. These docked panels are then displayed as tabbed panels in the view, along with the explorers.

The Status Bar

The Status Bar, below the view, provides information about Cross-Site system progress. Messages are displayed on the Status Bar and an indication light informs you when the Cross-Site server is accessing the management server.

The Status Bar also provides a visual representation of the number of events in the system. The red icon on the left indicates the number of critical events in the event log. The yellow icon (in the middle) provides a count of the warning events. The green icon (on the right) gives a count of status events received by the event service. The total count will not exceed the number of events allowed in the event log at one time. For more information, see “Viewing Real-time Events” on page 84.

Elements of the Explorer

When you select a view, the Cross-Site console displays the view's explorers below the view icons. An explorer is a panel that enables you to review and set properties for resources, policy, tasks, and events in a view. It provides a tree for navigation on the left and table or configuration panel on the right. The following graphic shows the major elements of an explorer.

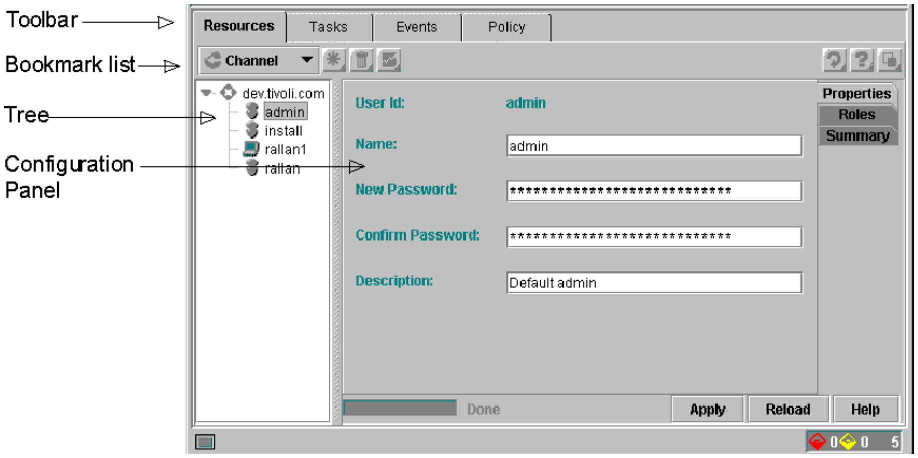

User Id: admin

Name: admin

New Password: ******************************

Confirm Password: ******************************

Description: Default admin

The items on the toolbar initiate specific functions such as creating, editing, and deleting a resource. The tree, always located on the left, varies from view to view; because views act as filters, only resources that apply to that view are displayed in the tree. The view on the right side of the explorer displays information about the selected tree item.

The **Resources** explorers also provide a Reload resources button (in the upper right of the explorer). This button enables you to view changes that other users have made to the management server, without restarting the Cross-Site console. This button refreshes the **Resources** explorer in all views.

Cross-Site displays the following explorers in the Global, Security, Availability, and Deployment views. They are collectively referred to as application explorers.

Resources Enables you to create and configure Cross-Site resources.

Tasks Enables you to view all tasks and their status. You can also start, stop, and configure tasks from this explorer.

Events Provides a central bulletin board of events that are generated by the Cross-Site server and applications.

Policy Enables you to view and configure policies for each Cross-Site resource.

The Admin view provides explorers and panels that enable you to configure and maintain the Cross-Site server and applications. The Admin view provides the following:

Roles Enables you to create roles in the Cross-Site domain, which enables you to control user access at a more granular level.

Event Log Enables you to view real-time events. To configure events and view past events, use the **Events** explorer in the Global view.

Permissions Enables you to assign roles to access control list (ACL) entries for the purpose of access control.

Certificates Enables you to manage the certificates needed to run signed Deployment channels.

The Help view is a specialized explorer that offers information about using the Cross-Site product suite. This explorer provides navigation tools and a browser for the HTML help provided for the Cross-Site server, console, and applications. You can view help using the table of contents, which is displayed in the explorer as a tree. You can also search for a specific term in the index. Additionally, the **About** tabbed panel details the version and license information about the Cross-Site management server.

In addition to explorers, you can dock and float other panels in the view. When docked, these panels appear as tabs.

Config Enables you to view and configure tasks and events. Cross-Site displays a floating Config panel when you press the Edit button in the **Tasks** explorer.

Report Enables you to retrieve information about resources from the Cross-Site management repository. Cross-Site displays a floating report panel when you click on the Reports button in the **Resources** explorer.

The Explorer’s Tree

A tree is displayed on the left side of each explorer. The trees displayed on the Cross-Site console are similar to the tree on the Windows Explorer.

A tree can contain *instances* and *groups*. An instance represents a resource, policy, event, or task, depending on the explorer you have displayed. A group can be a collection or application group, such as those on the **Policy** explorer, and can contain instances and other groups.

In the **Resources** explorer, you create groups called *collections*. You can apply policy to collections. When you assign a policy to a collection, the policy applies to all instances in the group. Refer to each application’s Resources section for instructions on how to create collections.

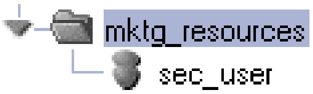

mktg_resources
sec_user

To expand a group, click on the arrow. The following icons may be listed in the tree, depending on the view you have displayed.

agent	user	collection
domain	task	Internet
policy	table-based report	chart-based report
ACL	server	event
role		

This example shows a bit of the tree on the **Policy** explorer. The **dev.tivoli.com** domain is expanded, as well as the **Availability** and **Site Scan Policy** groups. (In this case, **Availability Agent Policy** and **Site Scan Policy** are policy groups.) The **Default** policy is the only instance displayed in this example.

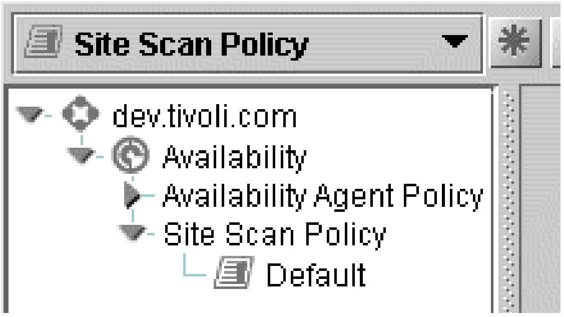

Site Scan Policy
dev.tivoli.com
Availability
Availability Agent Policy
Site Scan Policy
Default

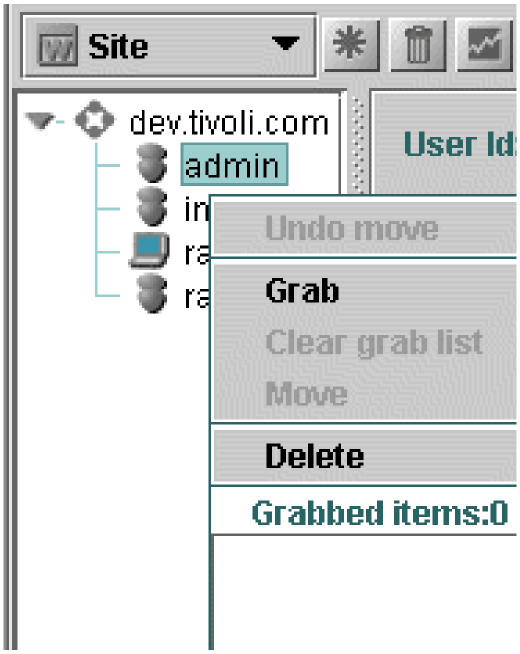

The tree on the Resources explorer provides a menu for moving and deleting resources, including collections. To view the menu, select a resource in the tree and right-click on it.

To move a resource, select it in the tree, right-click on it, and select **Grab** from the menu. The selected resource is added to the **Grabbed items** list. You can grab as many resources as you would like to move. When you select a collection or domain and select **Move** from the right-click menu, all grabbed items are moved to that collection.

To delete a resource from a tree, select it and press the Delete button. You can also select it, right-click on it, and select **Delete** from the menu.

Tables

Tables list instances, such as resources, tasks, policies, and events. When you select a group in the tree, a table is displayed on the right side of the explorer. (Some tables, such as the Event Log, do not have a tree associated with them.) Groups are only displayed in trees and are not displayed in the table. However, if a group contains instances as well as nested groups, the instances are listed in the table. You can then double-click on a table listing to view its configuration information.

You can sort the table columns. Click on the column heading to sort it in ascending or descending alphanumeric order. Click on the column heading a second time to sort the column in the reverse order. An arrow next to the column title indicates how the column is sorted.

You can resize columns. Grab the edge of the column heading and enlarge or reduce the column size as you like. Lastly, you can change the order of the columns. Press the **Ctrl** key, select the column head, and slide it either left or right, until it is positioned where you want it.

Configuration Panels

All resources, tasks, policies, and events have configuration information beyond that which is displayed in the view. To review or set configuration information in the Resources, Events, and Policy explorers, click on the instance in the tree. When you select an instance in the tree, its configuration panel is displayed. In the Tasks explorer, you view and set configuration information by selecting the instance in the table and pressing the **Edit** button. A floating configuration panel is displayed outside of the explorer.

To save your changes, click the **Apply** button. If you wish to reset the panel and ensure that you are viewing the most recent values for the instance, you can click the **Reload** button. This button forces the console to retrieve values from the management repository. The console behaves much like a Web browser, caching values and only updating them when the console is restarted or when you click the **Reload** button. Use this button if you are concerned that the console has old values cached or to cancel any changes you might have made and reset the panel.

Buttons and Gadgets

The following buttons are available on various explorers and panels of the console.

? The Help button, located in the upper-right corner of the view, enables you to display help for the current explorer or table. Click this button to find out what you can do from the current explorer. In the case of tables, find out what the various column headings mean and what values you might see in a particular column. You can find the Help button on all explorers and tables.

⧉ The Float button, located in the upper right of the view, enables you to display the explorer or panel outside of the console, in its own window, as a floating dialog. You find the Float button on the Resources, Tasks, Events, Policy, Roles, Event Log, Contents (Deployment only), Certificates (Admin only), and Permissions explorers.

▣ The Dock button, located in the upper right of the view, enables you to display a floating window as a tabbed panel on the console. You find the Dock button on the Resources, Tasks, Events, Policy, Roles, Event Log, Contents (Deployment only), Certificates (Admin only), and Permissions explorers.

X The Close button, located in the upper-right corner of the view, to the right of the Float or Dock button, enables you to close a floating report or configuration panel.

The Edit button, located in the toolbar of the Tasks explorer, displays the configuration dialog for the selected task or content bundle.

The Generate Reports button enables you to generate reports based on resources. When you click this button, the Report Generator dialog is displayed. You can find the Generate Reports button in the toolbar of the Resources explorer.

The Save to File button enables you to save reports to a file, which you can then print. You can save a report as an HTML file or a text file. When you click this button, the Save Report dialog is displayed. You can find the Save to File button in the toolbar of a Report.

The Delete button enables you to delete the selected item. You can find the Delete button in the upper left side of the view. You can only delete resources, empty collections, tasks, policies, and roles that you created. You find the Delete button on the Resources, Tasks, Policy, and Roles explorers.

The Create a New Certificate button enables you to import certificate files into Cross-Site. Cross-Site uses certificates to sign Deployment channels. When you sign a channel, end users can see that your channel is secure to download. You can find this button on the Certificates explorer in the Admin view.

The Create New Foreign Domain button enables you to create a new foreign domain for making Deployment channels available to users in that domain. A foreign domain is simply a domain associated with a management server other than your local management server. You can find this button on the Roles explorer in the Admin view.

The Create a New Role button enables you to create a new role for using Cross-Site services and applications. You find this button on the Roles explorer in the Admin view.

The Update Events button enables you to retrieve real-time events from the management repository and view them in the event log. You find this button on the Event Log in the Admin view.

The Clear All Events button enables you to purge the event log of all events. (The events, however, remain in the management repository and can be viewed on the Events explorer.) You find this button on the Event Log in the Admin view.

The Create New button enables you to create new instances. You select which instance to create from the drop-down list in the explorer toolbar. You find this button on the Resources, Tasks, and Policy explorers.

The Terminate Task button enables you to stop a task that is running. You find this button on the Tasks explorer.

The Start Task button enables you to start a task immediately. You find this button on the Tasks explorer.

The Fetch Events button enables you to manually retrieve events from the management repository. You find this button on the Events explorer.

The Reload button enables you to reload the information displayed on the console. This button is useful if another user has made changes in the Cross-Site management repository since you accessed the information. You can find this button on the Resources explorer.

The Show Availability Events button enables you to filter events in the event log such that only Availability-related events are displayed. You find this button on the Event Log in the Admin view.

The Show Deployment Events button enables you to filter events in the event log such that only Deployment-related events are displayed. You find this button on the Event Log in the Admin view.

The Show Security Events button enables you to filter events in the event log such that only Security-related events are displayed. You find this button on the Event Log in the Admin view.

The Show General Events button enables you to filter events in the event log such that only global events are displayed. You find this button on the Event Log in the Admin view.

In addition to buttons, the explorers provide drop-down lists to enable you to use Cross-Site functions. The following drop-down lists are located on the application explorers' toolbars.

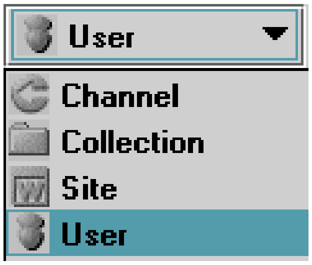

The Resource drop-down list enables you to choose which resource you want to create. After selecting a resource, click the Create New button to display the resource’s configuration panel. You find this drop-down on the Resources explorer.

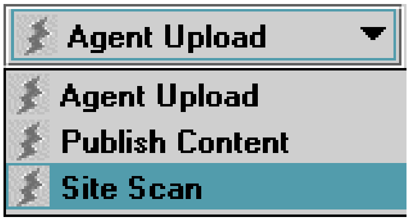

The Task drop-down list enables you choose which task you want to create. After selecting a task, click the Create New button to display the configuration panel. You find this drop-down on the Tasks explorer.

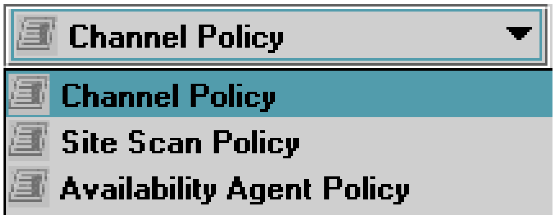

The Policy Type drop-down list enables you to choose which policy type you want create. After selecting a policy type, you click the Create New button to display the policy on the configuration panels. You find this drop-down list on the Policy explorer.

The Core Services

The Cross-Site server and console provide some basic services that all of the Cross-Site applications use. Most of the services are visible and configurable from the console.

The services provided by the Cross-Site server include the following:

- The Event Service
- The Task Manager
- The Policy Manager
- The Data Manager
- The Channel Manager
- Framework Security
- Auto-updating

Some of the services correspond to elements on the console. Others, such as the data manager, auto-updating, and most of the security service, are not exposed through the console. All of the services provided by the Cross-Site server are collectively referred to as the Cross-Site framework, the framework services, or the core services.

In addition to the services provided by the Cross-Site server, the Cross-Site console provides several additional services including:

- The Report Generator
- Administration

The console relies on information stored in the Cross-Site management repository to provide these services. It also relies on core services such as the data manager to process this information.

The Event Service

An *event* is any alert, error message, or status message generated by the Cross-Site server or applications. The source of an event can be the Cross-Site server itself or a Cross-Site agent. The Cross-Site server receives the event, stores it in the Cross-Site management repository, and sends it to the defined destinations.

Each event is assigned a priority. An event can be one of three priority levels:

- Critical, such as when the server goes down or when the Task manager cannot access the database
- Warning, which indicates an incident that is not currently critical but that could progress to that level
- Status, such as when the task manager posts a message when a task begins

By default, the event service posts events to the **Events** explorer on the console. You can view detailed event information here. From the Admin view, you can also configure the event service to send events to other destinations, such as an e-mail address or pager.

For information on viewing events, event priorities, and how the event service works, see the event information in each Cross-Site application's help. This section also provides information on customizing and purging events from the management repository.

The Task Manager

The Cross-Site server provides a scheduling facility that enables you to launch and track the progress of a *task*. A task is any operation that can be scheduled, such as uploading data or running a scan. Tasks can be internal and external processes (to the server). The facility that schedules and tracks tasks is referred to as the task manager.

The table on the **Tasks** explorer lists Cross-Site tasks and their statuses. A task's status may be "scheduled," "running," "succeeded," or "failed." An audit trail of all tasks is stored in the Cross-Site management repository.

For detailed information, see the tasks section in each Cross-Site application's help.

The Policy Manager

A *policy* is a set of rules that defines the behavior of a resource. The policy manager is the service that keeps track of and enforces the policies that are assigned to resources. When the task manager launches a task, it does so according to the policy applied to the resource on which the task will run.

Policy is a method of configuring a resource. It encapsulates configuration parameters, or rules, which “live” in their own container. You can then reuse the rules for any resource that needs a particular set of parameters by applying a policy to the resource.

Cross-Site provides default policy settings for each policy type. You can create policy through the console from the **Policy** explorer in various views. Refer to each Cross-Site application’s help for instructions and further information on policy.

The Data Manager

The data manager is a facility that gives the Cross-Site server and applications access to the management repository. The management repository is the database that acts as a persistent store for all Cross-Site management information.

Cross-Site does not require a third-party database to store information. Instead, Cross-Site uses the data manager, which provides a layer between the Cross-Site server and the database. The data manager uses a Java database connectivity (JDBC) driver, because the Cross-Site server is Java-based.

The Channel Manager

The channel manager is responsible for publishing channels. The Cross-Site console and agents rely on this service to download updates. Cross-Site, using the channel manager, can automatically update itself. This ensures that your Cross-Site installation is up-to-date and using the latest binaries. The self-updating feature of Cross-Site is provided by Marimba’s Castanet technology.

Framework Security

The Cross-Site framework provides three methods of security: authorization, authentication, and confidentiality through secure socket layer (SSL) connections.

With authorization, you can restrict or grant access to services and applications for all Cross-Site principals. A principal is any user, agent, or group that is granted roles and, therefore, permissions to perform Cross-Site operations. To configure authorization, you can create Cross-Site users through the console and give them roles for performing operations. The **Resources** explorer in the Global view enables you to create and configure new users. See “Creating and Configuring Users” on page 159 for instructions.

Cross-Site also authenticates all Cross-Site clients that request access to Cross-Site services on the server. The Cross-Site server verifies the identity of the Cross-Site client and ensures that it has rights to the service.

Finally, when the Cross-Site server communicates with Cross-Site agents, it uses the secure sockets layer (SSL) protocol (if it is configured to do so). The Cross-Site server is dependent on a web server, which provides support for SSL. The installation of the web server is a prerequisite step in the installation of Cross-Site. See the installation guide for information on how to configure the Cross-Site server.

Auto-updating

The Cross-Site server periodically checks the Tivoli fulfillment server for updates to itself, and for updates to any consoles and agents connected to the Cross-Site server. The fulfillment server is the same one from which your Cross-Site server was originally installed. The Cross-Site server checks the Tivoli fulfillment server at a regular interval. If it finds an update, the Cross-Site server automatically downloads it. However, the Cross-Site server does not automatically install the update; a Cross-Site administrator must initiate the installation.

Cross-Site agents and consoles, in turn, periodically check the Cross-Site server for updates, which were downloaded from the fulfillment server. If an agent or console finds an updated copy of itself, it notifies the user. The user can then choose to update the Cross-Site agent or console immediately or at a later time.

This auto-updating feature is provided by Marimba’s Castanet technology.

The Report Generator

With the report generator, you can create reports from the data collected from Cross-Site agents. The data is retrieved from the Cross-Site management repository using the data manager. The console provides several types of reports, including graphs, charts, and text-based reports. You simply need to select a report type and generate the report.

For information on each report type and steps on how to generate them, see the reports section in each Cross-Site application’s help.

Administration

Being able to administer an application is crucial and essential to its use. Cross-Site gives you a way to centrally manage users. You can create and configure users from the **Resources** explorer of the Global view. User access is enforced by the authentication and authorization security methods, which are described in “Managing Access Control” on page 155.

You can configure events to specify where they are sent after the Cross-Site server processes them. Configure events from the Admin view. See “Managing the Event Service” on page 163 for more information.

3

Security Procedures

This section of the guide provides conceptual and procedural information about the Cross-Site for Security resource. It provides the basic information you need to use Cross-Site for Security, through the console.

This section describes:

- the Security agent and how you can configure it
- Security policy, which is the most important part of configuring Security agents to detect activity that could affect your environment
- how to configure and run the upload task, which uploads data from agents to the Cross-Site server

This section also includes procedures for viewing and configuring events, and generating reports.

The Security Resource

The Cross-Site for Security agent continually monitors a network in order to collect all network packets that are “seen” by the system on which it is installed. The Cross-Site agent filters these packets, looking for various *signatures* that might represent suspicious activity. A signature is sometimes a string of text, which might include numbers and symbols that uniquely identify contents of a packet. Other signatures are more closely tied to the protocol they are monitoring, and may or may not contain an actual string. Certain signatures indicate attacks on the network, referred to as *incidents*. Incidents have a priority associated with them that is used to determine whether the agent generates an immediate critical alert or simply saves it to a local log file for subsequent upload. The signatures that the Security agent detects and associates with incidents are described in “Security Policy” on page 57.

When a Cross-Site agent is installed, it registers with its Cross-Site management server. A representation of the agent is then created in the server’s management repository. An agent icon is displayed on any Cross-Site console that connects to that management server. Cross-Site agents are displayed in the tree on the left side of the **Resources** explorer.

You can install Cross-Site for Security agents on UNIX and AIX, as a daemon, or on Windows NT, as a service. Cross-Site for Security starts automatically when the system reboots. It runs on the system in the background, regardless of whether or not a user is logged on.

Each time a Security agent detects an incident, it creates an alert in a log file on the agent’s host machine. If the alert is assigned a priority of 1, it is considered critical. The Security agent *immediately* sends critical alerts to the alert log in the Cross-Site server’s management repository. For more information on the alert log and types of Security events, see “Security Events” on page 75. All other alerts are uploaded to the alert log periodically. See “Security Tasks” on page 69 for information on scheduling these alert log uploads through a task.

In addition to detecting incidents and generating alerts, the Cross-Site agent regularly contacts the Cross-Site server to let it know that the agent is running. This is called sending a *heartbeat*. When the Cross-Site agent sends a heartbeat to the Cross-Site server, the server responds. It notifies the agent of any updated configuration information and new signatures. The agent automatically downloads and installs these updates. If the Cross-Site agent does not send a heartbeat to the management server, the server generates an event and displays it on the **Events** explorer on the Cross-Site console.

You can specify how often each Cross-Site agent sends a heartbeat in the **ids.cfg** file. For information on how to do this, see “Configuring Scan and Flood Parameters” on page 63.

Using The Resources Explorer

The **Resources** explorer is displayed in the Global, Deployment, Security, and Availability views.

This explorer enables you to do the following:

✳ Create new resources by choosing a resource from the Resource drop-down list. Then press the New button shown here. Cross-Site resources are configured on the **Resources** explorer, in the Global or application view. (You create users *only* in the Global view.)

(no icon) Edit a resource by selecting it in the tree.

🗑 Delete a resource by selecting it and pressing the Delete button. You can also delete a resource by selecting it and right-clicking on it. You can then select **Delete** from the menu.

(no icon) Move a resource by selecting it and right-clicking to display a menu. Select the **Grab** option to add it to the list. Select the collection or domain where you are moving the resource, right-click on it, and select **Move**.

📈 Generate reports by pressing the Reports button. The report generator, which has a drop-down list of reports, is displayed.

↻ This explorer also provides a Reload button. This button enables you to view changes that other users made since you accessed the information currently displayed on the console.

For more detailed information about resources and how to configure them, see the resources sections of each application user's guide. For more information about reports, see the reports sections of each application user's guide. For information about configuring with users, see "Access Control Lists and Entries" on page 158.

Configuring the Security Agent

Each Cross-Site agent is configured during installation. You can view a Cross-Site agent's properties and change its policy by completing these steps from the Cross-Site console:

1. Select the **Resources** tab in either the Global or Security view. A list of domains is displayed in the tree on the left side of the explorer. For this release, only your Cross-Site server domain is displayed.

2. Expand the domain by clicking on the arrow next to it in the tree.

3. Select either the domain or collection in which the agent resource resides. When you do so, each resource in the domain or collection is listed in the tree.

4. Select the agent you want to edit. Its configuration panels are displayed on the right.

5. Select the **Properties** tab to view information specific to the Cross-Site agent, including the following:

   Name The name of the Cross-Site agent (This field is editable.)

   UUID The agent's Universally Unique Identifier (UUID), which is generated during installation

   Description A description of the agent, which is set during installation.

   IP address The IP address of the computer on which the agent resides

   Owner User ID
   The Cross-Site user name of the person using the Availability agent. This user name is created when the agent is installed.

   Certificate ID
   The ID number of the security certificate used by the agent

6. Select the **Policy** tab to view and assign policy to the agent.

7. Select a policy from the **Intrusion Detection Policy** drop-down list.

You should reassign policy in any of the following situations:

- You want the agent to detect a new set of signatures
- You want to establish new scan and flood parameters for the agent
- You have updated the list of alert messages
- You want to change the interval at which the agent sends heartbeats to the management server

Create new policy using the configuration files and the **Policy** explorer in either the Global or Security views. For more information on policy, see “Security Policy” on page 57.

8. Select the **Summary** tab to display the following information about the agent:

   Name The name of the agent resource. This is the name provided during installation.

   Resource type
   The agent’s resource type

   Resource ID The resource ID of the agent, which was generated when the agent was installed

   Creator The user name of the person who installed the agent (This field is not supported in this release.)

   Last modified
   The date and time the agent’s properties were last modified. (This field is not supported in this release.)

9. Click the **Apply** button to save your changes to the agent’s configuration or click the **Reload** button to clear the fields in the configuration panel.

The next time the agent contacts the management server, it downloads the updated configuration information. Any changes are applied to the Security agent.

When you delete a Cross-Site agent, its record is deleted from the management repository. However, if you do not delete the agent software from the client machine, Cross-Site for Security generates a **Security agent *X* unknown** event when it detect’s the agent’s heartbeat. (The priority of this event is warning.) Therefore, delete the agent from the client machine first, then delete its icon from the console.

Note: *The Cross-Site for Security agent software is released under the Library GNU Public License (LGPL). If you would like to relink the agent with a replacement library, please contact Tivoli Cross-Site customer support to obtain the object files of the agent (which are required for relinking).*

Creating a Collection

A collection is a container for resources or other collections. Collections enable you to organize related items into logical groups. You can also apply policy to a collection, which applies the policy to all appropriate resources in the collection. See “Applying Policy to a Collection” on page 67 for more information.

To create a collection, complete the following steps from the Cross-Site console:

1. Select the **Resources** explorer in either the Global or Security view.
2. Expand your Cross-Site server domain by clicking on the arrow next to it in the tree.
3. To create a new collection, select **Collection** from the drop-down list below the **Resources** tab and click the **Create New** button. The **Create new collection** dialog is displayed.
4. Enter a name for the collection. Assign any name, but remember that it should be meaningful to other Cross-Site users
5. Press the **Create** button. The collection is displayed in the tree and its configuration panels are displayed on the right.
6. Select the **Properties** tab to edit the name of the collection. The following additional information is displayed on the **Properties** panel:

   Domain The domain name of the management server on which the collection was registered

   Server ID The unique ID of the collection’s management server
7. Select the **Policy** tab to select a policy to assign to the collection. You can assign a policy for each type listed. All available policies are listed in each of the policy drop-down lists. Following is the list of policy types, which will vary depending on your configuration of Cross-Site:
   - Availability Agent Policy
   - Channel Policy
   - Site Scan Policy
   - Intrusion Detection Policy

You apply policy to a collection the same way you apply policy to any other resource. There are, however, additional things to consider when applying policy to a collection. See “Applying Policy to a Collection” on page 67 for more information.

8. Select the **Summary** tab to view information about the collection. You cannot edit the information in the **Summary** panel. The following additional information is displayed in this panel:

   Name The name of the collection

   Resource type
   The collection’s resource type

   Resource ID The unique ID of the collection, which is generated when the collection is created

   Creator The user name of the person who created the collection. (This field is not supported in this release.)

   Last modified
   The date and time the collection’s properties were last modified. (This field is not supported in this release.)

9. Click the **Apply** button to create the collection or click the **Reload** button to reset the fields in the configuration panel.
10. To edit a collection, select it from the tree and edit the collection’s configuration information in the panels on the right.
11. To delete a collection, select it from the tree and press the **Delete** button. You cannot delete a collection that has resources in it, you must delete each resource in the collection first.

You can create new resources within a collection by selecting the collection before clicking the **Create New** button. You can also move existing items into the collection. Finally, you can create collections within collections.

Security Policy

Cross-Site for Security policy specifies the actions that should occur when a particular incident is detected. By configuring Security policy and applying it to a Cross-Site for Security agent, you define when alerts are generated, based on the signatures detected by that agent. Each agent consults its policy to determine how it will operate.

You configure policy for the Security agent by creating alert messages and editing two different files: **ids.cfg**, and **ids.rules**. When a Security agent starts, it contacts the Cross-Site management server and downloads these configuration files. These files establish the various rules, or policy, used to define the behavior of the agent to which they are applied.

Using The Policy Explorer

The **Policy** explorer, found in the Global, Security, Availability, and Deployment views, enables you manage policy. From this explorer you can do the following:

✳	Create policy by selecting a policy from the Policy Type drop-down list and clicking the **New** button shown here.
🗑	Delete policy by selecting a policy and clicking the **Delete** button.
(no icon)	Edit policy by selecting a policy from the tree on the left. The configuration information is displayed in the table on the right.

For more detailed information about policy and how to configure policy, see the policy section of the application user's guide.

Creating Policy Files

Cross-Site provides two utilities for setting Security policy. The first utility creates a new policy and installs the configuration files that support the policy in the proper directory. The second utility notifies the management server that you have modified the configuration files.

The configuration files provide default policy settings only. Tivoli recommends that you test these default settings in your environment before you edit them. After you have used the product, update the configuration files to accommodate your needs. For example, you might realize that you are missing active ports, or you might be getting a lot of network traffic from a trusted machine that you really do not need to monitor.

The **xs_createpol** command is provided for you to create new policy files. You must log in to the machine on which the Cross-Site server is installed to use this command. The **xs_createpol** command is located in the *base_dir***/XSITsagt/bin** directory, where *base_dir* is the base directory specified when the Cross-Site server was installed. You can set the PATH environment variable to include this directory, or you can run the command by specifying its full path.

The **xs_createpol** command creates a directory under the **policy.root** directory. Typically, **policy.root** points to *base_dir***/XSITsagt/lib/policyfiles/**. An empty **ids.cfg** file and **ids.rules** file are created in this directory if no arguments are specified.

The syntax of the **xs_createpol** command is as follows:

xs_createpol [**-cfg** *cfg_file*] [**-rules** *rules_file*] *name*

where:

-cfg *cfg_file*
Specifies *cfg_file* as the file to be copied into the new *policy_ID***/ids.cfg** file.

-rules *rules_file*
Specifies *rules_file* as the file to be copied into the new *policy_ID***/ids.rules** file.

name
Indicates the name of the policy to create. This name is displayed in the **Policy** explorer on the Cross-Site console.

The following example creates a directory, identified by a unique Policy ID number (numbers start at 101), and copies the contents of the default policy configuration files (in folder **5**) to new files in this new directory. It also names the new policy **SecurityPolicy**. Note that this command is on a single line. It has been split to accommodate the page size.

```
xs_createpol -cfg /management_server/policy/5/ids.cfg -rules
/management_server/policy/5/ids.rules SecurityPolicy
```

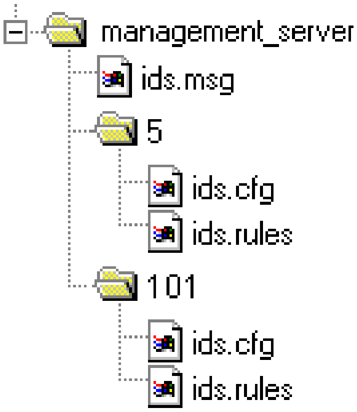

Remember the ID number of this policy subdirectory. After you have modified the new policy’s configuration files to suit your organization's security needs, you will need to apply your changes by referring to this unique policy ID. See “Notifying the Server of Policy Changes” on page 65 for more information. An entry for the new policy is also created in the management repository.

You may have noticed that a new **ids.msg** file is not created for the new policy. That is because the alert messages in this file are global—all Cross-Site agents reference the same **ids.msg** file.

There may be times when you want to move the files under **policy.root** to a new location. For example, you may want to store these files on a drive that is being backed up regularly. Creating new policy can be labor intensive, and no one wants to have to duplicate effort due to a disk failure.

If you decide to change the location of your policy files, Cross-Site requires that you update the **policy.root** property in the **dbconfig.properties** file. The **policy.root** property specifies where your Security policy files are held in the management server’s file system. During installation, the **dbconfig.properties** file is copied into the *base_dir*/**XSITsagt/lib/properties** directory. Open and scroll to the bottom of the **dbconfig.properties** file to find the **policy.root** property. Edit the default value for this setting to match the new location of your policy files.

Note: *Always shut down the management server before moving the policy files. For more information see, “Restarting the Cross-Site Server” on page 175.*

Changing the Priority of an Alert

After testing the default policy, if you wish to display a particular Security alert on the Cross-Site console (as an event), you must edit the alert’s priority. The way Cross-Site treats an alert is directly related to the alert’s priority. For example, if the alert’s priority is 3, the Security agent logs the alert to a local log file, which the agent then periodically uploads to the Cross-Site server. If you want the alert to generate an event on the console, it must have a priority of 1.

Following are the basic steps required to change the priority of an alert. Each step provides a reference to the procedure that you need to perform.

1. Edit the priority of the alert message using the Cross-Site console. See “Configuring Alert Messages” for instructions.

2. Edit the priority of the related alert in the **ids.rules** file of the appropriate agent's policy. For more information, see "Configuring Intrusion Alert Signatures" on page 64.

   a. Find the relevant protocol section in the **ids.rules** file.

   b. Edit the priority of the alert so that is matches the priority of the alert message. For example, if you changed the priority of the alert message from 2 to 1, make the same change in the **ids.rules** file.

3. Run the **xs_updatepol** command to notify the Cross-Site server that you changed the policy. See "Notifying the Server of Policy Changes" on page 65 for instructions.

Adding a New Rule to a Policy

After testing the default policy in your environment, you may decide that you want to add a new rule to a Security agent policy. For example, perhaps you have discovered a new signature that you want to add to an agent's policy. You need to update the policy with this new rule. Each step provides a reference to the procedure that you need to perform.

The general steps required to add a new rule to policy are as follows:

1. Create a new alert message for the signature using the Cross-Site console. For instructions, see "Configuring Alert Messages" on page 61.

   Make note of the priority you assign the alert message and the message token generated by Cross-Site. You need this information for the next step.

2. Edit the **ids.rules** file of the agent's policy to which you want to add the new rule.

   a. Find the relevant protocol section in the **ids.rules** file.

   b. Enter the new rule using the same syntax as the other rules in this section. If necessary, refer to "Cross-Site for Security Features" on page 97 for information on each protocol. The various alert signatures (including syntax examples), are documented here.

Run the **xs_updatepol** command to notify the Cross-Site server that you have changed the policy. See "Notifying the Server of Policy Changes" on page 65 for instructions.

If you entered the rule incorrectly, an error message is displayed. The error message explains what input was expected and where you can find the error. The file name and line number where Cross-Site encountered the problem are provided.

Make the necessary changes to correct the problem.

Run the **xs_updatepol** command again to notify the Cross-Site server that you have changed the policy.

Any new agents to which you apply this policy refer to this new rule, as well as the existing rules in the agent policy.

Configuring Alert Messages

You can modify the custom alert messages provided with Cross-Site, and create your own alert messages from the Cross-Site console. Cross-Site provides these editable alert messages to get you started; you can edit both the message text and the priority of each alert message.

Use the **Policy** explorer on the Cross-Site console to display alert messages.

1. In the Policy tree, select the Security application group. The policy categories for Security are displayed in the tree.

2. Select the **Intrusion Detection Policy** category and then a specific security policy. The **Reserved Alert Message Table** and the **Custom Alert Message Table** components are displayed.

   Note: *You cannot edit the information in the Reserved Alert Message Table.*

3. Select the **Custom Alert Message Table** component in the Policy tree. In the table on the right side of the **Policy** explorer, the list of custom alert messages is displayed.

Editing Alert Messages

Cross-Site provides customizable alert messages that you can edit according to your needs. After you have displayed the custom alert messages on the console, complete the following steps to edit an alert message:

1. Select an alert message by clicking on any field in its row.

2. Change the alert message priority by deleting the existing number and replacing it with a different value. Alert message priorities are ranked from 1 to 5, with 1 representing the highest priority level.

   The priority of an alert message determines how Cross-Site handles the alert. For example, priority 1 alerts are sent to the Cross-Site server immediately, and displayed on the **Events** explorer of the console. For more information, see “Logging Alerts” on page 76.

   Note: *If you change the priority of an alert message, you must also make the same change to the alert in every **ids.rules** file that references that alert.*

3. Edit the text of the alert message. Change the message as you like, but keep in mind that it must be meaningful to other Cross-Site users.

4. When you are satisfied with your changes, click **Apply**. Your changes to the policy messages are saved in the management repository.

   If you do not wish to keep your changes, click the **Reload** button to refresh the fields in the table with the latest data in the management repository.

Creating Additional Alert Messages

In addition to providing alert messages that you can tailor to your needs, Cross-Site also enables you to create your own alert messages. Anytime you wish to create a new rule in the **ids.rules** file, you must first create an associated alert message in the file. Complete the following steps to create a new alert message:

1. Open the list of custom alert messages and press the **New** button located in the top left corner of the panel. The row immediately after the last entry in the table is highlighted.

2. Change the priority of the new alert message, or accept the default value of **5**. Alert message priorities are ranked from 1 to 5, with 1 representing the highest priority level.

   The priority of an alert message determines how Cross-Site handles the alert. For example, priority 1 alerts are sent to the Cross-Site server immediately, and displayed on the **Events** explorer of the console. For more information, see “Logging Alerts” on page 76.

   Note: *You do not need to enter a value in the* ***ID*** *field. When you click the* ***Apply*** *button, Cross-Site assigns the next available number to the alert message. This ID becomes permanently associated with the alert message. It is stored in the alert log each time a Security agent generates an alert associated with the new rule added to the* ***ids.rules*** *file.*

3. Enter the text of the alert message. This can be any message you like, but keep in mind that it must be meaningful to other Cross-Site users.

4. Click the **Apply** button. Cross-Site adds the new alert message to the list of user-defined, custom messages.

   If you do not wish to save the alert message, click the **Reload** button to refresh the fields in the table with the latest data in the management repository.

The next time a Cross-Site agent contacts the Cross-Site server, the server notifies the agent that the alert messages have been updated. The Cross-Site agent then automatically downloads and installs these changes.

Configuring Scan and Flood Parameters

Use the **ids.cfg** file to configure scan and flood parameters. This file lists the parameters for the types of scans and floods detected by the Cross-Site for Security agent, and provides a default value for each parameter. For a table of scan and flood parameters, see “Scan and Flood Summary” on page 102.

Use the first section of the **ids.cfg** file to define the range of port numbers to check for scans and to specify *peak time.* Peak time represents the hours and days that you consider to be legitimate business hours. It is a range of hours during a single 24-hour period, set for particular days of the week.

The Security agent can only accept one PEAK statement per policy. You can also only define one peak time range in that statement. For example, you cannot specify PEAK as 8 a.m. – 12 p.m. and 2 p.m. – 6 p.m. When specifying peak time, the value for the beginning second variable must be less than the ending second. For example, you cannot have a peak time ranging from 6 p.m. to 2 a.m. This range of time spans midnight, which marks the beginning of the next day.

The PEAK time variable is defined in the **ids.cfg** file as follows:

```
PEAK BegSec EndSec DaysOfWeek
```

where:

BegSec Specifies the time PEAK begins, in seconds. Calculate this value as follows: (hour*3600) + (min*60) + seconds.
For example, 9:00 a.m. is calculated as (9*3600) + (0*60) + 0, which is 32400 seconds.

EndSec Specifies the time PEAK ends, in seconds. Calculate this value as follows: (hour*3600) + (min*60) + seconds.
For example, 5:00 p.m. is calculated as (17*3600) + (0*60) + 0, which is 61200 seconds.

DaysOfWeek Specifies the days of the week to which PEAK time applies. The days of the week are numbered zero through six, with zero representing Sunday, one representing Monday, and so on.

The following example defines PEAK time as 7:00 a.m. to 6:00 p.m. (18:00), for Monday through Friday.

```
PEAK 25200 64800 1 2 3 4 5
```

The rules listed in the **ids.rules** file reference the range of port numbers and peak values, which the Security agent then uses when detecting intrusion signatures.

The next sections of the **ids.cfg** file specify the values that indicate the types of scans and floods that Cross-Site for Security detects. You can change these values or accept the defaults provided. For information on the values in the **ids.cfg** file and the type of floods and scans they represent, see “Detecting Scans and Floods” on page 99.

Use the last sections of the **ids.cfg** file to specify time-out parameters for inactive hosts and sessions, as well as how long Cross-Site for Security waits before periodically purging memory.

- The *idlehost* variable specifies the length of time for which the agent counts packets and stores packet statistics in memory for a particular host.
- The *idlesession* variable specifies how often the agent stops monitoring TCP/IP connections.
- The *idlepurge* variable specifies the intervals at which inactive hosts and sessions are actually purged. This parameter is also used to specify the interval between agent heartbeats. For more information on heartbeats, see “Monitoring Traffic” on page 13.

For example, when *idlehost* is 200 seconds, *idlepurge* is 100 seconds, and the packet counting starts at 60 seconds, the agent stops counting packets and storing packet statistics at 260 seconds. However, the packets are not actually purged from memory until the next *idlepurge* interval occurs: at 300 seconds.

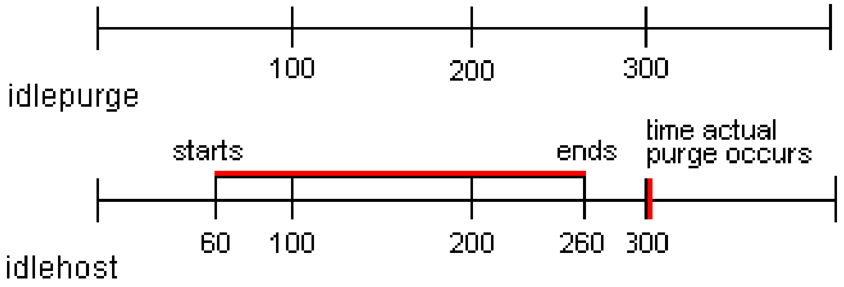

Note: *The value of the idlepurge parameter must be no more than half the value of the idlehost and idlesession parameters.*

Configuring Intrusion Alert Signatures

Use the **ids.rules** file to configure alert signatures. The **ids.rules** file lists the actual rules for Security agents to follow when detecting attacks by intruders, and when sending notifications of potential security breaches. The order and precedence of rule processing follows the order of appearance of the rules in the **ids.rules** file. For example, if there are two rules in the **ids.rules** file, the first looking for “cat” and the second looking for “dog”, and data is processed that contains both of these strings, the alert corresponding to the first rule (the one looking for “cat”) is triggered. The second rule is effectively ignored. This is not considered a weakness in the system because signatures are rarely combined (few select cases).

ALLOW rule processing does not follow this same rule order precedence. In the case of an ALLOW rule, the associated alert is not fired until the last ALLOW statement is processed. For more information. see “The ALLOW Command” on page 105.

The sections of the **ids.rules** file correspond to the types of services and attacks you can detect with Cross-Site for Security. Each section outlines the rules you can choose from and the syntax to use for each. Example rules are also provided.

After you install the Cross-Site server, print out the **ids.rules** file so you can easily reference it while reading the signature explanations in “Cross-Site for Security Features” on page 109. The file contains commented-out rules that enable you to setup protocol-level ACLs. The protocols that Cross-Site for Security supports are discussed in detail in “Cross-Site for Security Features” on page 109. You can edit the rules in the **ids.rules** file according to your security needs, or use the default values provided.

Tivoli recommends that you modify the default rules provided with Cross-Site only if you know the topology of your network environment well enough to properly adjust the protocols. It is best to create a new policy to incorporate your changes, rather than change the default policy. See “Creating Policy Files” on page 58 for instructions on how to create a new policy.

For information on the various alert signatures (including syntax examples), see “Cross-Site for Security Features” on page 109.

Notifying the Server of Policy Changes

After you have modified the **ids.cfg** or **ids.rules** files to suit your security needs, you must run the **xs_updatepol** command. This command notifies the Cross-Site server that you have modified the policy configuration files. The next time a Cross-Site agent contacts the Cross-Site server, the server notifies the agent that updated policy configuration files are available. The agent retrieves these updated files from the server if the updated policy applies to the particular Cross-Site agent. New policies and policy changes will not be detected and “picked up” by the Cross-Site server until you run the **xs_updatepol** command.

You must log in to the machine on which the Cross-Site server is installed to use this command. The **xs_updatepol** command is located in the *base_dir***/XSITsagt/bin** directory, where *base_dir* is the base directory specified when the Cross-Site server was installed. You can set the PATH environment variable to include this directory, or you can run the command by specifying its full path.

The syntax of the **xs_updatepol** command is as follows:

xs_updatepol *polID*

where:

polID The ID number of the policy directory that was created when you ran the **xs_createpol** command.

The following example notifies the Cross-Site server to look for policy updates in directory **101**, in the **policy.root** directory.

```
xs_updatepol 101
```

Applying Policy to an Agent

After creating a policy, you must apply it to the agents for which it is intended. Otherwise, the agent uses the default Intrusion Detection policy. To apply a policy to an agent, complete the following steps from the Cross-Site console:

1. Select the **Resources** explorer in either the Global or Security view. Your Cross-Site server domain is displayed in the tree on the left.

2. Expand this domain by clicking on the arrow next to it in the tree. A list of resources and collections is displayed.

3. Select the agent to which you want to apply policy. The agent’s configuration panels are displayed on the right.

   You can also select a collection to apply policy to all the resources within the collection. You apply policy to a collection the same way you apply policy to any other resource. However, there are additional things to consider when applying policy to a collection. For more information, see “Applying Policy to a Collection” on page 67.

4. Select the **Policy** tab.

5. Select a policy from the **Intrusion Detection Policy** drop-down list.

6. Click the **Apply** button to assign the policy to the agent, or click the **Reload** button to clear the fields in the configuration panel.

After policy is applied to an agent, it becomes effective the next time the agent contacts the server to check for policy updates. (The Security agent contacts the Cross-Site server based on the interval specified for the *idlepurge* variable in the **ids.cfg** file.) If there are policy updates available, the Security agent downloads these changes.

Events are generated when the policy applied to the Security agent is violated. You can view a list of events on the **Events** explorer, or generate reports based on event data and data collected from Cross-Site agents. Both events and reports are discussed in greater detail in their own sections.

Applying Policy to a Collection

When you apply policy to a collection, it is applied to all resources in the collection, including nested collections. The policy applied to the parent collection is also applied to new resources that are added to the collection.

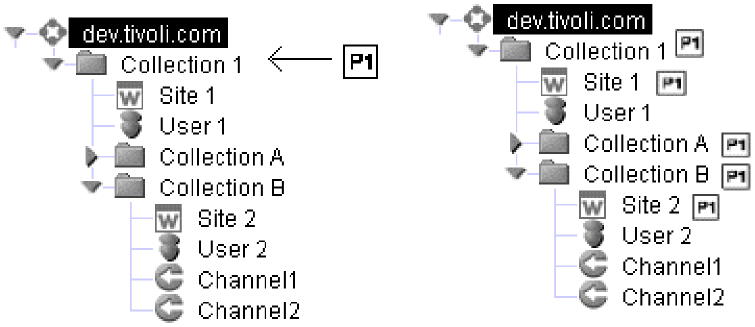

You can, however, override the policy applied to a resource or nested collection within a parent collection. Then, if you change the policy applied to the parent collection, the policy override is maintained: the new policy is applied only to those resources whose policy matches the collection's policy.

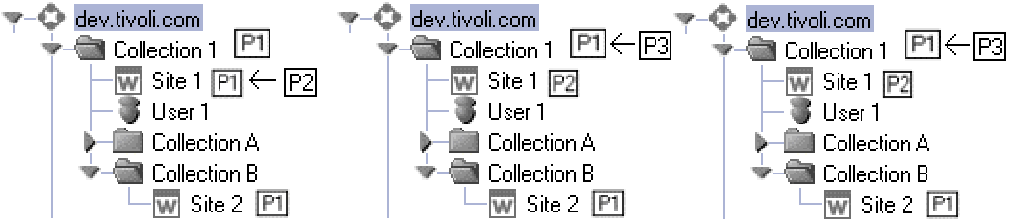

Ordinarily, if you remove policy from a resource, the appropriate default policy is applied to the resource. However, if you remove an overriding policy from a resource in a collection, the resource inherits the policy applied to it's parent collection. In the example above, if **P2** is removed from **Site1**, the site resource inherits **P3** from its parent collection.

To apply policy to a collection, complete the following steps from the Cross-Site console:

1. Select the **Resources** tab in either the Global or Security view. Your Cross-Site server domain is displayed in the tree on the left of the explorer.

2. Expand this domain by clicking on the arrow next to it in the tree. A list of resources and collections is displayed.

3. Select the collection to which you want to apply policy. The collection's configuration panels are displayed.

4. Select the **Policy** tab.

5. Select a policy from the **Intrusion Detection Policy** drop-down list. You can assign a policy for each type listed. All available policies are listed in each of the policy drop-down lists. Following is the list of policy types, which will vary depending on your configuration of Cross-Site:

   - Availability Agent Policy

   - Channel Policy

   - Site Scan Policy

   - Intrusion Detection Policy

   When you apply policy to a collection, the policy is applied only to the appropriate resources. For example, the Intrusion Detection policy is applied to all Security agents in the collection, and the Channel policy is applied to all Deployment channels in the collection.

6. Click **Apply** to save the collection's configuration settings and assign the policy to the collection, or click **Reload** button to reset the configuration panel.

When applied to a collection that contains Security agents, the Intrusion Detection policy takes effect the next time the agents in the collection contact the server. Security agents contact the Cross-Site server based on the interval specified for the *idlepurge* variable in the **ids.cfg** file. If there are policy updates available, each Security agent downloads these changes.

Security Tasks

Cross-Site provides *tasks* to enable you to collect information from Cross-Site resources. A task is an executable to which you assign various options and a schedule. This schedule may run the task one time or repeatedly. All tasks are listed in the tree on the **Tasks** explorer.

Each Cross-Site application provides predefined *task categories* from which you can create tasks. Think of the task categories as the explicit operations that you can perform with Cross-Site. You can see a list of *all* task categories in the tree on the **Tasks** explorer in the Global view. You can see a list of application-specific task categories in the respective application views, on the **Tasks** explorer.

Cross-Site for Security uses the upload task to retrieve information from Security agents. Each Security agent records all non-critical intrusion attempts as alerts in a local log file. Alert information is forwarded to an alert log on the Cross-Site management server, according to the priority of the alert. For example, critical alerts are sent to the server's alert log immediately, while non-critical alerts are forwarded periodically, according to the schedule set for the upload task. For more information on the type of information generated by Security agents, see "Storing Alerts in the Alert Log" on page 77.

In Cross-Site for Security, when you create a task, it runs as scheduled for all relevant resources. For example, when you create and run an upload task, the task runs and uploads information from all Security agents.

Using the Tasks Explorer

Tasks enable you to perform actions on Cross-Site resources. You can scan web sites for broken links using the Cross-Site for Availability task. You can upload information from Security agents using the Cross-Site for Security task. The Cross-Site for Deployment task enables you to publish channels, making applications available to end-users that have the Deployment agent installed. The **Tasks** explorer enables you to manage these tasks.

From the **Tasks** explorer you can do the following:

Create tasks by selecting a task type from the drop-down list and pressing **New** button shown here.

Each application provides predefined task categories from which you create tasks. Think of the task categories as all of the possible operations that you can do using Cross-Site. You can see a list of all task categories in the tree on the **Tasks** explorer in the Global view. You can see a list of application-specific task categories in the respective application views, on the **Tasks** explorer.

Delete tasks by selecting a task and clicking the **Delete** button.

Enable and disable tasks. By default, all tasks are enabled, or "turned on," when they are first created. To "turn off," or disable, a task's schedule, you can select the **Enabled** check box, thereby unchecking the box. The task is disabled until you click on the check box again

Terminate a running task by clicking the **Terminate Task** button. Terminating a task enables you to override the task's schedule without reconfiguring the task.

Start a task immediately by clicking the Start Task button.

For more detailed information about tasks and how to configure them, see the tasks section of the application user's guide.

Uploading Alerts from Agents

To create and run an upload task, complete the following steps from the Cross-Site console:

1. Select the Security or Global view icon. The explorers for the selected view are displayed below the view icons.

2. Select the **Tasks** explorer. A list of all tasks for the resources in the view is displayed.

3. Expand the domain in which you want to create the task.

4. Select **Agent Upload** from the drop-down list above the tree and click the **Create New** task button. A task configuration panel is displayed.

To edit a task, select it from the table on the **Tasks** explorer and click the **Edit** button.

To delete a task, select it from the table on the **Tasks** explorer and click the **Delete** button.

In general, you need to create only one task to upload alert information from all Security agents. This task notifies all agents to upload their data to the management server.

5. Select the **Info** tab to enter a name for the task. This panel briefly describes the purpose of the agent upload task.

6. Select the **Schedule** tab to specify when you want to run the upload task. You can schedule the task to run one time only, weekly, or monthly.

   Cross-Site stores time in Universal Time Coordinate (UTC), otherwise known as Greenwich mean time (GMT). This enables Cross-Site to display times that are meaningful to you, in your local time. All tasks, therefore, are scheduled and run according to your local time.

   a. Select **Run now** if you want to run the Agent Upload task when you click **Apply**.

   b. Select **Run once, later** if you want Security agents to upload their information to the management server only once. Use the drop-down calendar to specify a year, month, day, and exact time.

      To set the time that you want to run the task, enter values for the hour, minute, and second settings. Click on each part of the time setting to activate that section (hour, minute, second, and AM/PM), and use the up and down arrow keys on your keyboard to set each value. You can also edit each field directly.

      ▲▼ Use the up and down arrows to scroll through the lists of months and years.

      In the calendar, select a day of the month. After you select a day, it is set and the calendar closes.

   c. Select **Weekly** if you want Security agents to upload their data weekly. Click the check box next to the day or days you want to run the Agent Upload task.

      To set the time that you want to run the task, enter values for the hour, minute, and second settings. Click on each part of the time setting to activate that section (hour, minute, second, and AM/PM), and use the up and down arrow keys on your keyboard to set each value. You can also edit each field directly.

▲▼ Use the up and down arrows to scroll through the lists of months and years.

d. Select **Monthly** if you want Security agents to upload their data to the Cross-Site server on a certain day of the month. Enter the day of the month you want Cross-Site to run the Agent Upload task. You can schedule the task to run twice a month by entering a number in each of the two **Day** fields and adjusting the time for each day. For example, if you enter **1** in the first **Day** field and **15** in the second, Cross-Site will run the task on the first and the 15th of that month.

Note: *If the number entered is greater than the last day of a given month, the task will be run on the last day of that month. For example, if you schedule a task for the 31st, the task will run on the November 30th.*

e. Select **No scheduling** if you do not want to schedule the task now. The task is saved only. You can then run the task at any time by selecting the task from the table and clicking the **Start Task** button.

7. The **Summary** panel displays the following information for the selected task:

- The name of the task.
- The ID number of the task. Cross-Site creates a unique number for each new task you create.
- The task's definition ID, which represents the task's category. A Security Agent Upload task is represented by the number **3**.
- The target resource. Since this task uploads information from all Security agents, no specific targets are listed.
- The current state of the task. An upload task can be in any one of the following states:
    - Not scheduled
    - Scheduled for one time only
    - Scheduled for periodic execution
    - Running
    - Run ASAP
    - Stop ASAP
    - Completed

- The date and time the task will run next, if the task is scheduled, in the following format:
  Day Month DD, YYYY HH:mm:ss Time zone
  (For example, Wed Mar 03, 1999 14:13:25 CST)

- The task's last session ID. Cross-Site creates a unique number each time a task runs.

- The result of the task the last time it ran. This can be **SUCCESSFUL**, **FAILED**, or **CANCELED**.

- The date and time that the task completed the last time it ran. This is displayed in your local time, in the following format:
  Day Month DD, YYYY HH:mm:ss Time zone
  (For example, Wed Mar 03, 1999 14:13:25 CST)

- The amount of time (in seconds) the task took to complete the last time it ran.

8. Click the **Apply** button to save the task settings or click the **Reload** button to clear the fields on the configuration panel.

9. Close the configuration panel when you are done. You can also dock the panel for future use.

The task runs as scheduled on all Security agents. Once a task's scheduled time is reached, it waits for a heartbeat from each Security agent. The next time a Security agent contacts the Cross-Site server, the task runs and the agent's data is uploaded to the server.

Cross-Site enables you to run a task immediately. Select the task in the table on the **Tasks** explorer and click the **Start Task** button.

You can also stop a task that is currently running. To do so, select the task in the table on the **Tasks** explorer and click the **Terminate Task** button. The task is terminated immediately.

Security Events

An *event* is any critical alert, warning, error message, or status message generated by the Cross-Site management server and applications. Security events are generated when a policy applied to a Cross-Site for Security agent is violated. Policy violations occur when an incident that represents a possible intrusion is detected by a Security agent. These policy violations are also referred to as *incidents*.

Events are stored in the management repository and displayed in two places on the console: in the event log, which is available from the Admin view, and on the **Events** explorer in each application view. Cross-Site stores time in Universal Time Coordinate (UTC), otherwise known as Greenwich Mean Time (GMT). Each Security agent converts its event data to UTC before it uploads this data to the management server. This ensures that all events are displayed on the Cross-Site console in your local time.

You can use the event log to view a list of the most recent events generated by Cross-Site. This is a real-time log of events: events are displayed according to the event polling interval you specify. The event log is located in the Admin view. For information on viewing and customizing the event log, see “Viewing Real-time Events” on page 84.

You can use the **Events** explorer to view a historical list of all events, displayed by application and type. This explorer is located in all application-related views. You can also forward events from the **Events** explorer to email addresses, a Simple Network Management Protocol (SNMP) management application, or a Tivoli Enterprise Console (TEC). For more information, see “Managing the Event Service” on page 163.

The event icons at the bottom right of the Cross-Site console show the breakdown of events in the event log and **Events** explorer. The left-most icon indicates the number of critical events. The middle icon provides a count of the number of warning events. The right-most icon gives a count of status, or informational, events posted to the events log. These icons provide a visual indication of the number and type of events. (The number of events displayed will never exceed the maximum value set for the event log.)

Finally, you can analyze event data by using the report generator to create reports based on events and data collected from Cross-Site agents. For more information, see “Security Reports” on page 87.

Logging Alerts

Each time a Cross-Site for Security agent detects an incident, it logs an alert to a local log file. If the alert is assigned a priority of 1, it is considered *critical*. An alert typically has a priority of 1 if the traffic detected represents activity that you should not see under normal circumstances. The Security agent *immediately* sends critical alerts to the alert log on the Cross-Site management server. You designate which incidents should generate critical alert messages through *policy*, as described in “Changing the Priority of an Alert” on page 59. All other alerts are uploaded to the management server’s alert log periodically. See “Security Tasks” on page 69 for information on scheduling these alert log uploads using an upload *task*.

The alert log stores all alerts in the management repository. You can use the reports generator to create reports based on the alerts in the management repository. For more information, see “Security Reports” on page 87. Cross-Site stores time in Universal Time Coordinate (UTC), otherwise known as Greenwich Mean Time (GMT). Each Security agent converts its alert data to UTC before it uploads this data to the alert log. This ensures that alert data displayed in a report is in your local time.

Critical alerts are converted to events, forwarded to the Event service, and assigned an event priority. Critical alerts are treated as high priority events and are assigned an event priority of 5. Events are displayed in the event log in the Admin view as they occur. You can use the Cross-Site console to display a historical list of *all* events on the **Events** explorer in the Global and Security views. Remember, critical Security alerts are displayed as an event with a priority of 5 on the **Events** explorer and in the event log.

The following illustration shows the relationship between the Security agent and Cross-Site server, and explains how events are logged.

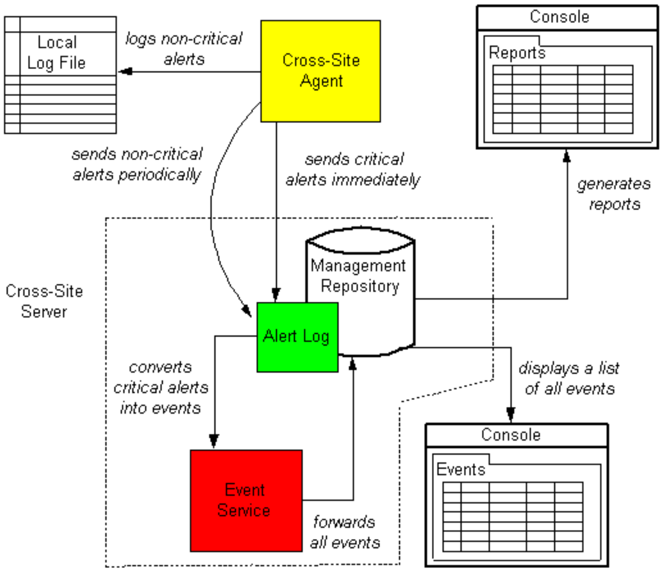

Storing Alerts in the Alert Log

The alert log is stored in the management repository. It records all alerts generated by the Cross-Site for Security agents. Information in the alert log is uploaded to the Cross-Site server by each Security agent, either when requested by the upload task or when sent as an event. Each entry in the alert log contains the following information:

- Alert log ID, which is a unique number assigned to the alert.
- Alert timestamp in the format of YYYY MM DD HH:mm:ss. Alert timestamps are displayed in your local time.
- Alert message ID, which is the number associated with the alert message. Alert messages and their associated IDs are configured as part of the policy assigned to a Security agent.
- Alert source ID, which is the universally unique identifier (UUID) of the Cross-Site agent that generated the alert

Other information, depending on the type of alert, is included in the alert log:

- The intrusion signature, which indicates the actual signature string that was detected along with any other information that may provide additional context for the signature (a predetermined number of bytes of data are captured).

- The user name of the person logged on to the computer or port on which the incident was detected. (This information is tracked for authenticated non-HTTP protocols only.)
- The IP address of the source of the incident. The source of an incident is the computer that performed, or attempted to perform, a transaction that signaled a potential intrusion.
- The IP address of the computer on which the signature was detected. This computer is often referred to as the destination computer or target.
- The port number of the source computer.
- The port number of the destination computer.
- The user ID of the user who performed the transaction that initiated the incident. [This information is tracked for Network File System (NFS) attacks only.]
- The group ID of the user who performed the transaction that initiated the incident. [This information is tracked for Network File System (NFS) attacks only.]

Types of Security Events

The following is a list of the events that are generated by Cross-Site for Security. They are displayed in the **Events** explorer and in the event log.

Security agent policy modified: *X*

This event is generated when a Cross-Site user changes a Security policy. The *X* variable identifies the policy ID. The priority of this event type is status.

Security agent *X* unknown

This event is generated when an agent’s heartbeat is detected by the management server, but the server has no record of the agent. The *X* variable identifies the agent ID. This could mean that another Cross-Site user deleted the agent from the Cross-Site server but has yet to remove the agent from the client machine. The priority of this event type is warning.

Security agent *X* not responding

This event is generated when the management server does not receive a heartbeat from an agent for more than ten minutes. The *X* variable identifies the agent ID. This could signal that the connection between the Cross-Site server and the Security agent has been compromised. The priority of this event type is critical.

Security agent X response resumed

This event is generated when an agent contacts the management server after a **Security agent *X* not responding** event was generated for that agent. The *X* variable identifies the agent ID. This event indicates that contact between the management server and agent has resumed. The priority of this event type is status.

Critical Security alert: *X*

This event is generated when a Security agent detects a serious incident that requires immediate attention. The *X* variable identifies the alert ID. The priority of this event type is critical.

You designate which incidents should generate critical alert messages through policy, as described in “Changing the Priority of an Alert” on page 59.

Upload task for Security agent failed: *X*

This event is generated when the task that notifies the Security agent to upload its alert data fails. In general, this event is generated when an internal error occurs, such as if the event service is unavailable. The *X* variable identifies the task ID. The priority of this event type is status.

Critical alerts are sent to the Cross-Site server where they are converted to events. Critical, or priority 1, alerts are displayed on the **Events** explorer and in the event log as critical events.

Types of Global Events

In addition to events generated by Cross-Site for Security, the management server generates events for services. The following is a list of all service events, which are displayed in the **Global** group on the **Events** explorer in the Global view.

Task manager cannot reschedule task *X*

This event is generated when an internal error occurs or if the application finds an environmental or resource constraint. The *X* variable identifies the task ID. The priority of this event type is critical.

Task manager cannot launch task *X*

This event is generated when an internal error occurs or if the application finds an environmental or resource constraint. The *X* variable identifies the task ID. The priority of this event type is critical.

Task manager successfully launched task *X*

This event is generated when the task identified by *X*, which is the task ID, launches successfully. The priority of this event type is status.

Task manager *X*

This event is generated when the management server starts, which in turn starts the Task manager, or when the web server shuts down, causing the Task manager to stop running. The *X* variable identifies stopped or started. The priority of this event type is status.

Task manager failed: *X*

This event is generated if the Task manager cannot access the database or if the task that the Task manager was attempting to start no longer exists. The *X* variable identifies the error. The priority of this event type is critical.

Authorization failure: *X*

This event is generated when the management server fails to authorize a client request. Events of this type are generated when an invalid password was specified, a user does not have permission to use the requested service, or any other authorization request is invalid. The *X* variable prints a message indicating the error. The priority of this event type is critical or warning, depending on the exception.

Authorization service succeeded: *X*

This event is generated when authorization or authentication succeeds. Events of this type are generated, for example, when a user is successfully created, roles are successfully set for a user, an agent is created, or a management server is defined. The *X* variable identifies the successful operation. The priority of this event type is status.

Authorization service failed: *X*

This event is generated when the authorization or authentication service fails. Events of this type are generated when an invalid principal or password is specified, an exception occurs during the creation of a principal or group, an attempt to delete a principal fails, or any other operation on a principal or role fails. The *X* variable identifies the error. The priority of this event type is critical.

Server authentication failed due to exception *X*

This event is generated when authentication fails while defining a foreign management server (server-to-server communication). The *X* variable identifies the error. The priority of this event type is critical.

The SMTP service failed to startup due to exception *X*

This event is generated when an internal error occurs or if the service finds an environmental or resource constraint while starting. The *X* variable identifies the error. The priority of this event type is critical.

The SNMP service failed to startup due to exception *X*

This event is generated when an internal error occurs or if the service finds an environmental or resource constraint while starting. The *X* variable identifies the error. The priority of this event type is critical.

The TEC service failed to startup due to exception *X*

This event is generated when an internal error occurs or if the service finds an environmental or resource constraint while starting. The *X* variable identifies the error. The priority of this event type is critical.

The Channel manager failed to startup

This event is generated when the Channel manager fails to startup. The Channel manager is responsible for publishing channels so that Deployment agents can run them. It is also responsible for managing the channels that update Cross-Site agents and the management server. The priority of this event type is status when the management server is shutting down but critical in all other situations.

Task *X* completed with status: SUCCESSFUL/FAILED

This event is generated when the task identified by *X*, which is the task ID, completes successfully or when the task fails. The priority of this event type is status.

Events are displayed on the **Events** explorer of each application-specific view and in the event log in the Admin view.

Viewing Events on the Events Explorer

You can view Cross-Site events using either the **Events** explorer or the event log. The **Events** explorer lists all events, old and new. The event log, which is displayed in the Admin view, lists real-time events. For more information about viewing real-time events, see “Viewing Real-time Events” on page 84.

The tree on the left of the **Events** explorer lists each event type, organized by application group. You can view all events generated by the event service and stored in the management repository on the **Events** explorer. Such a list of events is useful in creating policy for Security agents, tracking problems, and troubleshooting. Also, it enables you to view a historical listing of events stored in the management repository, one event type at a time.

From the **Events** explorer you can do the following:

- View an event’s details by double-clicking on it in the table.
- Configure an event’s forwarding information by selecting the **Config** tab on the right of the **Events** explorer.

- Retrieve events from the management repository by using either the **Auto-load** check box or the **Fetch Events** button shown here. The **Auto-load** check box forces the console to retrieve events from the management repository each time you click on an event type.
- You can also manually retrieve events by selecting an event type and clicking the **Fetch Events** button.

Configuring the Events Explorer

You can choose to retrieve events from the management repository using either the **Auto-load** check box or the **Fetch Events** button. The **Auto-load** check box, which is located at the top of the explorer, forces the console to retrieve events from the management repository each time you click on an application group or event type. Automatically loading events ensures that the list of events on the **Events** explorer is always current.

If you are concerned about console performance, you can use the **Fetch Events** button to control when the console retrieves events from the management repository. You can manually retrieve events by selecting an application group or event type and clicking the **Fetch Events** button. The **Fetch Events** button is located on the **Events** panel on the right of the explorer.

You can limit the events displayed in the **Events** explorer to those that were generated during a specific time period.

1. Specify a time and date, using the **From** drop-down calendar, for the beginning of the time period.

   To specify a date and time with the calendar, click on each part of the time setting (hour, minute, second, and AM/PM), and use the arrows to set each value. After you select a day, the time is set and the calendar closes.

2. Click the **To** button to toggle between **latest** and another drop-down calendar.

   If you choose **latest**, the explorer displays all events generated after the time and date specified on the **From** calendar.

   If you choose a date using the drop-down calendar, only events generated between the **From** and **To** dates are displayed.

You can also enter a number in the **Max Number of Events Shown** field to specify the maximum number of events that can be displayed on the explorer.

Viewing Events

To view events on the **Events** explorer, perform the following steps from the Cross-Site console:

1. Select the Security or Global view by clicking its icon button at the top of the console. The explorers and panels of the selected view are displayed.

2. Select the **Events** tab.

3. From the tree on the left side of the explorer, expand an application group in a domain to display a list of event types.

4. Select an event type in the tree to view a list of those generated and stored in the management repository. The events are displayed in the table on the right of the explorer. The following information is displayed for each event in the table:

    - The service or application that generated the event

    - The time and date that the event was generated

    - The priority of the event. The priority icons correspond to the event icons at the bottom of the console. A red event icon indicates a critical event. A yellow event icon indicates a warning event, and a green event icon indicates a status event.

    - The IP address of the Cross-Site server or agent that generated the event

    - The event message

5. To view additional details about an event, double-click on an event in the table. A floating **Event details** panel is displayed. Event information is displayed on the **General** panel as follows:

    - The log ID, which is auto-generated and indicates the ID of the event in the management repository

    - Status, which indicates whether the event was received during the current session of the console (**New**) or if it was an existing event (**Prev**)

    - Priority, which includes **0**–**9** for critical events, **10**–**19** for warning events, or **20**–**29** for status events

    - When occurred, which indicates the last occurrence of the event and is in the following format:
      Day of the week Month DD HH:mm:ss Time zone YYYY
      (e.g., Wed Mar 03 14:13:25 CST 1999)

- Source, which indicates the IP address of the Cross-Site agent or server on which the event occurred
- The plain-text message that explains the event

The **Raw info** panel includes the following information:

- The event ID, which identifies the severity of the event, the facility that generated it, and its unique ID
- The ID of the facility that generated the event, as follows.

  2 indicates the Cross-Site server and its services.

  4 indicates Cross-Site for Security.

  6 indicates Cross-Site for Availability.

  8 indicates Cross-Site for Deployment.
- An annotation that summarizes the event

To close the **Event details** panel, click the **Close** button at the top, right-hand side of the panel.

6. Optionally, you can configure an event's forwarding information by selecting the **Config** tab on the right of the **Events** explorer. You cannot forward individual events, but rather types of events. See "Managing the Event Service" on page 163 for more information on forwarding events and how the event service works.

For more detailed information about the various types of events, see "Types of Security Events" on page 78 and "Types of Global Events" on page 79.

Viewing Real-time Events

The event log, which is located in the Admin view, displays a list of the most recent events generated by Cross-Site. This is a real-time log of Cross-Site events and contains events generated during the current console session. You can configure the event log in the following ways:

- Specify the polling interval, which sets how often the Cross-Site console contacts the management repository to check for new events.
- Clear the event log of old events.

- Configure the event log to display only a certain number of events at one time.
- Filter the event log to display events generated by a particular application.

Perform all of these procedures from the event log in the Admin view.

To configure the event log, complete the following steps from the Cross-Site console:

1. Select the Admin view by selecting the **Admin** icon button at the top of the console.
2. Select the **Event Log** tab.
3. To set the polling interval of the console, click the **Update** button (labeled **60** by default). Specify a value, in seconds, in the **Input** panel that is displayed and press the **OK** button.
4. Click the **Update Events Now** button to check for new events immediately.
5. To remove the events listed in the event log, click the **Clear All Event**s button.
6. To set the maximum number of events you want to store and display in the event log, click the **Events Stored** button (labeled **400** by default). Specify a value in the **Input** panel that is displayed and click the **OK** button.

   When the event log contains the maximum number of events, new events generated by Cross-Site begin replacing old events using the “first in, first out” method.
7. To filter the type of events you want to view in the log, select one or more of the application icons. To view events generated by all applications, select all of the icons provided. By default, all application icons are selected.
8. To view detailed information about an event, double-click on it in the table. Event information is displayed on the **General** panel as follows:
   - The log ID, which is auto-generated and indicates the ID of the event in the management repository
   - Status, which indicates whether the event was received during the current session of the console (**New**) or if it was an existing event (**Prev**)
   - Priority, which includes **0**–**9** for critical events, **10**–**19** for warning events, or **20**–**29** for status events
   - When occurred, which indicates the last occurrence of the event and is in the following format:
     Day of the week Month DD HH:mm:ss Time zone YYYY
     (e.g., Wed Mar 03 14:13:25 CST 1999)

- Source, which indicates the IP address of the Cross-Site agent or server on which the event occurred
- The plain-text message that explains the event

The **Raw info** panel includes the following information:

- The event ID, which identifies the severity of the event, the facility that generated it, and its unique ID
- The ID of the facility that generated the event, as follows.

  2 indicates the Cross-Site server and its services.

  4 indicates Cross-Site for Security.

  6 indicates Cross-Site for Availability.

  8 indicates Cross-Site for Deployment.
- An annotation that summarizes the event

To close the **Event details** panel, click the **Close** button at the top, right-hand side of the panel.

The console checks for new events according to the interval you specified. New events are displayed based on the maximum number you specified and the filters you applied.

Security Reports

Cross-Site enables you to generate reports based on the data collected by Cross-Site agents. In the case of Security, reports are generated based on the security alerts that occur for each Cross-Site for Security agent. The alert information is stored in the Cross-Site server's management repository. When you generate a report, the Cross-Site server retrieves the data from the management repository. The report generator then generates and displays the report on the console.

Several types of reports are provided by Cross-Site for Security, including graphs and table-based reports. You simply need to select the report type, agents, and the time period on which to base the report. This section describes the types of reports you can generate and how to generate them.

You can generate multiple reports and display as many as you like on the console. You can also generate multiple copies of the same report. This enables you to compare data as you refine it. You can also float and dock reports in order to compare them.

Overview of Security Reports

Each time a Security agent detects an incident, it generates an alert. It saves each alert in a local log file. If the alert is assigned a priority of 1, it is considered critical. The Security agent *immediately* sends critical alerts to the alert log on the Cross-Site management server. All other alerts are uploaded to the management server's alert log periodically. See "Security Tasks" on page 69 for information on scheduling these alert log uploads using an upload task.

The illustration on the following page shows the relationship between the Security agent and Cross-Site server, and explains how events are sent to the alert log in the management repository. You can then generate reports based on the information in the management repository.

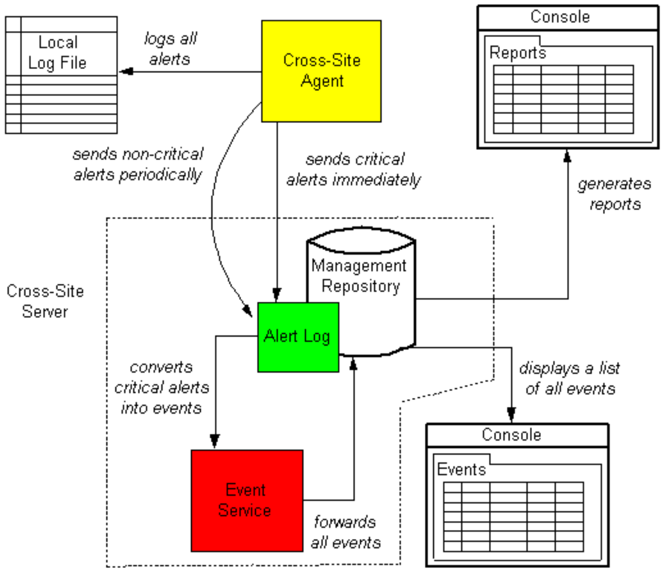

To view data collected from Security agents, you can generate reports based on the following information in the management repository:

- Types of incidents that occurred
- Severity of the alerts that were generated by the incidents
- Group ID of the person who initiated the incident. This information is tracked for Network File System (NFS) attacks only.
- User ID of the person who initiated the incident. This information is tracked for NFS attacks only.
- Port number used by the computer for which the potential attack was intended. (This is also referred to as the destination host.)
- IP address of the destination host
- Port number used by the originating host (computer) of the incident. (This is also referred to as the source host.)
- IP address of the source host
- User name of the person logged onto the destination host computer or port. This information is tracked for authenticated non-HTTP protocols only.
- The intrusion signature that was detected
- Universally unique identifier (UUID) of the Cross-Site agent that generated the alert based on the incident

- Date and time each incident occurred. This is displayed in your local time.
- ID associated with the alert message. Alert messages are configured as part of the policy assigned to a Security agent.
- Number of incidents, which represent possible intrusion attempts, that occurred

You can use this information to decide whether the incident is a false alarm or is a potential intrusion attempt.

Cross-Site for Security provides several types of reports that enable you to view a summary of the information in the management repository, organized according to different variables. Two of the most important variables are the source host and destination host of each incident.

The source host represents the computer that may have performed, or attempted to perform, the transaction that signaled a potential intrusion. However, in "spoofing" attacks, an attacker replaces the source host IP address with another host’s IP address. In these attacks, the source host detected by Cross-Site may not be the actual origin of the attack. The destination host is the computer for which the potential attack was intended. This information is in the packet traffic detected by the Cross-Site for Security agent.

You can generate the following types of reports to view Security agent data:

Incident VS Time

A graph of security incidents by hour, for each day in the specified date range. Data is displayed for every hour of every day in which incidents were detected.

Severity VS Time

A graph of security incidents, organized by severity and time, for each day in the specified date range. Incident severities range from 1 to 5, with 1 being the most severe. Data is displayed for every hour of every day in which incidents were detected.

Incident VS Source Host

A graph of incidents that displays the IP address of each source host that initiated an incident, and the number of incidents initiated by each source host

Incident VS Destination Host

A graph of incidents that displays the IP address of each destination host for which an attack was intended, and the number of attacks intended for each destination host

Destination Host VS Source Host
A graph of incidents that displays the IP address of each destination host for which an attack was intended, and the source host that initiated each attack

Source Host VS Destination Host
A graph of incidents that displays the IP address of each source host that initiated an attack, and each destination host for which an attack was intended

Source Host VS Incident Type
A graph of incidents that displays the total number of times each type of incident occurred, and the IP address of each source hosts that initiated an incident

Time Table A table that lists incidents, including the date and time each incident occurred. For each incident, the IP addresses of the source host and destination host are also listed.

Target Table A table that lists incidents, including the IP addresses of the source host and destination host of each incident. The total number of times the same incident occurred is also displayed.

Authentication Services
A table that lists incidents, including the user names and IDs associated with each incident. The time of each incident, and the IP addresses of the source host and destination host of each incident are also listed.

User Defined Columns
A table that lists incidents and the columns you select. You select the columns to be displayed, according to the type of data you want to view.

Information that is not accessible with the other Security reports is available when you select this report type. For example, you can generate a report based on the port numbers of both the source host and the destination host.

For instruction on how to generate these Security reports, see “Generating Security Reports” on page 91. You can refine report results as discussed in “Refining Reports” on page 94.

Generating Security Reports

You generate all Security reports using the same procedure. The type of report you select, along with the agents and time period you specify, determine what information is displayed in the report. For descriptions of the different types of Security reports, see "Overview of Security Reports" on page 87.

To generate a Security report, complete the following steps from the Cross-Site console:

1. From the **Resources** explorer in the Security or Global view, click the **Generate Reports** button. The **Report Generator** panel is displayed, floating.

2. Select a report type from the drop-down list. A description of the report is displayed in the **Description** area to the right when you rest the cursor on the name of the report. After you select a report type, the settings for the report are displayed.

   Reports are organized by the application to which they apply. For example, reports for Security are listed under the **Security** icon in the report tree.

3. Specify a time and date in the **Start** and **End Time** fields. The default time range ends at the current time and starts 24 hours earlier. For both fields, use the drop-down calendar to set the dates and times. Do not edit the fields directly.

   Note: *Cross-Site stores time in Universal Time Coordinate (UTC), otherwise known as Greenwich Mean Time (GMT). Each Security agent converts its alert data to UTC before it uploads this data to the alert log. Therefore, the time period you specify is the local time zone of your machine. This ensures that alert data in the report is displayed in your local time.*

   To specify a date and time with the calendar, click on each part of the time setting (hour, minute, second, and AM/PM), and use the arrows to set each value. After you select a day, the time is set and the calendar closes.

   Tivoli recommends that you start with short intervals of time to avoid retrieving too much information from the management repository at once.

4. From the drop-down list in the Resource area, select an agent on which to base the report. This drop-down list is a tree structure, organized by domain. When you choose an agent from the drop-down list, it is added to the list box below. Choose as many agents as you like.

   You can also select a collection from the drop-down list. When you select a collection, the report is based on all agents in the collection.

   Select the **all** option to generate a report based on all Cross-Site for Security agents.

5. If you selected the **User Defined Columns** report type from the drop-down list, select the columns you want displayed in your table report from the **Table Columns** list. To select multiple columns at once, press the **Ctrl** key or the **Shift** key while selecting column headings in the list. **Ctrl** enables you to select multiple items individually. **Shift** allows you to select a series of items. You can display any combination of available table columns, or select all of them.

   The following lists the available table column headings and a brief description of what each heading represents:

   Incident Type
   The type of incident that occurred

   Severity The severity level of the alert generated by the incident, ranging from 1 to 5, with 1 being the most severe

   Group ID The group ID of the person who initiated the incident. This is part of the credentials presented during an Network File System (NFS) transaction.

   User ID The user ID of the person who initiated the incident. This is part of the credentials presented during an NFS transaction.

   Dest Port The port of the destination host for which the attack was intended

   Dest IP The IP address of the destination host

   Src Port The port of the computer what was the source of the incident

   Src IP The IP address of the source of the incident

   User The user name of the person logged onto the destination host computer or port for which the attack was intended. (This information is tracked for authenticated non-HTTP protocols only.)

   Signature The intrusion signature. This indicates the actual signature string that was detected along with any other information that may provide additional context for the signature.

   Agent ID The universally unique identifier (UUID) of the Cross-Site agent that generated the alert.

   Time The time the incident occurred, in the format:
   Day of the week Month DD HH:mm:ss Time Zone YYYY.
   For example, Mon Mar 01 09:15:26 CST 1999. This is always displayed in your local time.

Alert Msg ID The number associated with the alert message. Alert messages are configured as part of the policy assigned to a Security agent.

Count The total number of times this particular type of incident (with the same parameters) has occurred, within the specified time range.

The **Count** heading condenses the length of the table. Without it, the same item would be listed in the table each time it occurred. **Count** is displayed in every **User Defined Columns** report by default. This column heading does not appear in the **Table Columns** list.

6. You can specify the maximum number of results you want displayed for the report in the **Max Result Set** field. This option prevents you from trying to retrieve too much information from the management repository at once (which could affect the performance of both the console and the network), and enables you to focus on more manageable subsets of information.

   If you specify a value on the **Max Result Set** field (or accept the default value) and results beyond this maximum setting are available, a note is displayed in the top right corner of the report to inform you that the data has been limited, or “cropped,” to suit your maximum result setting.

   If you enter zero as the value or leave the **Max Result Set** field blank, all relevant data that is available in the management repository is displayed in the report.

7. Click the **Generate** button to generate the report. The report is displayed on a panel on the console. You may have to move the **Report Generator** panel in order to view the report.

When you generate a table report, the total number of rows and columns in the table is displayed.

When you generate a graph report, you may have to enlarge the graph to see all of each IP address. See “Refining Reports” on page 94 for instructions.

If you incorrectly or incompletely filled out the report generator, an error message is displayed. See “Refining Reports” on page 94 to change the chart type or data set of the report to correct the problem.

Each type of incident is displayed in a different color. From the report’s **Display Control** panel, you can choose to display a legend that lists the color of each type of incident.

Note: *The **Display Control** panel is not available for the User Defined Columns report type.*

You can save the report as an HTML or text file. See “Saving Reports” on page 94 for instructions.

Refining Reports

You can refine the results displayed in a report using the report's **Data Control** panel. You can choose a different resource on which to base the report, or change the start and end times for the report. You can also specify a maximum result setting which also limits the amount of data displayed in the report. Click the **Update** button to update the data on the report's **Display** panel.

Note: *The **Data Control** panel is not available for the User Defined Columns report type.*

When you generate graph reports, you can change how the graph is displayed in each report in the following ways:

- To enlarge a set of data, press the **Shift** key while selecting and dragging the cursor over the area of the graph. Press the **R** key to reset the view.
- To pan across a wide graph, press the **Ctrl** key while clicking anywhere on the graph. Then drag the mouse to the left or right to view the rest of the graph. Press the **R** key to reset the view. It may take a few seconds to pan the graph.
- To change the type of graph used to display the report results, select the report's **Display Control** tab. The available graph types include the following:
    - Stacked Bar
    - Scatter Plot
    - Pie
    - Plot
    - Bar

  You can also specify to display a legend and to display three-dimensional bar and pie chart reports.

Changes you make to the options in the **Display Control** panel are reflected when you return to the **Display** panel.

Saving Reports

After you generate a report, Cross-Site enables you to save it as a file, which you can then print. You can save a report as an HTML file or a text file. You can also choose whether you want to separate the fields in the text file with tabs or commas.

Text files that are delimited by tabs or commas can then be imported into another program, such as Microsoft Excel. To save a graph report type as an image, select **HTML** as the file type. Otherwise, the data that comprises the report is saved only.

Complete the following steps to save a report:

1. With the report displayed on the console, click the Save to File button. The **Save Report** panel is displayed.

2. Select a type of file. The default file name that is displayed in the **File** field changes according to the type of file you select.

3. In the **File** field, enter or browse for a file name.

4. Click the **Save** button or click the **Cancel** button to close the **Save Report** panel without saving the report.

After you save a report, you can view and print the report in any browser or text editor.

4

Cross-Site for Security Features

Cross-Site for Security offers a variety of features you can use to monitor your network for unusual access patterns that may represent hacker activity. You may choose to use various combinations of these features, depending on your organization’s network configuration and your security needs.

The intrusion-detection features provided with Cross-Site for Security are discussed in this section of the guide. Specifically, the Security agent detects or monitors the following:

- Scans and floods
- IP traffic
- Port services
- Domain Name System (DNS), mount service, and network file system (NFS) requests and replies
- Portmapper service request and reply dumps
- RStatd calls
- Requests for specific map names and file names
- SMB-based attacks on PC file servers
- Internet control message protocol (ICMP) attacks

In addition to the features listed above, the Security agent (UNIX only) can monitor ONC-based RPC protocols. Most of these features include RPC Service Notify ACL alerts along with specific Notify ACL-enabled rules for each of the calls.

Detecting Scans and Floods

Typically, the first task of an intruder is to determine where the holes on a network exist using scanning techniques. The intruder scans a network for the systems (IP addresses) and ports that respond to connection requests. Cross-Site for Security monitors your network, and recognizes this type of scan behavior.

Cross-Site for Security also monitors your network for flood attacks. Flood attacks are considered “classic” Denial of Service (DoS) attacks because they do not directly harm the target system, but make it difficult for other systems to use the target system. The Cross-Site for Security agent generates an event whenever it detects scan and flood activity.

Following is a list of some of the DoS attacks that Cross-Site for Security agents detect:

- teardrop—IP Fragmentation attack
- newtear—IP Fragmentation attack
- bonk—IP Fragmentation attack
- syndrop—IP Fragmentation attack
- jolt—Possible Ping of Death (POD) attack

The following topics are discussed in this section:

- ICMP Ping Sweeps
- TCP and UDP Port Scanning
- Flood Detection
- Scan and Flood Summary
- Scan and Flood Summary

ICMP Ping Sweeps

Internet Control Message Protocol (ICMP) ping sweeps find systems on a network by sending a sequence of pings to IP addresses on the network, and looking for the IP addresses that respond.

If the scan is *fast* and hits many IP addresses, it is more likely an intruder using a programmed ping sweeper. A system administrator legitimately pinging various machines manually would naturally be much slower.

A *slow* scan could also indicate an intruder, one who is aware of typical fast-window scan-detection algorithms and is therefore purposely scanning slowly, so as not to trigger a scanning alert.

TCP and UDP Port Scanning

A hacker can perform various scans over both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports. A *wide* scan is executed on a fixed and well known port over several IP addresses. A *narrow* scan is performed over different ports on the same system.

Scans can also be done slowly and stealthily using TCP packets, such as SYN, FIN, and RST, in a non-standard way. For example, a hacker could find out if a daemon is running by taking advantage of the way the TCP protocol works.

Following is a list of the different types of scans that Cross-Site for Security agents detect. The parameters, which appear in italics, for each of the scans are set in the **ids.cfg** file. Depending on the type of scan, see either the **Fast Scan Values** section or the **Slow Scan Values** section for a list of the appropriate parameters and their values. You can edit these values for each Security policy you create. For more information, see “Configuring Scan and Flood Parameters” on page 63.

TCP Wide Scans

The Security agent generates a TCP Wide Scan alert if a single host attempts a TCP connection to more than *ehostmax* hosts in *epochtime* seconds.

TCP Wide Slow Scans

The Security agent generates a TCP Wide Slow Scan alert if a single host attempts a TCP connection to more than *whostmax* hosts in *wintime* seconds.

TCP Narrow Scans

The Security agent generates a TCP Narrow Scan alert if a single host attempts a TCP connection to more than *etcpscanmax* ports on the same system within *epochtime* seconds.

UDP Narrow Scans

The Security agent generates a UDP Narrow Scan alert if a single host attempts a UDP connection to more than *eudpscanmax* ports on the same system within *epochtime* seconds.

UDP Narrow Slow Scans

The Security agent generates a UDP Narrow Slow Scan alert if a single host attempts a UDP connection to more than *wudpscanmax* ports on the same system within *wintime* seconds.

A mixture of scans that are not associated with a specific packet type can also generate a generic alert. For example, if an intruder attempts a classic port scan using stealthy techniques that use RST or FIN packets to scan a given host, the Security agent generates a TCP Narrow Scan alert if more than *etcpscanmax* ports on the system are scanned within *epochtime* seconds.

Flood Detection

Even if an intruder cannot directly affect a system, he or she can still interfere with a system through a Denial Of Service (DoS) attack. A DoS attack is done by *flooding* a port with connection requests or other types of packets. Cross-Site for Security can identify this type of attack and generate an alert when an attack is detected.

The parameters for each of the following floods are set in the **Fast Scan Values** section of the **ids.cfg** file:

SYN Floods
The Security agent generates a SYN Flood alert if a single host sends more than *esynmax* packets to the same port in *epochtime* seconds.

RST Floods
The Security agent generates an RST Flood alert if a single host sends more than *erstmax* packets to the same port in *epochtime* seconds.

FIN Floods
The Security agent generates a FIN Flood alert if a single host sends more than *efinmax* packets to the same port in *epochtime* seconds.

UDP Floods
The Security agent generates a UDP Flood alert if a single host sends more than *eudpmax* packets to the same port in *epochtime* seconds.

Scan and Flood Summary

The following table provides a summary of the parameters used to configure alerts for each type of scan or flood detected. Each type of scan is defined in “TCP and UDP Port Scanning” on page 100. You can edit the values defined for these parameters in the general configuration file (**ids.cfg**).

The Fast Scan and Slow Scan speed parameters are shared across the different types of scans. For example, a TCP Wide Fast scan and a TCP Narrow Fast scan both use the *epochtime* parameter, as well as their own parameters for number of hosts (*ehostmax*) and number of ports (*etcpscanmax*).

Scan Speed	TCP Wide	TCP Narrow	UDP Narrow	SYN Flood	RST Flood	FIN Flood	UDP Flood
Fast Scan epochtime	ehostmax	etcpscanmax	eudpscanmax	esynmax	erstmax	efinmax	eudpmax
Slow Scan wintime	whostmax	wtcpscanmax	wudpscanmax	n/a	n/a	n/a	n/a

Flood Squelching

Each flooding attack can result in a considerable number of alerts being sent to the server, which, in turn, can affect system availability. Whenever the server receives five alerts having the same traits (basically describing the same attack), all subsequent alerts issued with the same traits are put in the local alert log. The additional alerts are available for retrieval, if necessary.

Monitoring Network Traffic

Because hackers can target specific addresses (e.g., IP, FTP, DNS) and subnets, you may want to set monitoring options for network traffic. You can configure Cross-Site for Security agents to send events to the Cross-Site management server when they detect packets from a particular host or network by setting the commands in the appropriate sections the **ids.rules** file.

The following commands are discussed in this section:

- The IGNORE Command
- The NOTIFY Command
- The ALLOW Command

The IGNORE Command

Use the IGNORE command to instruct a Security agent to ignore packets from a specific host or entire network. This command takes precedence over the NOTIFY and ALLOW commands. The IGNORE command is particularly useful for filtering out packets that are generated by internally deployed network scanners or by machines that are simply “noisy”, like a backup server.

The IGNORE command enables you to tell the agent to ignore all traffic from or to any specified host or network. While this is a useful command to use if several false alerts are being generated, it should be used sparingly; this command effectively creates a “blind spot” where the agent does not check for intrusions from the specified host or network. You can apply the IGNORE command to a single host or an entire network.

The syntax of the IGNORE command is as follows:

```
IGNORE {SRC|DST|SRCDST} {NET|HOST ip_addr} [NETMASK]
```

where:

IGNORE Specifies the IGNORE command.

SRC Causes the agent to ignore packets coming from the specified network or host.

DST Causes the agent to ignore packets going to the specified network or host.

SRCDST Causes the agent to ignore packets coming from *and* going to the specified network or host.

NET Specifies that the network the agent is to ignore packets coming from (or going to).

HOST Specifies that the host the agent is to ignore packets coming from (or going to).

ip_addr Indicates a single machine's IP address if used with the HOST parameter, or the network IP address if used with the NET parameter.

NETMASK The netmask of the network for which the agent is to ignore packets coming from or going to (e.g., 255.255.255.0).

The following example tells the Cross-Site agent to ignore packet traffic coming from the IP address 129.34.40.176.

```
IGNORE SRCDST HOST 129.34.40.176
```

The NOTIFY Command

Use this command to generate an alert when a Cross-Site agent receives a packet from the specified host or network.

The syntax of the NOTIFY command is as follows:

```
NOTIFY [PEAK|OFFPEAK] [ANY|NEVER] {NET|HOST ip_addr} [NETMASK] msg_id
priority
```

where:

NOTIFY Specifies the NOTIFY command.

PEAK Specifies to generate an alert if packets are received during peak hours.

OFFPEAK Specifies to generate an alert if packets are received during off-peak hours.

`ANY` Specifies to generate an alert if packets are received any time.

`NEVER` Specifies to never generate an alert if an IP packet is received.

`NET` Causes the agent to send an alert when packets are received from the specified network.

`HOST` Causes the agent to send an alert when packets are received from the specified IP address.

`ip_addr` Indicates a single machine's IP address if used with the HOST parameter, or the network IP address if used with the NET parameter.

`NETMASK` Specifies the netmask of the network for which the agent is to generate an alert (e.g., 255.255.255.0).

`msg_id` Specifies the ID number associated with the alert message.

`priority` Specifies the priority of the alert. This should be the same priority indicated in the console.

The following example instructs the Cross-Site agent to issue an alert when traffic is detected from the network 9.2.13.0.

```
NOTIFY NET 9.2.13.0 255.255.255.0 1954 1
```

Note: *You can NOTIFY an entire network and ALLOW specific hosts and you can ALLOW an entire network and NOTIFY specific hosts.*

The ALLOW Command

The ALLOW command permits a specified host or network to send packets without an agent issuing the corresponding alert. This is useful for allowing local network traffic and traffic from trusted networks, while alerting on any other traffic.

The syntax of the ALLOW command is as follows:

```
ALLOW [PEAK|OFFPEAK] [ANY|NEVER] {NET|HOST ip_addr} [NETMASK] msg_id
priority
```

where:

`ALLOW` Specifies the ALLOW command.

`PEAK` Specifies not to generate an alert if packets are received during peak hours.

`OFFPEAK` Specifies to generate an alert if packets are received during off-peak hours.

ANY Specifies to generate an alert if packets are received any time.

NEVER Specifies to never generate an alert if packets are received.

NET Instructs the agent not to send an alert when packets are received from the specified network.

HOST Instructs the agent not to send an alert when packets are received from the specified host IP address.

ip_addr Indicates a single machine’s IP address if used with the HOST parameter, or the network IP address if used with the NET parameter.

NETMASK Specifies the netmask of the network for which the agent is not to generate an alert (e.g., 255.255.255.0).

msg_id Specifies the ID number associated with the alert message.

priority Specifies the priority of the alert.

Note: *You can NOTIFY an entire network and ALLOW specific hosts and you can ALLOW an entire network and NOTIFY specific hosts.*

In the following example, the ALLOW rule instructs the Cross-Site agent *not* to issue an alert when traffic is detected locally, or from your business partners' networks. In this scenario, business partner A could represent a supplier, while business partner B could be a distributor. Any other traffic detected from outside of these three networks will cause the Cross-Site agent to generate an alert.

```
ALLOW NET 7.8.0.0 255.255.0.0 1000 1 (business partner A)

ALLOW NET 1.2.304.0 255.255.0.0 1001 1 (business partner B)

ALLOW NET 263.28.0 255.255.0.0 1002 1 (your company)
```

The ALLOW rule searches packet traffic in the order specified by the rule. It continues to be processed until the Cross-Site agent encounters the last ALLOW statement or an applicable NOTIFY rule.

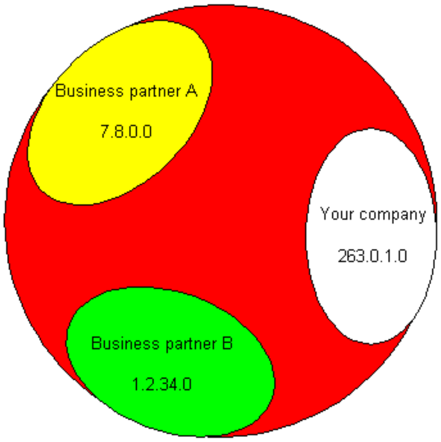

In the example, when a packet is detected, Cross-Site for Security first checks to see if it originated from business partner A's network. If not, it then verifies whether it came from business partner B's network. Finally, Cross-Site for Security confirms whether or not the packet came from within your own network.

If the packet came from any of the IP addresses listed, no alerts are generated. If, after checking the packet against each of the addresses specified, Cross-Site for Security determines that the source address does not match one of the IP addresses listed in the ALLOW rules, the alert in the last ALLOW rule is generated.

Monitoring Port Services

Cross-Site for Security intrusion detection includes several “default port service” signature configurations. By the term “default port service” we mean, for example, signature configurations for the FTP protocol on the FTP port.

Some of these configurations are network-specific ACL signatures, which will need to be configured by a Cross-Site Administrator. The remaining signatures can be used as part of the default configuration of Cross-Site for Security.

The following topics are discussed in this section on monitoring port services:

- Default Port Services
- Generic Port Services

Default Port Services

Default signature configurations are provided for many typical services. For descriptions and examples of how to apply the feature settings for each of the services listed below, see “Generic Port Services” on page 110. The topics that follow are in the same order as they appear in the **ids.rules** file.

See the related “Signature Configuration” section in the **ids.rules** file for a list of the default signatures provided for each of the following:

- Telnet/R Services—Several of the more sensitive commands such as "rpcinfo -p" and "finger @" are included.
- FTP and FTP Data—The more sensitive authentication checks exist; however you will want to add some signature checks for "/.cshrc", "/.profile", and "/.login".
- Finger—Finger of root and more detected.
- TFTP—Attempts to grab sensitive files over the network through this unauthenticated service are detected.

- WWW—Detects weaknesses of "bin/finger" and "bin/phf" among others.
- SENDMAIL—Mail alias expansion and address verifications that are known weaknesses in sendmail are detected.
- X11—Detects the same signatures as Telnet and FTP.
- Gopher—Monitors for signatures with password files.
- IMAP—Detects signatures associated with buffer overflow attacks and for signatures with password files.
- POP—Detects signatures associated with buffer overflow attacks.
- IDENT—Detects signatures associated with buffer overflow attacks.
- NNTP—Attempts to include shell commands in messages are detected.
- IRC—Detects signatures associated with buffer overflow attacks.
- LPR—Detects signatures associated with buffer overflow attacks.
- TALK—Detects traffic from this interactive messaging system.
- UUCP—Detects traffic from this UNIX copy-e-mail protocol.
- Kerberos—Monitors for dictionary attacks using “krbtgt” signature.
- WRITESRV—Detects signatures associated with buffer overflow attacks.

Generic Port Services

Several port service features are generic, and can therefore be applied to any of the default port services (if the protocol merits their use). The “FTP Signature Configuration” section of the **ids.rules** file offers a summary of these generic port service features. Following is a sample of the “FTP Signature Configuration“ section of the **ids.rules** file, with references to the topic in the documentation in which each section is discussed.

```
#####################################
#  FTP Signature Configuration
#####################################
PORTS 21 (See “Port Numbers” on page 111)
PTYPE TCP AUTH SRC DST (See “Port Types” on page 111)
#PROXY HOST 128.119.40.219 (See “Proxy Hosts and IP Addresses” on page 113)

# Service ACL (See “Service ACLs” on page 112)
PROMPT USER     "USER"
PROMPT PASSWD   "PASS"
PROMPT AUTHFAIL "ogin incorrect"

# Authentication Section (See “Authentication Features” on page 113)
AUTH USER "root"  NOTIFY     OFFPEAK                    3042    2
AUTH USER "guest" NOTIFY     ANY                        3043    2
AUTH USER "lp"    NOTIFY     ANY                        3044    1
AUTH USER "sync"  NOTIFY     ANY                        3045    1
AUTH USER "demos" NOTIFY     ANY                        3046    1
AUTH USER "ftp"   NOTIFY     ANY                        3043    1
AUTH USER "anonymous"   NOTIFY     ANY                  3043    1
#
# Signature List (See “Signature Alerts” on page 114)
#
SIG DST   "[Ee][Xx][Ee][Cc]"                            3047    2
SIG SRC   "uest login"                                  3048    3
SIG DST   "CWD ~root"                                   3049    3
SIG DST   "SYST"                                        3050    3
SIG DST   "SITE"                                        3051    3
SIG DST   "passwd"                                      3052    2
SIG DST   "MKD"                                         3053    2
SIG DST   "[Pp][Aa][Ss][Vv]"                            3054    2
SIG DST   ".rhosts"                                     3055    2
SIG DST   "hosts.equiv"                                 3056    2
SIG DST   "../.."                                       3057    2
END
#####################################
#  END FTP Signature Configuration
#####################################
```

Port Numbers

To monitor a port or range of port numbers, use the following command:

```
PORTS 23 512 513 514 600 - 800
```

where:

`23 512 514 514`
Specifies individual ports to monitor.

`600 - 800` Specifies a range of port numbers to monitor.

Note: *There can only be one PORTS line per section of the* ***ids.rules*** *file. The PORTS line begins each section and applies to everything up to the END line. Also, the same port number cannot appear in more than one PORTS line.*

Port Types

Designate the type of port and how it is monitored, using the following command:

```
PTYPE {TCP|UDP|TCP UDP} [AUTH] {SRC|DST|SRCDST}
```

where:

`PTYPE` Specifies the Port Type command.

`TCP` Specifies that a TCP port is being monitored.

`UDP` Specifies that a UDP port is being monitored.

`TCP UDP` Specifies that a TCP port and a UDP port are being monitored.

`AUTH` Instructs Cross-Site to check user authentication, through the use of at least one prompt. User IDs are processed and, depending on their validity, an authentication alert is generated.

`SRC` Monitors inbound data streams.

`DST` Monitors outbound data streams.

`SRCDST` Monitors inbound and outbound data streams.

The following example instructs Cross-Site for Security agents to monitor TCP ports and enable authentication checks. See “Authentication Features” on page 113 for more information on authentication checks. Both inbound and outbound data streams are monitored.

```
PTYPE TCP AUTH SRC DST
```

Service ACLs

Service ACLs are similar to the IP-level ACLs explained in “Monitoring Network Traffic” on page 103. However, the IGNORE command does not apply to Service ACLs.

Service ACLS also allow time limitations. Service ACLs use PEAK time to create an ACL rule that will only fire at PEAK time or OFFPEAK time. Authentication rules can also include the Service ACL syntax. For more information, see “Authentication Features” on page 113.

In the following example, if a connection is received from the host with IP address 123.23.32.10 during PEAK hours, alert message number 101 is generated. Alert messages numbered from 1 to 100 are reserved by Cross-Site and, therefore, cannot be referred to in the **ids.rules** file.

PEAK time represents a range of hours during an entire day, set for particular days of the week. This time range is defined in the **ids.cfg** file. For information on how to calculate PEAK time, see “Configuring Scan and Flood Parameters” on page 63.

```
NOTIFY HOST 123.23.32.10 PEAK 101 1
```

Proxy Hosts and IP Addresses

In some network configurations, a particular machine may act as a proxy host, such as a mail server. It might seem that this machine is performing TCP Wide Scans, due to the number of TCP connections that occur from that machine to many hosts on the same port. This command suppresses TCP/UDP Wide Scan alerts that may be generated due to normal and expected traffic.

Authentication Features

Cross-Site for Security provides authentication features that enable you to parse the logon process, record a user's identity associated with a session, and alert on unsuccessful logon attempts. Security agents can then generate alerts based on use by a particular user.

Custom Prompts

You must supply the prompts used to identify authorization states and authorization failure. The following examples illustrate the definition of the user, password, and authentication failure strings that are used in authentication alerts.

```
PROMPT USER "login:"

PROMPT PASSWD "password:"

PROMPT AUTHFAIL "incorrect"
```

You can customize these authorization prompts in the "Prompt Section" for the TELNET/R Services, and in the "Service ACL" section for the FTP services of the **ids.rules** file.

Authentication Alerts

There are two authentication alerts that are always enabled. The first alert is activated by any authentication "Login Failure". The other alert is activated by any "Unknown User" authentication error.

The syntax for the AUTH USER command is as follows:

```
AUTH USER [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {HOST|NET
ip_addr} [NETMASK]
```

where:

`AUTH` Specifies an authentication alert.

`USER` Specifies the user login.

`ALLOW` Specifies to not issue an authentication alert.

NOTIFY Specifies to issue an authentication alert.

PEAK Specifies to allow or generate an alert if user authorization occurs during peak hours.

OFFPEAK Specifies to allow or generate an alert if user authorization occurs during off-peak hours.

ANY Specifies to allow or generate an alert if user authorization occurs any time.

NEVER Specifies to never allow or never generate an alert if user authorization occurs.

HOST Specifies the host to allow or notify on.

NET Specifies the network to monitor.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK The netmask of the network on which the agent is to allow or notify.

The following example instructs Cross-Site to generate alert number 204 for any authorization of the user "test1" from the network 9.2.75.0.

```
AUTH USER "test1" NOTIFY NET 9.2.75.0 255.255.255.0 204 1
```

This example instructs Cross-Site to generate alert number 203 for any authorization of user "root" during off-peak hours.

```
AUTH USER "root" NOTIFY OFFPEAK 203 1
```

Signature Alerts

Signature alerts allow you to scan inbound and outbound data streams for intrusion signatures (in text string form). If the specified string is detected, the message number that corresponds to the signature is sent as part of an alert.

The syntax for a signature alert is as follows:

```
SIG {SRC|DST|SRCDST} data msg_id priority
```

where:

SIG Specifies a signature alert.

SRC Looks for data in inbound traffic.

DST Looks for data in outbound traffic.

SRCDST Looks for data in inbound and outbound traffic.

`data` Specifies the data string to look for; the data can include the regular expression syntax **[ ]** to indicate match one of the characters contained within these brackets.

`msg_id` Specifies the ID associated with the alert message.

`priority` Specifies the priority of the alert.

The following example specifies to look for a signature of the string "**rpcinfo -p**" in both inbound and outbound traffic. Notice that the number of spaces between the "**rpcinfo**" and "**-p**" switch is incidental; white space is normalized to a single space when comparing against the signature.

This method of signature detection relies on the **rpcinfo** program being executed and the first switch being the **-p** switch.

```
SIG SRCDST "rpcinfo -p" 214
```

Detecting DNS Requests and Replies

Domain Name System (DNS) features not only support the standard service ACLs discussed in “Service ACLs” on page 112, but also the DNS protocol-specific features for DNS requests and replies. See the “DNS Signature Configuration” section in the **ids.rules** file for the default DNS signature configurations provided with Cross-Site.

The following alerts are specific to DNS requests and replies:

- Address Resolve Alerts
- Zone Transfer Request Alerts
- Reply Alerts with Signatures

Address Resolve Alerts

You can configure the Cross-Site for Security agent to generate an alert when someone from an internal network attempts to access an identified domain. The alert is generated when the server attempts any DNS resolve containing an identified pattern (usually keywords) in its URL. In the **ids.rules** file, you designate the pattern that you want to generate alerts by editing the list contained in the file. The alert rule can also have standard rule-based ACLs associated with it. This alert can be used with any protocol.

The syntax for this rule is as follows:

```
REQ RESOLVE [pattern] [ALLOW|NOTIFY] [ANY|NEVER] [HOST|NET ip_addr]
[NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`RESOLVE` Specifies an address resolve request.

`pattern` Specifies the pattern contained in the request.

`NOTIFY` Specifies to issue an alert.

`NET` Specifies the network to allow or notify on.

`HOST` Specifies the host IP address to notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to notify address resolve requests.

`msg_id` Indicates the ID associated with the alert message.

The following example Address Resolve rule instructs the Security agent to generate alert message number 123 for any address resolve request that contains the URL **bad.site.com**.

```
REQ RESOLVE "bad.site.com" NOTIFY NET 123.45.54.0 255.255.255.0 123 3
```

Zone Transfer Request Alerts

You can configure the Cross-Site for Security agent to generate an alert when someone requests a zone transfer. While this kind of alert could indicate a possible intrusion, it could also be used to verify that ISP-based DNS zone transfers that are supposed to be conducted to a secondary DNS on a periodic basis are, in fact, being done for a given customer site.

The rule that specifies this alert can further qualify the alert with a standard rule-based ACL.

The syntax for the DNS Zone Transfer Request command is as follows:

```
REQ ZONEXFER [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {NET|HOST
ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`ZONEXFER` Specifies a zone transfer request.

`ALLOW` Specifies not to issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a DNS zone transfer request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a DNS zone transfer request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a DNS zone transfer request occurs any time.

`NEVER` Specifies to never generate an alert if a DNS zone transfer request occurs.

`NET` Specifies the network to allow or notify on.

`HOST` Specifies the host IP address to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify zone transfer requests.

`msg_id` Indicates the ID associated with the alert message.

`priority` Indicates the priority of the alert.

The following example Zone Transfer rule instructs Security agents to allow DNS zone transfers *only* from the network 123.45.54.xxx and to generate alert message number 203 for any zone transfer requests outside this network.

```
REQ ZONEXFER ALLOW NET 123.45.54.0 255.255.255.0 203 1
```

Note: *You should only ever receive zone transfer requests from secondary domain name servers. Any other incident is a questionable request from an external party that relays a large quantity of information about the structure of the domain.*

Reply Alerts with Signatures

Any replies sent via DNS, with a particular signature, could also trigger the generation of an alert. The alert rule can also have standard rule-based ACLs associated with it.

The syntax for this rule is as follows:

```
REPLY SIG sig [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {NET|HOST
ip_addr} [NETMASK] msg_id priority
```

where:

`REPLY` Specifies a response.

`SIG` *`sig`* Specifies the signature or string to look for.

ALLOW Specifies to not issue an alert.

NOTIFY Specifies to issue an alert.

PEAK Specifies to allow or generate an alert if a DNS reply occurs during peak hours.

OFFPEAK Specifies to allow or generate an alert if a DNS reply occurs during off-peak hours.

ANY Specifies to allow or generate an alert if a DNS reply occurs any time.

NEVER Specifies to never generate an alert if a DNS reply occurs.

NET Specifies the network to allow or notify on.

HOST Specifies the host IP address to allow or notify on.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK Specifies the netmask of the network on which to allow or notify DNS replies.

msg_id Specifies the ID of the alert message.

priority Specifies the priority of the alert message.

The following example Reply rule instructs Security agents to generate alert message number 768 for any DNS reply with the forward slash contained within it.

```
REPLY SIG "/" NOTIFY ANY 768 2
```

Under no circumstance should a “/” character appear in the normal encoding of a DNS packet. However, many buffer overflow attacks against a DNS server would likely contain this character.

Detecting Portmapper Service Request and Reply Dumps

The Portmapper Service is responsible for connecting a Remote Procedure Call (RPC) client with the appropriate RPC service. This service is similar to the role of the inetd daemon on UNIX. See the “Portmapper Configuration” section in the **ids.rules** file for a list of the default Portmapper signature configurations provided with Cross-Site for Security.

Note: *The RPC-based services provided by the Cross-Site for Security agent are supported only on UNIX platforms.*

A Portmapper dump is typically done to find out which RPC services are available. The dump reply can be further qualified by one or more RPC program numbers that are contained in a reply. In addition, like many of the other services, standard service-level and rule-level NOTIFY/ALLOW ACLs are supported.

The syntax for the Portmapper rule is as follows:

```
REQ DUMP [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {NET|HOST ip_addr}
[NETMASK] msg_id priority

REPLY DUMP [prog_num] [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{NET|HOST ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`REPLY` Specifies a reply or response.

`DUMP` Specifies the DUMP RPC list when used with REQ, and the reply from the DUMP command when used with REPLY.

`prog_num` Specifies the RPC program number contained in a reply.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a Portmapper dump occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a Portmapper dump occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a Portmapper dump occurs any time.

`NEVER` Specifies to never generate an alert if a Portmapper dump occurs.

`NET` Specifies the network to allow or notify on.

`HOST` Specifies the host IP address to allow or notify on.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify Portmapper Dumps.

msg_id Specifies the host IP address to allow or notify on.

priority Indicates the priority of the alert.

The following example Portmapper rule instructs the Security agent to generate alert message number 877 for any Portmapper dump that occurs, containing the program number 150001 (which represents pcnfs, a vulnerable service), performed at any time.

```
REPLY DUMP 150001 NOTIFY ANY 877 2
```

Detecting Mount Service Requests and Replies

The Remote Procedure Call (RPC) mount operation is a very sensitive procedure, because an entire file system is made available to an external user. The files are made available through the non-secure Network File System (NFS) protocol. Therefore, do not consider this operation in any situation where files are transported over the Internet.

To that end, an RPC Mount Service ACL can alert the Cross-Site for Security agent (UNIX only) when any mount operations are performed. However, if NFS is used over an intranet, individual mount requests can be monitored. Individual requests can also have their own set of NOTIFY/ALLOW ACLs.

Note: *The RPC-based services provided by the Cross-Site for Security agent are supported on UNIX platforms only.*

These mount services are discussed in this section:

- Mount Request Rule
- Mount Reply Rule

See the “Mount Configuration” section of the **ids.rules** file for a list of the default RPC Mount signature configurations provided with Cross-Site for Security.

Mount Request Rule

The Mount Request rule can be further qualified to specify different types of mount service requests. The following types of requests are discussed in this section:

- “Mount Request” on page 125
- “Dump Request” on page 125

- "Export Request" on page 125
- "Unmount Request" on page 125

The syntax for the Mount Request rule is as follows:

```
REQ {MOUNT|DUMP|EXPORT|UNMOUNT} [ALLOW|NOTIFY] [PEAK|OFFPEAK]
[ANY|NEVER] {HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`MOUNT` Requests the MOUNT procedure.

`DUMP` Requests the DUMP procedure.

`EXPORT` Requests the EXPORT procedure.

`UNMOUNT` Requests the UNMOUNT procedure.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a mount request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a mount request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a mount request occurs any time.

`NEVER` Specifies to never generate an alert if a mount request occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify mount requests.

`msg_id` Specifies the ID of the alert message.

`priority` Indicates the priority of the alert.

Mount Request

The Security agent can send an alert on any mount requests. You can set the normal rule-level ACL qualification, as well as file system mount request qualifications.

The following example Mount Request rule instructs the Security agent to issue alert message number 175 any time the file system "/cdrom0" is mounted during off-peak hours.

```
REQ MOUNT "/cdrom0" NOTIFY OFFPEAK 175 2
```

Dump Request

The Security agent can also send an alert if an RPC dump request is detected. You can also set the standard rule-level ACL. This request is typically generated by the “**showmount - a**” command to obtain a list of all machines mounting a server, and on which mount points they are doing so.

The following example Dump Request rule instructs the Security agent to generate alert message number 175 on any dump requests from any network.

```
REQ DUMP NOTIFY NET ANY 175 2
```

Export Request

The Export Request feature detects any export requests along with the standard set of rule-level ACLs. This request is typically generated by a “**showmount -e**” command to obtain a list of exported file systems from a server.

The following example Export Request rule configures Cross-Site for Security agents to generate alert message number 178 on any RPC export requests.

```
REQ EXPORT NOTIFY ANY 178 2
```

Unmount Request

The Security agent can also send an alert on any unmount requests. You can set the normal rule-level ACL qualification, as well as directory unmount request qualifications.

The following example Unmount Request rule instructs the Security agent to issue alert message number 172 any time the directory "/usr/local/bin" is unmounted during off-peak hours.

```
REQ UMOUNT "/usr/local/bin" NOTIFY OFFPEAK 172 2
```

Mount Reply Rule

The Mount Reply rule can also be used to specify different types of mount reply alerts that should be generated. The following types of replies are discussed in this section:

- “Export Reply” on page 127
- “Dump Reply” on page 127

The syntax for the Mount Reply rule is as follows:

```
REPLY {EXPORT|DUMP} [DIR|HOST] [ALLOW|NOTIFY][PEAK|OFFPEAK]
[ANY|NEVER] {HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REPLY` Specifies an EXPORT or DUMP reply.

`EXPORT` Specifies the reply from the EXPORT request.

`DUMP` Specifies the reply from the DUMP request.

`DIR` Specifies the directory mounted or exported.

`HOST` Specifies the host mounting directory or host exported to.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a mount reply occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a mount reply occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a mount reply occurs any time.

`NEVER` Specifies to never generate an alert if a mount reply occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify mount replies.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates priority of the alert.

Export Reply

The Export Reply feature detects any export replies, along with the standard set of rule-level ACLs. You can further qualify this rule with the directory exported and the hosts exported to.

The UNIX command **showmount -a** requests a list of the current exports on a server. The reply to this command provides the list of currently exported directories, along with whether or not the directory is exported to a specific host. Use this rule when you want to be notified if a particular directory is being exported to a specific host.

The following example Export Reply rule configures Cross-Site for Security agents to generate alert number 101 on any RPC Mount export replies of the directory “/usr1/users1” to the host “cax.bob.com”.

```
REPLY EXPORT DIR “/usr1/users1“ HOST “cax.bob.com“ NOTIFY ANY 101 2
```

Dump Reply

The Dump Reply feature detects any dump replies, along with the standard set of rule-level ACLs. You can further qualify this rule with the directory mounted and the host that has the directory mounted.

The UNIX command **showmount -e** requests a list of hosts that currently have a specific directory mounted. Use this rule when you want to be notified that a particular host has a specific directory mounted.

The following example Dump Reply rule configures Cross-Site for Security agents to generate alert number 101 on any RPC Mount dump replies of the directory “/usr1/users1” from the host “cax.bob.com”.

```
REPLY DUMP DIR “/usr1/users1“ HOST “cax.bob.com“ NOTIFY ANY 101 2
```

Detecting NFS Service Requests

The Network File System (NFS) service features include several request operations with their own specific operation qualifiers. In addition, the NFS rule set includes NFS Remote Procedure Call (RPC) protocol-level ACL support. See the “NFS Configuration” section of the **ids.rules** file for the default NFS signature configurations provided with Cross-Site for Security.

Note: *The Cross-Site for Security agent currently monitors only NFS v2 requests.*

The following NFS service requests are discussed in this section:

- Lookup Request
- SetAttribute Request
- ReadLink Request
- Create Request
- Remove Request
- Write Request

Note: Write Request*The RPC-based services provided by the Cross-Site for Security agent are supported only on UNIX platforms.*

Lookup Request

The Lookup Request feature detects any commands on a particular file or directory that would search for and open a given file. This command also supports the use of any rule-level ACLs.

The syntax for the Lookup Request rule is as follows:

```
REQ LOOKUP FILE file [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`LOOKUP` Requests to lookup a filename.

`FILE` *`file`* Specifies the filename.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a lookup request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a lookup request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a lookup request occurs any time.

`NEVER` Specifies to never generate an alert if a lookup request occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameter.

`NETMASK` Specifies the netmask of the network on which to allow or notify lookup requests.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example Lookup Request rule configures the Cross-Site for Security agent to generate alert number 123 on any attempt to do an NFS Lookup Request from the host with IP address 1.2.3.4 of the file “../..”.

```
REQ LOOKUP "../.." NOTIFY HOST 1.2.3.4 123 2
```

SetAttribute Request

The SetAttribute Request feature detects commands that would change the rights or other attributes on various files. You can further delineate this command by file, user ID (UID), group ID (GID), and permission MODE. The SetAttribute Request feature also supports the use of any rule-level ACLs.

The syntax for the SetAttribute Request rule is as follows:

```
REQ SETATTR FILE file [UID decimal] [GID decimal] [MODE octal]
[ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {HOST|NET ip_addr}
[NETMASK] msg_id priority
```

where:

REQ Specifies a request.

SETATTR Requests to set or change the attributes of a file.

FILE *file* Specifies the filename.

UID *decimal* Specifies the UID number of the owner of the file. This is a decimal number.

GID *decimal* Specifies the group ID number of the owner of the file. This is a decimal number.

MODE *octal* Specifies the file permission mode. This is an octal number.

ALLOW Specifies to not issue an alert.

NOTIFY Specifies to issue an alert.

PEAK Specifies to allow or generate an alert if a set attribute request occurs during peak hours.

OFFPEAK Specifies to allow or generate an alert if a set attribute request occurs during off-peak hours.

ANY Specifies to allow or generate an alert if a set attribute request occurs any time.

NEVER Specifies to never generate an alert if a set attribute request occurs.

HOST Specifies the host IP address to allow or notify on.

NET Specifies the network to allow or notify on.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK Specifies the netmask of the network on which to allow or notify set attribute requests.

msg_id Indicates the ID of the alert message.

priority Indicates the priority of the alert.

The following example SetAttribute Request rule configures Cross-Site for Security agents to generate alert message number 123 on any NFS SetAttribute Request to the file "/etc/aliases" except from the host with IP address 2.3.4.5.

```
REQ SETATTR FILE "/etc/aliases" ALLOW HOST 2.3.4.5 123 2
```

ReadLink Request

The ReadLink Request feature detects any read commands on a particular file or directory. This command also supports the use of any rule-level ACLs.

The syntax for the ReadLink Request rule is as follows:

```
REQ READLINK FILE file [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

REQ Specifies a request.

READLINK Requests to read a particular file or directory.

FILE *file* Specifies the filename.

ALLOW Specifies to not issue an alert.

NOTIFY Specifies to issue an alert.

PEAK Specifies to allow or generate an alert if a readlink request occurs during peak hours.

OFFPEAK Specifies to allow or generate an alert if a readlink request occurs during off-peak hours.

ANY Specifies to allow or generate an alert if a readlink request occurs any time.

NEVER Specifies to never generate an alert if a readlink request occurs.

HOST Specifies the host IP address to allow or notify on.

NET Specifies the network to allow or notify on.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK Specifies the netmask of the network on which to allow or notify readlink requests.

msg_id Indicates the ID of the alert message.

priority Indicates the priority of the alert.

The following example Readlink Request rule configures the Security agent to generate alert message number 123 on any NFS ReadLink Request to the file "/.cshrc" except from the host with IP address 1.2.3.4.

```
REQ READLINK FILE "/.cshrc" ALLOW HOST 1.2.3.4 123
```

Create Request

The Create Request feature detects any file creation commands on a particular host. You can further delineate the Create Request command by file, user ID (UID), group ID (GID), and permission MODE. This command also supports the use of any rule-level ACLs.

The syntax for the Create Request rule is as follows:

```
REQ CREATE FILE file [UID decimal] [GID decimal] [MODE octal]
[ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {HOST|NET ip_addr}
[NETMASK] msg_id priority
```

where:

REQ Specifies a request.

CREATE Requests to create a file.

FILE *file* Specifies the filename.

UID *decimal* Specifies the UID number of the owner of the file. This is a decimal number.

GID *decimal* Specifies the group ID number of the owner of the file. This is a decimal number.

MODE *octal* Specifies the file permission mode. This is an octal number.

ALLOW Specifies to not issue an alert.

NOTIFY Specifies to issue an alert.

PEAK Specifies to allow or generate an alert if a create request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a create request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a create request occurs any time.

`NEVER` Specifies to never generate an alert if a create request occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify create requests.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates priority of the alert.

The following example Create Request rule instructs Cross-Site to generate alert message number 123 when any node outside of the network 3.4.5.6 does a file creation operation on the file name "/.cshrc".

```
RPC NFS REQ CREATE FILE "/.cshrc" ALLOW NET 3.4.5.6
255.0.0.0 123 2
```

Remove Request

The Remove Request feature detects any file deletion commands on a particular file or directory. This command also supports the use of any rule-level ACLs.

The syntax for the Remove Request rule is as follows:

```
REQ REMOVE FILE file [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`REMOVE` Requests to remove/delete a file.

`FILE` *`file`* Specifies the filename.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

PEAK Specifies to allow or generate an alert if a remove request occurs during peak hours.

OFFPEAK Specifies to allow or generate an alert if a remove request occurs during off-peak hours.

ANY Specifies to allow or generate an alert if a remove request occurs any time.

NEVER Specifies to never generate an alert if a remove request occurs.

HOST Specifies the host IP address to allow or notify on.

NET Specifies the network to allow or notify on.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK Specifies the netmask of the network on which to allow or notify remove requests.

msg_id Indicates the ID of the alert message.

priority Indicates the priority of the alert.

The following example Remove Request rule instructs Cross-Site for Security agents to generate alert number 123 if the machine with IP address 1.2.3.4 attempts to do a remove request on the file named "/var/adm/messages".

```
REQ REMOVE FILE "/var/adm/messages" NOTIFY HOST
1.2.3.4 123 2
```

Write Request

The Write Request feature detects any write commands on a particular file. You can further qualify this command using the FILE and DATA qualifiers. The Write Request feature also supports the use of any rule-level ACLs.

The syntax for the Write Request rule is as follows:

```
REQ WRITE FILE DATA data [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

REQ Specifies a request.

WRITE Requests to write to a file.

FILE Specifies the filename.

DATA *data* Identifies the data to write to the specified file.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a write request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a write request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a write request occurs any time.

`NEVER` Specifies to never generate an alert if a write request occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify write requests.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example Write Request rule issues alert message number 123 on the request of any write operations to the ".rhosts" file of the data "+".

```
REQ WRITE FILE ".rhosts" DATA "+" NOTIFY ANY 123 2
```

Detecting RStatd Calls

The RStatd features are protocol-level ACLs that cause the agent to generate an alert upon decoding a Network Information Services (NIS) RStatd call. The Remote Procedure Call (RPC) program number is 10000. See the “RPC Statd Configuration” section in the **ids.rules** file for the default RStatd signature configurations provided with Cross-Site for Security.

Note: *The RPC-based services provided by the Cross-Site for Security agent are supported only on UNIX platforms.*

The syntax for the RSTAT command is as follows:

```
[ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {HOST|NET ip_addr}
[NETMASK] msg_id priority
```

where:

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if an RStatd call occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if an RStatd call occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if an RStatd call occurs any time.

`NEVER` Specifies to never generate an alert if an RStatd call occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify RStatd calls.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example rule set instructs the Security agent to report, with alert message number 1301, any RStatd request traffic detected from outside of the network 9.2.75.0.

```
RPC RSTATD

ALLOW NET 9.2.75.0 255.255.255.0 1301 2

END
```

Detecting Requests for Specific Mapnames

You can configure the Cross-Site for Security agent to send an alert when it detects a request for a specific mapname. When a request for the specified mapname is made, the agent generates the appropriate alert to notify the Cross-Site management server that this request has occurred.

The following features are used to detect requests for specific mapnames:

- YPupdated
- YPServ

YPupdated

The YPupdated feature detects and alerts upon the request of a specific mapname. The YPupdated features include the standard protocol-level ACLs, as well as the protocol-specific Mapname Request command. See the “YPupdated Requests” section in the **ids.rules** file for the default YPupdate signature configurations provided with Cross-Site for Security.

The syntax for a YPupdated Request is as follows:

```
REQ MAPNAME mapname [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`MAPNAME` *`mapname`*
Specifies the YP mapname.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a mapname request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a mapname request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a mapname request occurs any time.

`NEVER` Specifies to never generate an alert if a mapname request occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify mapname requests.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example rule set instructs the Security agent to report, with alert message number 1301, any YPupdated request traffic detected from outside of the network 9.2.75.0.

Any YPupdated requests, regardless of where they originate from, for the YP map "auto.direct" cause the agent to generate alert message number 1302.

```
RPC YPUPDATED

ALLOW NET 9.2.75.0 255.255.255.0 1301 2

REQ MAPNAME "auto.direct" NOTIFY ANY 1302 2

END
```

The **ids.rules** file includes a rule for one of the more common YPupdated mapname requests, referred to as the “Slammer Attack”. This rule instructs Security agents to notify on the request for mapname “|”.

YPServ

Cross-Site for Security also uses the YPServ feature to detect and alert upon the request of specific mapnames. The YPServ feature includes the standard protocol-level ACLs, as well as the protocol-specific Mapname request command. See the “YPServ Requests” section in the **ids.rules** file for the default YPServ signature configurations provided with Cross-Site for Security.

The syntax for a YPServ Request is as follows:

```
REQ MAPNAME mapname [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`MAPNAME` *`mapname`*
Specifies the YP mapname.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if a mapname request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if a mapname request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if a mapname request occurs any time.

`NEVER` Specifies to never generate an alert if a mapname request occurs.

`HOST` Specifies the host IP address to allow or notify on.

`NET` Specifies the network to allow or notify on.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network on which to allow or notify mapname requests.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example rule set instructs the Security agent to report, with alert message number 1301, any YPServ request traffic detected from outside of the network 9.2.75.0. Any YPServ requests, regardless of where they originate from, for the YP map "passwd" cause the agent to generate alert message number 1302.

```
RPC YPSERV

ALLOW NET 9.2.75.0 255.255.255.0 1301 2

REQ MAPNAME "passwd" NOTIFY ANY 1302 2

END
```

Detecting Requests for Specific File Names

The Status feature causes Cross-Site for Security agents to detect and alert upon the request for specific file names. Status features include the standard protocol-level ACLs, as well as the protocol-specific FILE request command. See the “Status Requests” section in the **ids.rules** file for the default RPC Status signature configurations provided with Cross-Site for Security.

The syntax for a Status Request is as follows:

```
REQ FILE file [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {HOST|NET
ip_addr} [NETMASK] msg_id priority
```

where:

`REQ`	Specifies a request.
`FILE` *`file`*	Specifies the file name.
`ALLOW`	Specifies to not issue an alert.
`NOTIFY`	Specifies to issue an alert.
`PEAK`	Specifies to allow or generate an alert if a file name request occurs during peak hours.
`OFFPEAK`	Specifies to allow or generate an alert if a file name request occurs during off-peak hours.
`ANY`	Specifies to allow or generate an alert if a file name request occurs any time.
`NEVER`	Specifies to never generate an alert if a file name request occurs.
`HOST`	Specifies the host IP address to allow or notify on.
`NET`	Specifies the network to allow or notify on.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK Specifies the netmask of the network on which to allow or notify file name requests.

msg_id Indicates the ID of the alert message.

priority Indicates the priority of the alert.

The following example rule set causes the agent to report, with alert message number 1301, any Status request traffic detected from outside of the network 9.2.75.0.

Any Status requests, regardless of where they originate from, for the path "/bin/" cause the agent to generate alert message number 1302.

```
RPC STATUS

ALLOW NET 9.2.75.0 255.255.255.0 1301 2

REQ PATH "/bin/" NOTIFY ANY 1302 2

END
```

Detecting SMB-based Attacks on PC File Servers

SAMBA is a PC file server networking protocol. This is sometimes referred to as the Server Message Block (SMB) protocol. Cross-Site for Security uses several SAMBA-related features to detect SMB-based attacks on PC file servers. See the “SAMBA/Netbios Configuration” section in the **ids.rules** file for the default SAMBA signature configurations provided with Cross-Site for Security.

The following types of alerts are discussed in this section:

- Authentication Alerts
- Service Request Alerts

Authentication Alerts

The authentication alert configuration line causes the Cross-Site for Security agent to generate alerts on specific user names or passwords. This alert rule can be further qualified with standard rule-based ACLs.

The syntax for the AUTH USER command is as follows:

```
AUTH {USER username|PASSWD password} [ALLOW|NOTIFY] [PEAK|OFFPEAK]
[ANY|NEVER] {HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`AUTH` Specifies an authentication alert.

`USER` *`username`*
Specifies the user login to look for.

`PASSWD` *`password`*
Specifies the password to look for.

`ALLOW` Specifies to not issue an authentication alert.

`NOTIFY` Specifies to issue an authentication alert.

`PEAK` Specifies to allow or generate an alert if user authorization occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if user authorization occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if user authorization occurs any time.

`NEVER` Specifies to never generate an alert if user authorization occurs.

`HOST` Specifies the host to monitor.

`NET` Specifies the network to monitor.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network to monitor.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example Authentication rule causes the agent to generate alert number 204 for any log on as Administrator without any password specified.

```
AUTH USER "ADMINISTRATOR" PASSWD "\0" NOTIFY ANY 204 2
```

This example Authentication rule causes the agent to generate alert message number 205 upon any attempted log on with a null user name and password. This logon attempt is associated with the Red Button attack.

```
AUTH USER "\0" PASSWD "\0" NOTIFY ANY 205 2
```

Service Request Alerts

The Cross-Site for Security agent can also send an alert when it detects an SMB Service request. Hackers make an SMB Service request to a particular service name. The service name is the parameter given to the SERVICE SMB command. See the “SAMBA Configuration” section in the **ids.rules** file for the default SMB signature configurations provided with Cross-Site for Security.

The syntax for the SMB Service Request command is as follows:

```
REQ SERVICE service [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REQ` Specifies a request.

`SERVICE` *`service`*
Specifies the service requested.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if an SMB Service request occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if an SMB Service request occurs during off-peak hours.

`ANY` Specifies to allow or generate an alert if an SMB Service request occurs any time.

`NEVER` Specifies to never generate an alert if an SMB Service request occurs.

`HOST` Specifies the host to monitor.

`NET` Specifies the network to monitor.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network to monitor.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example Service Request rule causes the agent to generate alert message number 204 for any SMB Service request to the service named "WINNT$". This is the default administrative service found on NT Workstations.

```
SMB REQ SERVICE "WINNT$" NOTIFY ANY 204
```

Detecting ICMP Attacks

Cross-Site for Security uses Internet Control Message Protocol (ICMP) features to detect PING-based attacks and routing indirection. See the “ICMP Configuration” section in the **ids.rules** file for the default ICMP signature configurations provided with Cross-Site for Security.

The following features used to detect PING-based attacks and routing indirection are discussed in this section:

- Echo
- Redirect

Echo

The Echo feature is used to detect oversized ICMP ping packets. These types of attacks are frequently called "Ping of Death" or “POD“ attacks. Use a threshold value (in bytes) to indicate, based on the length of the ping packet, whether or not you want the agent to generate an alert.

The syntax for an ICMP Echo signature is as follows:

```
ECHO length [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER] {HOST|NET
ip_addr} [NETMASK] msg_id priority
```

where:

`ECHO` *`length`*
Specifies the maximum length for packets (in bytes).

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if an ICMP packet is received during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if an ICMP packet is received during off-peak hours.

`ANY` Specifies to allow or generate an alert if an ICMP packet is received any time.

`NEVER` Specifies to never generate an alert if an ICMP packet is received.

`HOST` Specifies the host to monitor.

`NET` Specifies the network to monitor.

`ip_addr` Indicates an IP address if used with the HOST or NET parameters.

`NETMASK` Specifies the netmask of the network to monitor.

`msg_id` Indicates the ID of the alert message.

`priority` Indicates the priority of the alert.

The following example Echo rule causes the agent to generate alert message number 1800 for any ICMP ping packets that are greater than 1024 bytes in length.

```
ECHO 1024 NOTIFY ANY 1800 1
```

Redirect

The Redirect feature is used to detect ICMP routing indirection.

The syntax for an ICMP Redirect signature is as follows:

```
REDIRECT [ALLOW|NOTIFY] [PEAK|OFFPEAK] [ANY|NEVER]
{HOST|NET ip_addr} [NETMASK] msg_id priority
```

where:

`REDIRECT` Specifies a redirection.

`ALLOW` Specifies to not issue an alert.

`NOTIFY` Specifies to issue an alert.

`PEAK` Specifies to allow or generate an alert if an ICMP redirection occurs during peak hours.

`OFFPEAK` Specifies to allow or generate an alert if an ICMP redirection occurs during off-peak hours.

ANY Specifies to allow or generate an alert if an ICMP redirection occurs any time.

NEVER Specifies to never generate an alert if an ICMP redirection occurs.

HOST Specifies the host to monitor.

NET Specifies the network to monitor.

ip_addr Indicates an IP address if used with the HOST or NET parameters.

NETMASK Specifies the netmask of the network to monitor.

msg_id Indicates the ID of the alert message.

priority Indicates the priority of the alert.

The following example Redirect rule causes the agent to generate alert message number 1800 for any ICMP reroute responses that were received due to route table changes or an incorrectly configured system.

```
REDIRECT NOTIFY ANY 1800 1
```

5

Cross-Site Administration

Cross-Site administration includes maintaining Cross-Site resources, Cross-Site data, and the Cross-Site server itself. This section provides details about the following administrative tasks:

Access control

This section provides information about controlling access to Cross-Site resources. Cross-Site roles and principals are discussed here, as well as how to control user access to services.

Event handlers

The events generated by Cross-Site are stored in the management repository. You can configure events to be forwarded to various destinations. This section describes these destinations, which include SNMP, SMTP, and TEC. Event handlers are used to forward events to these destinations. This section also includes the procedures for configuring event handlers and forwarding events from the console.

Data management

Certain data in the management repository collects over time, thereby requiring that you monitor it. This data includes events, management data that is uploaded by Cross-Site agents, and data generated by tasks. However, Cross-Site does not monitor the management repository for space utilization and thresholds. This part of the documentation provides information about the **xs_purgedb** command, which enables you to delete unwanted or outdated data from the management repository.

Restarting and uninstalling Cross-Site

Finally, this part includes instructions for restarting and uninstalling Cross-Site. You may need to restart the Cross-Site server if, for example, you install another Cross-Site application or configure an event handler. The procedures for uninstalling Cross-Site are provided in case you no longer need an agent or if you wish to reinstall Cross-Site.

Managing Access Control

Access control is a way to restrict access to resources and functions. To manage access control, you manage three things:

- Roles that define privileges
- Principals who are assigned roles
- Channels or services that use roles for access control

By configuring roles, principals, and ACLs for a service or channel, you specify who can access Cross-Site resources.

A principal is a resource to which roles are assigned in order to grant privileges. Privileges enable Cross-Site to ensure that a user, agent, or Cross-Site management server has unique access to Cross-Site functions. Because you can assign different roles to different principals, the principals act as identifiers: they tell the management server who or what is requesting access to a Cross-Site function.

Roles grant privileges to principals. They also control access to services and, if installed, Deployment channels. You assign roles to principals and channels using each resource’s configuration panel. You assign roles to services using the **Permissions** explorer in the Admin view. When a principal’s role intersects with a service’s or channel’s role, access is granted.

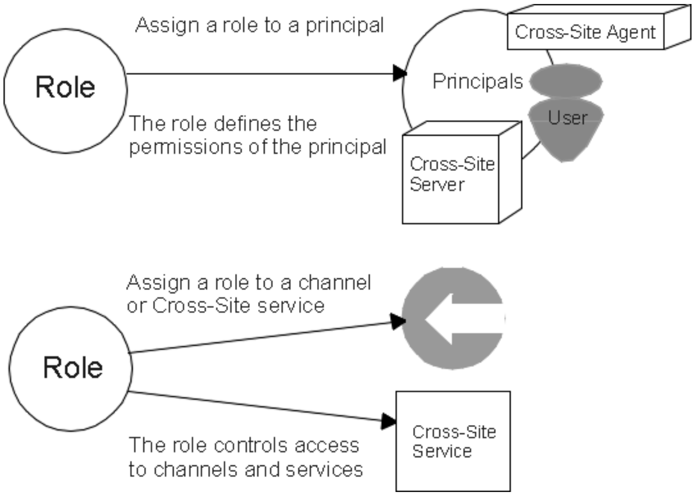
Role
Assign a role to a principal
The role defines the permissions of the principal
Principals
Cross-Site Agent
User
Cross-Site Server
Role
Assign a role to a channel or Cross-Site service
The role controls access to channels and services
Cross-Site Service

An ACL entry represents a specific Cross-Site function or channel. ACL entries are grouped in Access Control Lists (ACLs) by function. You can view ACLs and ACL entries on the **Permissions** explorer.

This section defines principals, roles, and ACL entries. Then, it takes you through the procedures for creating and configuring principals and roles.

Principals

Principals are assigned security roles and request access to secure Cross-Site administrative functions or Deployment channels. The following principals are provided by default:

Users A Cross-Site user is any user who requests access to Cross-Site functions. A user can be an administrator that manages Cross-Site or an end user of a Cross-Site agent.

Agents Agents contact the management server to send events and data, and to download information. The management server assigns the appropriate default roles to agents automatically during agent installation.

Management Servers

A management server has preconfigured security roles that will become important in future releases when Cross-Site management servers can communicate with each other.

The following sections explain how to assign roles to principals. However, most often you will assign roles to users, as described in “Creating and Configuring Users” on page 159.

Roles

Roles enable you to give principals access to Cross-Site functions and Deployment channels. You assign roles to principals and you assign roles to Cross-Site functions and Deployment channels. To assign roles to principals and channels, use the resource’s configuration panel. You can assign roles to services through the **Permissions** explorer by assigning roles to the ACL entries associated with that Cross-Site function. However, Cross-Site provides predefined roles that are assigned to ACL entries. The predefined roles enable you to assign varying levels of access, without having to manipulate ACL entries.

Cross-Site agents are automatically assigned the **agent** role and the Cross-Site server is automatically assigned the **mgmtserver** role. However to differentiate users from one another, you must assign roles during configuration.

The following roles are configured by default on your management server.

any Only used in ACL entries. It provides read-only access to non-critical information for any authenticated principal.

agent Assigned only to Cross-Site agents for Cross-Site operations.

user For end users who need to run agents and need read-only access to information and executable files on the Cross-Site server. This role should be assigned to end users who do not administer the product.

admin For administrators who need full access to all services and channels. Assign this role to a person who is responsible for making organizational and content changes to the management server. This role is equivalent to root on UNIX, or Administrator on NT.

mgmtserver

Automatically assigned to Cross-Site servers only; this role is for inter-domain operations. This role is not for users.

install For users who are responsible for installing Cross-Site for Security agents.

securityAgent
Automatically assigned to all Cross-Site for Security agents. This role is not for users.

availability For end users of Cross-Site agents who must install and use the agent component of the Cross-Site for Availability application.

deployment For end users of Cross-Site agents who must install and use the agent component of the Cross-Site for Deployment application.

For more information about how to assign a role to a user, see “Creating and Configuring Users” on page 159.

Access Control Lists and Entries

An ACL defines what roles are required to access Cross-Site services and Deployment channels. Each ACL entry represents a method or function you can perform or a Deployment channel in Cross-Site. You can set permissions for a role by assigning the role to ACL entries. You can then assign the role to a user to grant that user permission to access Cross-Site functions that require that role.

Each service ACL entry specifies the following:

- Name of the function; for example, **createUser** is the name of the function that creates a user in the Cross-Site management repository
- The roles that are assigned to the ACL entry
- Whether an unauthenticated user can invoke the function

The management server automatically creates a channel ACL when you create a channel resource. Channel ACLs have the same name as the channel they represent and have only one entry. You can configure channel ACLs from the **Permissions** explorer. However, Tivoli recommends that you assign roles to channels through the channel configuration on the **Resources** explorer.

Authentication and Authorization

The Cross-Site server authenticates a principal the first time it makes a request and when its credentials expire. Credentials consist of the principal’s name, roles, and expiration time.

The management server authenticates principals by verifying the principal’s name and password. It also checks whether the principal’s roles intersect with those of the requested channel or service. The following graphic depicts Cross-Site authorization.

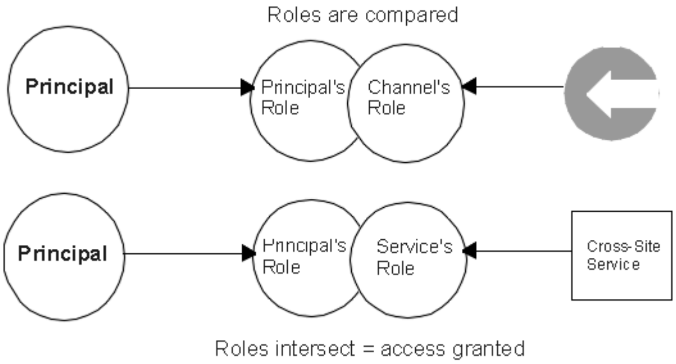

If the user name and password are valid and the principal's roles intersect with the roles of the requested service or channel, the server authorizes the principal and sends credentials to the client machine. Credentials are stored, by default, for four hours and used in any subsequent requests, enabling the server to authorize the principal.

Creating and Configuring Users

Any user who needs to view, configure, or maintain information in Cross-Site must be represented by a user resource on the management server. Users can include Cross-Site administrators and end users of the Cross-Site agent.

Complete the following steps from the Cross-Site console to create a new Cross-Site user. (You can also follow this procedure to edit users.)

1. Select the Global view by selecting the **Global** icon at the top of the console.
2. Select the **Resources** tab in the Global view.
3. Expand either the domain or collection where you intend the user to reside.
4. Select **User** from the drop-down menu below the **Resources** tab.
5. Click the **Create New** button, which is directly to the right of the drop-down list. The **Create new user** dialog is displayed.
6. Enter a name and password for user in the appropriate fields. Confirm the password and click **Create**. The new user is displayed in the tree, under your server's domain.

   Note: *When creating or editing a user with the* ***admin*** *role, do not specify a user's password that contains the* ***$*** *(dollar sign) in any position other than the last character of the password.*

7. Change the user’s properties on the **Properties** tab (on the right side of the console). You can enter the full name of the user in the **Name** field. You can also provide a description of the user, which is optional. A description helps other Cross-Site users understand the new user’s responsibilities.

8. Select the **Roles** tab to specify a role for the user.

9. Click the **Enabled** check box for each role you want to assign to the user. Following is a list of available roles.

   Take note that some of the roles introduced earlier are omitted in this list. Tivoli recommends that you do not use the **agent** or **mgmtserver** roles. These roles are intended for principals other than users. Also, you do not need to assign the **any** role to a user, the management server automatically assigns this default role to new users. For more information on roles and how they work, see “Roles” on page 157.

   user For end users who run agents and need read-only access to information and executable files on the Cross-Site server. This role should be assigned to end users that do not administer the product.

   admin For administrators who need full access to all services and channels. Assign this role to a person who is responsible for making organizational and content changes to the management server. This role is equivalent to root on UNIX, or Administrator on NT.

   install For users who are responsible for installing Cross-Site for Security agents.

   deployment For end users of Cross-Site agents who must install and use the agent component of the Cross-Site for Deployment application.

   availability For end users of Cross-Site agents who must install and use the agent component of the Cross-Site for Availability application.

   You can also create custom roles. For more information about creating roles, see the following section “Creating Roles”.

10. Select the **Summary** tab to view information about the user. You cannot edit the information in the **Summary** panel. The following information is displayed in this panel:

    Name The name of the user. This is the name entered when the user resource was created or edited.

    Type

    The user’s resource type

ID The resource ID of the user, which was generated when the user resource was created

Creator The user name of the person who created the user resource. (This field is not supported in this release.)

Last modified
The date and time the user’s properties were last modified. (This field is not supported in this release.)

11. Click the **Apply** button to create the user resource or click the **Reload** button to refresh the fields in the configuration panels with the latest data in the management repository.

New Cross-Site users are displayed in the tree on the console.

Creating Roles

If the default roles provided with Cross-Site do not suit your needs, you can create custom roles. Creating a role involves the following steps:

- Creating a new role
- Assigning the role to ACL entries

Creating a New Role

To create a new role, complete the following steps from the Cross-Site console:

1. From the Admin view, select the **Roles** tab.
2. Click the **Create a New Role** button. The new role fields are displayed on the right.
3. Enter a name for the role in the **Name** field.
4. Type a description of the role in the **Description** field. This is optional. Ensure that descriptions are meaningful to other administrators.
5. Click the **Apply** button to create the role in the management repository.

   Note: *The **Apply** button is not available for predefined roles.*
6. Click **Reload** to refresh the fields in the **Roles** tab with the latest data in the management repository.

After you click the **Apply** button, the new role is displayed in the tree.

Assigning Roles to ACL Entries

To assign a role to an ACL entry, complete the following steps from the Cross-Site console:

1. In the Admin view, select the **Permissions** tab. The tree on the left lists the Cross-Site ACLs.

2. Expand or double-click on an ACL to display its entries.

3. Select an ACL entry. A table is displayed on the right that lists roles, their domains, and a check box indicating whether the role is able to perform the function associated with the selected ACL entry.

   Note: *If you cannot see this table, click the* ***Authenticate*** *check box. If the* ***Authenticate*** *check box is left unchecked, anyone can perform the Cross-Site function associated with the selected ACL entry.*

4. Click the **Allowed** check box beside a role to assign the role to the ACL entry.

   Note: *Use extreme caution when assigning roles to ACL entries. If you accidently select a default role and remove the allowed check mark beside a role, you could corrupt some of the roles and, therefore, the default access control settings. For a complete list of Cross-Site ACL entry defaults, see "Cross-Site ACLs and Roles" on page 225.*

Repeat these steps to assign roles to all the necessary ACL entries.

After you create a role and associate ACL entries with it, the role has a meaningful security identity. Now you can assign the role to principals. For example, you can assign the role to a user to grant that user permission to access Cross-Site functions requiring that role. For information about how to assign roles to users, see "Creating and Configuring Users" on page 159.

Managing the Event Service

The Cross-Site management server and applications generate events. An event is any warning, error message, or status message generated by the Cross-Site server and applications. Events are typically generated when the policy applied to a resource is violated, which usually occurs during the execution of a related task.

The events generated by Cross-Site are stored in the management repository. You can configure events to be forwarded to a destination console by one of the following handlers on the Cross-Site server:

- To one of more e-mail addresses, through the Simple Mail Transfer Protocol (SMTP) handler
- To a Simple Network Management Protocol (SNMP) management application, through the SNMP handler
- To a Tivoli Enterprise Console (TEC), through the TEC handler

By forwarding events, you can ensure that the appropriate people are notified of each type of event.

Cross-Site stores time in Universal Time Coordinate (UTC), otherwise known as Greenwich Mean Time (GMT). The event data stored in the management server has a UTC timestamp. This ensures that all events are displayed on the Cross-Site console in your local time.

Forwarding Events to an E-mail Account

You can configure the event service to forward events to one or more e-mail addresses, using the Simple Mail Transfer Protocol (SMTP) handler. Use the SMTP forwarding feature to ensure that the appropriate people are notified for each type of event generated. Automatic e-mail notification is a simple and efficient way to communicate that an issue needs to be addressed.

Configuring the Cross-Site Server

Before the Cross-Site server can forward events to an SMTP server, you must configure the Cross-Site server as follows:

1. Using a text editor, edit the **event.properties** file, which is located in the *base_dir*/**XSITsagt/lib/properties** directory. The *base_dir* variable represents the installation directory of the Cross-Site server. The contents of the file are as follows:

```
# # # # # # # # # # # # # # # # # # # # # # # # # # # # # # #
# Configuration file for Tivoli Cross-Site SMTP & SOCKS
# Wed Mar 10 15:49:23 CST 1999
# # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # #

trace.service=com.tivoli.xtela.core.util.DefaultTraceService
trace.switches=1-4.0

# SMTP host name
mail.smtp.host=

# socksProxyHost
socksProxyHost=

# socksProxyPort
socksProxyPort=
```

2. Specify the fully qualified host name of the SMTP server with the **mail.smtp.host** key. For example, if the SMTP server is **rabbit.mycompany.com**, the key looks like this:

```
mail.smtp.host=rabbit.mycompany.com
```

3. If your network uses a SOCKS server to access the SMTP server, specify the fully qualified name of the SOCKS server with the **socksProxyHost** key. For example, if the SOCKS server is **socks23.mycompany.com**, the key looks like this:

```
socksProxyHost=socks23.mycompany.com
```

4. If your network uses a SOCKS server to access the SMTP server, specify the port number of the SOCKS server with the **socksProxyPort** key. For example, if the SOCKS server's port is **144**, the key looks like this:

   ```
   socksProxyPort=144
   ```

5. Save and close the **event.properties** file.

6. Restart the Cross-Site server as described in "Restarting the Cross-Site Server" on page 175.

Configuring Events to Use the SMTP Handler

To forward all events of a particular type to an e-mail address, complete the following steps from the Cross-Site console:

1. Select an application view and select the **Events** explorer.

2. Select an event type from the tree. The configuration panels for the event type are displayed on the right.

3. Select the **Config** tab to view the options available for forwarding events.

4. In the **Email addresses** field, enter the e-mail address of the first person to which you want to forward instances of this type of event.

   Note: *The **Add** button becomes enabled when you type the @ ("at" symbol) in an e-mail address.*

5. Click the **Add** button. The address is added to the list.

6. Continue to enter and add e-mail addresses to the list until you have included all addresses to which you want to forward this event type.

   To remove an e-mail address from the list, select it and click the **Delete** button.

7. Click the **Apply** button to save the configuration settings.

All new events of the specified type are sent to the e-mail addresses specified.

Forwarding Events to an SNMP Management Application

You can configure events to be forwarded to a Simple Network Management Protocol (SNMP) management application such as Tivoli NetView. To forward Cross-Site events to an SNMP management application, you must have an SNMP management application installed. If your company uses Tivoli NetView, Tivoli recommends that you contact the Tivoli NetView administrator before forwarding events to the SNMP management application.

Configuring the Cross-Site and SNMP Servers

Before the Cross-Site server can forward events to an SNMP server, you must complete the following steps:

1. Using a text editor, edit the **snmp.properties** file, which is located in the *base_dir*/**XSITsagt/lib/properties** directory. The *base_dir* variable represents the installation directory of the Cross-Site server. The contents of the file are as follows:

   ```
   # # # # # # # # # # # # # # # # # # # # # # # # # # # # # #
   # Configuration file for Tivoli Cross-Site / SNMP integration
   # Wed Mar 10 15:49:23 CST 1999
   # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # #

   # Location (IP address or hostname) of SNMP platform
   # e.g. thaison.dev.tivoli.com
   ServerLocation=

   # SNMP management platform incoming SNMP trap port number
   # e.g. 162
   ServerPort=
   ```

2. Specify the fully qualified host name of the SNMP server with the **ServerLocation** key. For example, if the SNMP server is **fodder.mycompany.com**, the key looks like this:

   ```
   ServerLocation=fodder.mycompany.com
   ```

3. Specify the port number of the SNMP server with the **ServerPort** key. For example, if the SNMP server’s port is 162, the key looks like this:

   ```
   ServerPort=162
   ```

4. Save and close the **snmp.properties** file.

5. Restart the Cross-Site server as described in “Restarting the Cross-Site Server” on page 175.

6. Copy the **Cross-Site.mib** file, which is located in the *base_dir***/XSITsagt/lib** directory, to the SNMP server. This file enables the SNMP server to understand the Cross-Site events.

7. Load the **Cross-Site.mib** file on the SNMP server. Refer to the SNMP server's documentation for instructions.

Configuring Events to Use the SNMP Handler

To configure all events of a particular type to be forwarded to an SNMP management application, complete the following steps from the Cross-Site console:

1. Select an application view, and then select the **Events** tab.

2. Select an event type from the tree. The configuration panels for the event type are displayed on the right.

3. Select the **Config** tab to view the options available for forwarding events.

4. Click the **SNMP** check box to enable SNMP event forwarding.

5. Click the **Apply** button to save the configuration settings.

Repeat this procedure to configure each event type you want to forward to the SNMP management application. After you enable SNMP event forwarding, all new events of each selected type are automatically sent to the SNMP management application that is configured to receive Cross-Site events.

Forwarding Events to the Tivoli Enterprise Console

You can configure events to be forwarded to the Tivoli Enterprise Console (TEC). To forward Cross-Site events to the TEC, you must have the Tivoli Framework and TEC installed. (See the *Cross-Site Integration with Tivoli Enterprise Applications User's Guide* for version information.) Tivoli also recommends that you contact the Tivoli Enterprise administrator before configuring the TEC server.

Configuring the Cross-Site and TEC Servers

Before the Cross-Site server can forward events to a TEC server, you must complete the following steps:

1. Using a text editor, edit the **tec.properties** file, which is located in the *base_dir***/XSITsagt/lib/properties** directory. The *base_dir* variable represents the installation directory of the Cross-Site server. The contents of the file are as follows:

```
# # # # # # # # # # # # # # # # # # # # # # # # # # # # # #
# Configuration file for Tivoli Cross-Site / TEC integration
# Wed Mar 10 15:49:23 CST 1999
# # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # #

# Location (IP address or hostname) of TEC server
# e.g. foo.development.mycompany.com
ServerLocation=

# TEC server port number e.g. 5529
ServerPort=

# Event group for display on the TEC console e.g. Cross-Site
Application=

# Event subgroup for display on the TEC console e.g. Cross-Site
Component=

# Debug test mode active (yes/no)
TestMode=
```

2. Specify the fully qualified host name of the TEC server with the **ServerLocation** key. For example, if the TEC server is **robot.mycompany.com**, the key looks like this:

   ```
   ServerLocation=robot.mycompany.com
   ```

3. Specify the port number of the TEC server with the **ServerPort** key. For example, if the TEC server's port is **5529**, the key looks like this:

   ```
   ServerPort=5529
   ```

   If the TEC server is installed on an NT machine, simply consult the **$BINDIR/TME/TEC/.tec_config** file to find out the port on which the TEC server is listening.

   If the TEC server on UNIX can use portmapper to decide which port to listen on for events. However, Cross-Site cannot talk to portmapper. Therefore, if the TEC server is installed on a UNIX machine, you must modify the **$BINDIR/TME/TEC/.tec_config** file to force the TEC to a port. This should not affect deployed TEC adapters that communicate with the TEC server.

4. Specify an event group that will display the Cross-Site events on the TEC. Specify the event group with the **Application** key. For example, if you wish to view Cross-Site events in the **Cross-Site** group, the key looks like this:

   ```
   Application=Cross-Site
   ```

5. Specify an event subgroup on the TEC with the **Component** key. For example, if you wish to view Cross-Site events in the **Cross-Site** subgroup, the key looks like this:

   ```
   Component=Cross-Site
   ```

6. Activate tracing for TEC events generated by Cross-Site events with the **TestMode** key. Tivoli recommends you set this key to **no**, as follows.

   ```
   TestMode=no
   ```

7. Save and exit the **tec.properties** file.

8. Restart the Cross-Site server as described in “Restarting the Cross-Site Server” on page 175.

9. Copy the **Cross-Site.baroc** file, which is located in the *base_dir*/**XSITsagt/lib** directory, to the TEC server. This file includes the classes for each Cross-Site event.

10. Compile and load the **Cross-Site.baroc** file on the TEC server. Refer to the *Tivoli Enterprise Console Rule Builder’s Guide* for instructions.

11. Create an event group on the TEC where you can view Cross-Site events. Name the event group the same name specified in step 4 on page 168. See the *Tivoli Enterprise Console User’s Guide* for more information.

Configuring Events to Use the TEC Handler

To configure all events of a particular type to be forwarded to a TEC, complete the following steps from the Cross-Site console:

1. Select an application view, and then select the **Events** explorer.

2. Select an event type from the tree. The configuration panels for the event type are displayed on the right.

3. Select the **Config** tab to view the options available for forwarding events.

4. Click the **TEC** check box to enable TEC event forwarding.

5. Click the **Apply** button to save the settings.

Repeat this procedure to configure each event type you want to forward to the TEC. After you enable TEC event forwarding, all new events of the type selected are automatically forward to the specified TEC server.

How the Event Service Works

The Cross-Site event service uses a “publish and subscribe” model. In this model, the three components involved in posting an event include the event generator, which is the source of the event, the event dispatcher, and the event handler.

The event generator can be the Cross-Site server itself, such as the task manager which posts task status messages. It can also be a Cross-Site agent. The event generator “publishes” an event to the event dispatcher.

The event dispatcher receives an event and sends it to its defined event handlers. Event handlers, thus, “subscribe” to an event dispatcher. This defines which event handlers are available when forwarding events.

An event handler is responsible for posting an event to the correct destination. The event handlers for Cross-Site include the following:

- Log event handler, which posts an event to the management repository. The event is displayed in the event log. By default, the event dispatcher sends all events to this event handler.
- SMTP event handler, which sends mail to an e-mail address
- SNMP handler, which forwards events to a Simple Network Management Protocol
- Tivoli Enterprise Console, which forwards events to a TEC server

All types of events are preregistered with the Cross-Site server. You define the correlation between the type of event and the event handlers. Use the **Config** panel on the **Events** explorer to define which events should be forwarded to an SMTP, SNMP, or TEC handler, in addition to the log event handler.

Managing Data in the Management Repository

Certain data in the management repository is expected to collect over time, thereby requiring that you monitor it. This data includes events, management data that is uploaded by Cross-Site agents, and data generated by tasks. However, Cross-Site does not monitor the management repository for space utilization and thresholds. It is your responsibility to monitor the health of the management repository and to regularly purge events, unneeded collections, tasks, and so on.

Cross-Site provides a utility to perform this maintenance. The **xs_purgedb** command enables you to delete unwanted or outdated data from the management repository. The **xs_purgedb** command also enables you to view statistics about the amount of data stored in the management repository at any one time.

Of course, how often you need to purge data depends on the size of the database. Consult your database administrator if you are unfamiliar with database maintenance. Failure to monitor and maintain the management repository could result in the Cross-Site management server becoming unusable or ineffective because it cannot write to the database.

Syntax and Usage of the xs_purgedb Command

You must log in to the machine on which the Cross-Site server is installed to use the **xs_purgedb** command. The **xs_purgedb** command is located in the *base_dir*/**XSITsagt/bin** directory, where *base_dir* is the base directory specified when the Cross-Site server was installed. You can set the PATH environment variable to include this directory, or you can run the command by specifying its full path.

The general syntax statement for the **xs_purgedb** command is as follows:

xs_purgedb [**–user** *user_name* **–passwd** *password*] {**–event** | **–alert**}
[**–help** | **–stats** | **–delete** [**–before "***time***"**]]

where:

–user *user_name*

Specifies a Cross-Site user name. The command uses this name to log in to the Cross-Site management server. If you do not specify this argument, the command prompts you for a user name.

Note that this logon name must be assigned the **admin** role. Because this command enables you to delete potentially important Cross-Site data, you must have this administrative authority.

–passwd *password*

Specifies a Cross-Site password that is used by the command, along with the user name specified by the **–user** argument, to log in to the Cross-Site server. If you do not specify this argument, the command prompts you for a password.

–event Specifies that subsequent arguments apply to event data in the management repository. If no additional arguments are specified after **–event**, statistical information about events is printed to standard output by default.

–alert Specifies that subsequent arguments apply to alerts that are uploaded by Security agents to the Cross-Site server and stored in the management repository. If no additional arguments are specified after **–alert**, statistical information about alerts is printed to standard output by default.

–help Prints the usage statement of this command to the standard output.

–stats Prints the total number of data entries in the management repository to standard output. This statistical information is printed to standard output by default if you do not specify the **–help**, **–stats**, or **–delete** argument.

delete [**–before "***time***"**]

Deletes all data in the management repository. The following options are available with the **–delete** argument:

If you specify the **–before** argument, data stored in the management repository on and before the specified date and time is purged. You must enclose the date and time in quotation marks and the format must be YYYY/MM/DD HH:mm:ss. The year must be a four-digit number and the time must be specified as military time (Midnight is 00:00:00, noon is 12:00:00, and so on).

Refer to the following sections for examples of this command. Each section also lists the syntax statement for the **xs_purgedb** command that is specific to the type of data you are deleting from the management repository.

Deleting Events

The Cross-Site server and applications continually generate events, which are stored in the management repository. Events are generated by console operations, by tasks, when policy is violated, and based on information from Cross-Site agents. Events provide an audit trail that you can use to monitor daily operations but that you might need to purge when it becomes outdated.

To delete events from the management repository, use this command:

xs_purgedb [**–user** *user_name* **–passwd** *password*] **–event**
[**–help** | **–stats** | **–delete** [**–before “***time***”**]]

Refer to “Syntax and Usage of the xs_purgedb Command” on page 171 for a description of the arguments.

The following example purges all events in the management repository that were received on and before October 11, 1997 at 7 pm.

```
xs_purgedb -user raghu -passwd clever!! -event
-delete -before "1997/10/11 19:00:00"
```

The following example prints statistical information about the events in the management repository to standard output, but does not delete the events from the management repository:

```
xs_purgedb -event -stats
```

Because the user name and password were not specified, the command prompts you for them, as follows:

```
User:
Password:
```

The output of the command is as follows:

```
Event Admin Utility
There are x event entries.
```

where *x* is the total number of events in the management repository.

Deleting Alert Log Data

The Cross-Site for Security agent continually monitors packet traffic to and from the system on which it is installed. Periodically, the agent's activity log is uploaded to the alert log, which is a set of tables in the management repository. (Security agents generate alerts, of which a subset is converted to events on the Cross-Site server.) Agents can generate a large amount of alert information, so you might need to purge the Security agent data that is stored in the management repository from time to time.

To delete alerts from the management repository, the arguments for the command are as follows:

xs_purgedb [**–user** *user_name* **–passwd** *password*] **–alert**
[**–help** | **–stats** | **–delete** [**–before "***time***"**]]

Refer to "Syntax and Usage of the xs_purgedb Command" on page 171 for a description of the arguments.

The following example purges alerts in the management repository that were uploaded by Security agents on or before November 1, 1998 at midnight. Because a Cross-Site user name and password are not specified, the command displays a prompt.

```
xs_purgedb -alert -delete -before "1998/11/01 00:00:00"
```

The following example prints statistical information about the alerts in the management repository to standard output:

```
xs_purgedb -user admin -passwd n8tion -alert
```

Though the **–stats** argument was not specified, it is the default behavior of the command to print the following to standard output:

```
Alert Admin Utility
There are x entries.
```

where *x* is the total number of alerts in the management repository.

Restarting the Cross-Site Server

For troubleshooting, maintenance, or emergency recovery, you might need to restart the Cross-Site server. You can restart the Cross-Site server by restarting the web server itself. (The Cross-Site server restarts automatically every time you restart the web server.)

Note: *Remember that restarting the web server not only affects the Cross-Site server but also any other applications that rely on the web server.*

Before restarting the web server, obtain the user name, password, port number, and certificate password (the Key File Password) for the web server. Cross-Site requires that you use a Netscape web server; you must perform the following procedure using the web server that the Cross-Site server uses.

You can perform this procedure through a web browser or from the command line. If you have trouble restarting the web server, the web server or servlet environment might be hung. You might have to stop the web server's process manually, using the **kill** command, though this should be a last resort. See the Netscape Enterprise Server documentation for troubleshooting information about Netscape's web server.

Using the Web Browser

Complete the following steps to restart the Cross-Site server through a web browser:

1. To access the Netscape Enterprise Server administration page, enter the following URL in a web browser:

   ```
   protocol://server_name:admin_port
   ```

   If the Netscape Enterprise Server was set up with a signing certificate, specify **https** instead of **http** for *protocol*. The *server_name* and *admin_port* variables are the values specified when the web server was installed. If the page does not load, ensure that the Netscape administrative server is running.

2. If prompted, enter your Netscape administrative user name and password for the server. The Netscape Server Administration page is displayed:

   http://*server_name*:*admin_port*/admin-serv/bin/index.

   This page lists the web servers installed on the specified host.

3. Locate the name of the web server to which the Cross-Site server is connected.

4. To ensure that the web server is running, rest the cursor over the **OFF/ON** button next to the server name. The tool tip should read **Turn off**.

5. Click the **OFF/ON** button next to the server name to shut it down. A new page is displayed indicating that the state of the server has changed.

6. Click the **OFF/ON** button again to restart the web server. The web server prompts you for the certificate password (the Key File Password) if it is running HTTPS.

If you try to run the console immediately after restarting the server and it fails, wait a few minutes and try again. The web server may need time to complete its initialization.

Using the Command Line

To restart the web server from the command line, enter the following commands:

```
install_dir/https-server/stop
install_dir/https-server/start
```

or

```
/install_dir/https-server/restart
```

where:

install_dir Specifies the directory path to where the web server is installed.

server Specifies the name of the web server.

The Cross-Site server is shut down and restarted automatically when you enter these commands. If you try to run the console immediately after restarting the server and it fails, wait a few minutes and try again. The web server may need time to complete its initialization.

Pinging the Cross-Site Server

To ensure that the Cross-Site management server is up and running, you can “ping” the server. Enter the following URL in a web browser:

```
protocol://server_name:port/servlet/CrossSiteServlet
```

This utility issues an HTTP request to the Cross-Site server. If the Cross-Site server is a secure server (running over SSL), specify **https** instead of **http** for *protocol*. The *server_name* and *port* variables are the values specified when the Cross-Site server was installed.

If the server is running, the URL displays a page containing the following message:

Cross-Site Management Server Status: ALIVE
Version: *x*

If, however, the page contains text other than this, the server is down.

6

Reference

This section contains information that relates to Cross-Site for Security, but is not integral to its use. This information is provided for your reference while using the product.

Security agent policy uses three files: **ids.msg**, **ids.cfg**, and **ids.rules**. When a Security agent starts, it contacts the Cross-Site management server and downloads these files. These files establish the various rules, or policy, used to define the behavior of the agent to which it is applied. Refer to this part for an example of each file.

This part also provides reference information about the Cross-Site roles and access control lists (ACLs). Each ACL is described, including the ACL's purpose and the roles that are assigned to the ACL by default. The preconfigured roles that are installed with Cross-Site are outlined, as well as the ACLs to which each role is assigned.

Finally, a list of reference information is provided in case you wish to further investigate attacks and signatures. A glossary is also provided to describe terms used throughout the documentation.

The ids.msg Policy Configuration File

```
# ids.msg
##########################################################
#
#  HAXOR - Messages Configuration File
#
#  Elements:
#
#  TOKEN  : Numeric ID..Corresponding to TOKEN value in ids.rules file
#  VALUE  : Severity value 1-5
#
#  Layout:
#
#  TOKEN      PRIORITY
#
#  Examples:
#
#  1          1
#  100      5
#
# The maximum token value in this file should be 1000
#
#
#########################################################
# Message "ICMP - Oversized Packet ... Possible Ping-of-death (POD) attack"
```

```
1          2
# Message "ICMP - Fast Wide Scan"
2          1
# Message "ICMP - Slow Wide Scan"
3          2
# Message "TCP - Wide Scan"
4          2
# Message "TCP - Wide Scan Slow"
5          2
# Message "TCP - Port Scan - SYN"
6          1
# Message "TCP - Port Scan Slow"
7          3
# Message "TCP - FIN SCAN"
8          4
# Message "TCP - RST SCAN"
9          2
# Message "TCP SLOW FIN SCAN"
10        1
# Message "RST SLOW SCAN"
11        1
# Message "UDP - WIDE SCAN"
12        1
# Message "UDP - WIDE SCAN SLOW"
13        5
# Message "UDP - PORTSCAN"
14        2
# Message "UDP - SLOW PORT SCAN"
15        3
# Message "FIN FLOOD"
16        2
# Message "SYN FLOOD"
17        2
# Message "RST FLOOD"
18        1
# Message "UDP FLOOD"
```

```
19        2
# Message "LOGIN FAILURE"
20        2
# Message "UNKNOWN USER"
21        3
# Message "Possible IP Frag attack"
22        3
# Message "Record Route Packet"
23        2
# Message "Source Routed Packet"
24        2
# Message "X11-Connection failed"
25        3
# Message "Win-nuke attack"
26        1
# Message "Land attack"
27        1
# Message "Smurf attack"
28        1
```

The ids.cfg Policy Configuration File

```
#####################################################
#
#HAXOR - Configuration File
#
#####################################################
#
#
# Define PEAK/OFFPEAK times: Anything not PEAK is OFFPEAK
#               Beg Sec = (hour*3600)+(min*60)+sec (9:00 am = 32400)
#               End Sec = (hour*3600)+(min*60)+sec (5:00 pm = 61200)
#               Days of Week = sun=0, mon=1, tues=2..etc
#               Beg Sec   END Sec    DAYS OF WEEK
PEAK            25200     64800      1 2 3 4 5
#
# LOWEST port to check for scans
#
lowscanport 1
#
# HIGHEST port to check for scans
#
highscanport 65530
#
# Fast Scan Values
#
# Sample Time for Fast window in seconds epochtime       10
#
# When single hosts sends data to different ports in ‘epochtime’ seconds
# issue a host portscan alert
#
etcpscanmax      25
eudpscanmax      25
#
# When single host sends data to different hosts in ‘epochtime’ seconds
# issue a wide scan alert
#
ehostmax       25
#
```

```
# Send # RSTS to port in ‘epochtime’ seconds - issue
RST flood alert
erstmax          200
#
# Single hosts send # SYNS to port in  in ‘epochtime’ seconds - issue
# SYN flood alert
esynmax           20
#
# hosts send # FINS to port in  in ‘epochtime’ seconds - issue
# SYN flood alert
efinmax             50
#
# hosts send # UDP packets to port in  in ‘epochtime’ seconds - issue
# UDP flood alert
eudpmax         200
#
# Slow Scan Values
#
# Sample Time for Slow window in seconds
wintime         1800
#
# When single hosts sends data to # different ports in ‘wintime’ seconds
# issue a host portscan alert
#
wtcpscanmax       50
wudpscanmax       50
#
# When single host sends data to #  different hosts in ‘wintime’ seconds
# issue a wide scan alert
#
whostmax        30
wsynmax         200
wrstmax         200
wfinmax         200
#
# alertlog = file to log alerts to
# Note: Required messages file
#
alertfile    ids.log
#
# msgfile = file containing token/message mappings
#
msgfile      ids.msg
#
# Host to issue alerts to  (syntax: alerthost <address> <port>)
#
#alerthost      127.0.0.1  6666
#
# Syslog flag (1) = issues SYSLOG alert messages (0) = no syslog alert messages
# Note: Requires msgfile
#
syslog          0
#
# Timeout for hosts we haven’t heard from in awhile .. in seconds
#
```

```
idlehost        600
#
# Timeout for inactive sessions.. idle telnet sessions, ftp
# session..etc seconds
#
idlesession     600
#
# idlepurge - interval between cleanup - in seconds
#
idlepurge      300
#
# Management server URL
#
manage_servr     tortuga:8000
```

The ids.rules Policy Configuration File

The signatures in the **ids.rules** file are organized into sections, by application or service (such as Telnet/R, FTP, and Gopher). Within many of these sections, there are signature lists that specify signature strings, followed by the alert message numbers associated with each signature, and the severity of the incident each signature represents. For example, the following is a signature for the Telnet/R service in the **ids.rules** file:

```
SIG SRCDST    "/phf?Qalias"      4024         1
```

where:

`SIG` Identifies the entry as a signature.

`SRCDST` Indicates that the Security agent is monitoring both inbound (SRC) and outbound (DST) traffic for this signature string.

`"phf?Qalias"` Specifies the character string being detected.

`4024` Represents the number associated with the alert message that Cross-Site generates when the signature is detected. This number and the corresponding alert message are displayed on the console's **Events** explorer when you select the **Messages** item in the explorer tree.

`1` Represents the severity of the incident that the signature represents.

The remainder of this document reflects the content of the **ids.rules** file that is installed with Cross-Site for Security.

```
##########################################################
#
# Tivoli Cross-site for Security - Rules Configuration File
#
##########################################################
##########################################################
#
#                        IP Level Filtering
#
##########################################################
#
# IGNORE  - Ignore traffic
#
# Example:
# NEW: 11/97
# IGNORE SRC NET  129.34.41.0   255.255.255.0 - Ignore traffic coming from NET
# IGNORE DST NET  129.34.41.0    255.255.255.0 - Ignore traffic going to NET
# IGNORE SRCDST NET  129.34.41.0 255.255.255.0 - Ignore traffic to/from net
# IGNORE SRC HOST  129.34.41.0     - Ignore traffic coming from host
# IGNORE DST HOST  129.34.41.0     - Ignore traffic going to host
# IGNORE SRCDST HOST  129.34.41.0  - Ignore traffic to/from host
#IGNORE  HOST  128.119.40.176
#
# NOTIFY HOST - Issue alert when traffic is detected from this host
# ALLOW HOST  - Do Not Issue alert when traffic is detected from this host
# NOTIFY NET  - Issue alert when traffic is detected from this network
# ALLOW NET   - Do Not Issue alert when traffic is detected from this host
# Note: You can NOTIFY an entire network and ALLOW specific hosts
#       You can ALLOW an entire network and NOTIFY specific hosts
#
#ALLOW NET  9.2.75.0  255.255.255.128           1000    1
#NOTIFY NET 9.2.13.0  255.255.255.0                1002    3
#NOTIFY HOST 146.84.32.109                            1003    2
#ALLOW  HOST 9.2.13.110                                 1004    1
#########################################################
#########################################################
```

```
#
#                       SERVICE/SIGNATURE Section
#
#########################################################
#########################################################
#
#
#  PORTS : Destination Ports to monitor -  beginning of session block
#
#  Example:
#       PORTS 23
#       Look for data to the Telnet Service port..
#
#  PTYPE : Port/Service type
#       TCP  - TCP port
#       UDP  - UDP Port
#       AUTH - Check Authentication
#       SRC  - Inbound traffic
#       DST  - Outbound Traffic
#
#  Example:
#       PTYPE TCP AUTH SRC DST
#       Look from TCP data TO and FROM specified port (in this case: 23)
#       and look at user authentication
#
# ALLOW/NOTIFY SERVICE ACLs:
#
# NOTIFY : Issue alert when traffic is detected from this host
#       NET  : Issue alert when traffic is detected from this network
#       HOST : Issue alert when traffic is detected from this host
#       PEAK : Issue alert when traffic is detected during peak hours
#       OFFPEAK: Issue alert when traffic is detected during offpeak hours
#       ANY:   Issue alert anytime traffic is detected
#       NEVER: Never issue alert
# ALLOW : Do not issue alert when traffic is detected from this host
#       NET  : Do not issue alert when traffic is detected from this network
#       HOST : Do not issue alert when traffic is detected from this host
#       PEAK : Do not issue alert when traffic is detected during peak hours
#      OFFPEAK: Do not issue alert when traffic detected during offpeak hours
#       ANY:   do not issue alert anytime traffic is detected
#      NEVER: Always alert
#
# Example:
#       NOTIFY HOST 123.23.32.10                   101
#       Issue Alert #101  when data to/from this host for specified
#      service is detected.
#       NOTIFY NET 123.23.32.0 255.255.255.0       102
#       Issue Alert # 102 when data to/from this host for specified
#      service is detected.
#
#       ALLOW HOST 123.23.32.10                   101
#       Issue Alert #101  when data to/from this host for specified
#      service is detected.
#       NOTIFY NET 123.23.32.0 255.255.255.0       102
#       Issue Alert # 102 when data to/from this host for specified
#      service is detected.
#
#
# PROMPT      : Prompts for Authentication
#       USER  : Prompt for login name
#       PASSWD: Prompt for passwords
#       AUTHFAIL: Login Failure messages
#
# AUTH        : Authentication Entry
#       USER  : look for user
#       ALLOW : Do not issue alert..
#               HOST - For this host
#               NET  - For this Network
#               PEAK - During PEAK times
#               OFFPEAK - during OFFPEAK times
#               ANY  - ANY time
#               NEVER - NEVER
#       NOTIFY: Issue Alert...
#               HOST - For this host
#               NET  - For this Network
#               PEAK - During PEAK times
#               OFFPEAK - during OFFPEAK times
#               ANY  - ANY time
#               NEVER - NEVER
#
# Examples:
#       AUTH USER "test2"      ALLOW NEVER                              107
#       Always report #107  when "test2" logs in
#
#       AUTH USER "root"       ALLOW NET 123.45.43.21 255.255.255.0     108
#       Allow root to login in ONLY from this network..else send alert  108
#
#       AUTH USER "root"       ALLOW HOST 123.45.43.21                  109
#       Allow root to login in ONLY from this host..else send alert 109
#
#       AUTH USER "root"       ALLOW PEAK                               110
#       Allow root to login in ONLY during PEAK hours..else send alert  110
#
#       AUTH USER "root"       ALLOW OFFPEAK                            110
#      Allow root to login in ONLY during OFFPEAK hours..else send alert 110
#
#       AUTH USER "test2"      NOTIFY ANY                               107
#       Always report #107  when "test2" logs in
#
#       AUTH USER "root"       NOTIFY NET 123.45.43.21 255.255.255.0    108
#       Issue alert when root logs in from from this network.           108
#
#       AUTH USER "root"       NOTIFY HOST 123.45.43.21                 109
#       Issue alert when root logs in from from this host.
#
#       AUTH USER "root"       NOTIFY PEAK                             110
#       Issue alert when root logs in during peak hours
#
```

```
#       AUTH USER "root"       NOTIFY OFFPEAK                          110
#       Issue alert when root logs in during OFFPEAK hours
#
#  SIG : data (with single whitespace) to search for in data stream
#      SRC   - look for data in INBOUND traffic
#      DST   - look for data in OUTBOUND traffic
#      SRCDST- look for data in both INBOUND and OUTBOUND
#
# Example:
#
#  SIG DST "cat/etc/passwd" 107
#  watch outbound data for text string - when found issue alert 107
#  SIG DST "bin/phf"  108
#  watch outbound data for text string - when found issue alert 107
#
# END : End of session block
#
#########################################################
############################################
#  TELNET/Rservices Signature Configuration
############################################
PORTS 23 512 513 514 1000-3000
PTYPE TCP AUTH SRC DST
PROXY HOST 9.2.75.0
# Service ACL
#ALLOW NET 9.2.75.0 255.255.255.0       2002    0
#NOTIFY  HOST 9.2.13.109                2005    3
#NOTIFY  NET  128.119.0.0 255.155.0.0   2006    0
#
# Prompt Section
#
PROMPT USER     "login:"
PROMPT PASSWD   "assword"
PROMPT AUTHFAIL "incorrect"
PROMPT AUTHFAIL "invalid login"
PROMPT AUTHFAIL "ission Denie"
```

```
#
# Authentication Section
#
AUTH USER "field" NOTIFY    ANY                 2010    1
AUTH USER "root"  NOTIFY    ANY                 2011    1
AUTH USER "guest" NOTIFY    ANY                 2010    1
AUTH USER "lp"    NOTIFY    ANY                 2010    1
AUTH USER "sync"  NOTIFY    ANY                 2010    1
AUTH USER "demos" NOTIFY    ANY                 2010    1
#
# AS/400 Default login ids
#
AUTH USER "qsecofr"     NOTIFY    ANY           2020    1
AUTH USER "qsysopr"     NOTIFY    ANY           2020    1
AUTH USER "qpgmr"       NOTIFY    ANY           2020    1
AUTH USER "ibm"         NOTIFY    ANY           2020    1
AUTH USER "qserv"       NOTIFY    ANY           2020    1
AUTH USER "qsrv"        NOTIFY    ANY           2020    1
AUTH USER "secofr"      NOTIFY    ANY           2020    1
#
# DEC Server Default Accounts
#
AUTH USER "ACCESS"      NOTIFY    ANY           2030    1
AUTH USER "SYSTEM"      NOTIFY    ANY           2030    1
#
# SGI Default Accounts
#
AUTH USER "4DGifts"     NOTIFY    ANY   2010    1
AUTH USER "nuucp"       NOTIFY    ANY   2010    1
AUTH USER "tour"        NOTIFY    ANY   2010    1
AUTH USER "tutor"       NOTIFY    ANY   2010    1
#
# Signature List
#
SIG SRCDST      "rpcinfo -p"                    2014    1
SIG SRCDST      "tprof -x"                      2015    1
SIG SRCDST      "expn"                          2016    1
SIG SRCDST      "vrfy"                          2017    1
SIG SRCDST      "mailfrom:|"                    2018    1
SIG SRCDST      "mailfrom:’|"                   2018    1
SIG SRCDST      "mailfrom:\"|"                  2018    1
SIG SRCDST      "finger root"                   2021    1
SIG SRCDST      "finger @"                      2022    1
SIG SRCDST      "bin/phf"                       2023    1
SIG SRCDST      "/phf?Qalias"                   2024    1
SIG SRCDST      "/finger"                       2025    1
SIG SRCDST      "/date"                         2026    1
SIG SRCDST      "/nph-test-cgi"                 2027    1
SIG SRCDST      "/test-cgi"                     2028    1
SIG SRCDST      "?%0A"                          2029    1
SIG SRCDST      "?%0a"                          2029    1
SIG SRCDST      "IFS[ ]=[ ]"                    2031    1
SIG DST         ".:$PATH"                       2045    1
SIG SRCDST      "ission [Dd]enie"               2057    1
SIG DST         "dig axfr"                      2040    2
SIG SRCDST      "chmod 2"                       2042    2
SIG SRCDST      "chmod 4"                       2043    2
SIG SRCDST      "chmod 6"                       2044    2
SIG DST         "lquerypv -h"                   2051    2
SIG DST         "workman -p"                    2053    2
SIG SRCDST      "invalid login"                 2055    3
SIG DST         "rlogin -froot"                 2050    3
END
################################################
#  END Telnet/Rservices Signature Configuration
################################################
################################################
#  FTP Signature Configuration
################################################
PORTS 21
PTYPE TCP AUTH SRC DST
#PROXY HOST 128.119.40.219
```

```
# Service ACL
PROMPT USER     "USER"
PROMPT PASSWD   "PASS"
PROMPT AUTHFAIL "ogin incorrect"
# Authentication Section
AUTH USER "lp"    NOTIFY     ANY                        3044    1
AUTH USER "sync"  NOTIFY     ANY                        3045    1
AUTH USER "demos" NOTIFY     ANY                        3046    1
AUTH USER "ftp"   NOTIFY     ANY                        3043    1
AUTH USER "root"  NOTIFY     OFFPEAK                    3042    2
AUTH USER "guest" NOTIFY     ANY                        3043    2
#
# Signature List
#
SIG DST   "[Ee][Xx][Ee][Cc]"                            3047    2
SIG DST   "passwd"                                      3052    2
SIG DST   "MKD"                                         3053    2
SIG DST   "[Pp][Aa][Ss][Vv]"                            3054    2
SIG DST   ".rhosts"                                     3055    2
SIG DST   "hosts.equiv"                                 3056    2
SIG DST   "../.."                                       3057    2
SIG SRC   "uest login"                                  3048    3
SIG DST   "CWD ~root"                                   3049    3
SIG DST   "SYST"                                        3050    3
SIG DST   "SITE"                                        3051    3
END
#####################################
#  END FTP Signature Configuration
#####################################
#####################################
#  FTP DATA Signature Configuration
#####################################
PORTS 20
PTYPE TCP SRC DST
SIG  SRCDST "bin:[*!]:"  3080   5
SIG  SRCDST "0:0:root"  3080
SIG  SRCDST "sys:[*!]:"  3080
END
########################################
#  END FTP DATA Signature Configuration
########################################
########################################
#  Finger Signature Configuration
########################################
PORTS 79
PTYPE TCP DST
# Service ACL
#ALLOW NET  128.0.0.0 255.0.0.0                  449    3
#NOTIFY  NET  129.0.0.0 255.0.0.0                449    2
#NOTIFY  NET  129.0.0.0 255.0.0.0                451    2
#NOTIFY  HOST  128.119.43.151                    452    2
#
# Signature List
#
SIG  DST  "root"                                453    2
SIG  DST  "."                                   455    2
SIG  DST  "0"                                   456    4
SIG  DST  "bin/"                                457    4
SIG  DST  "planner"                             458    5
END
######################################
#  END Finger Signature Configuration
######################################
######################################
#  TFTP Signature Configuration
######################################
PORTS 69
PTYPE TCP UDP SRC DST
# Service ACL
NOTIFY  ANY                             1900    1
#
# Signature List
```

```
#
SIG  DST  "passwd"                      1905    3
SIG  DST  "inetd.conf"                  1910    3
SIG  DST  "shadow"                      1905    3
SIG  DST        "master"                1905    3
SIG  DST        ".rhosts"               1905    3
SIG  DST        "config"                1910    3
SIG  DST        "hosts.equiv"           1910    3
SIG  DST        "../.."                 1910    3
SIG  SRCDST     "bin:[*!]:"             1905    3
SIG  SRCDST     "0:0:root"              1905    3
SIG  SRCDST     "sys:[*!]:"             1905    3
SIG  SRCDST     "userPassword"          1915    3
END
#####################################
#  END TFTP Signature Configuration
#####################################
#####################################
#  WWW Signature Configuration
#####################################
PORTS 80
PTYPE TCP DST
SIG DST   "bin/finger"                          5056    1
SIG DST   "bin/date"                            5057    1
SIG DST   ".url"                                5059    3
SIG DST   ".lnk"                                5060    3
SIG DST   ".bak"                                5061    3
SIG DST   "bin/perl"                            5062    3
SIG DST   "bin/tcsh"                            5062    3
SIG DST   "bin/tcl"                             5062    3
SIG DST   "-bin/sh%20"                          5062    3
SIG DST   "bin/csh"                             5062    3
SIG DST   "filename=\";"                        5073    3
SIG DST   "/uploader.exe"                       5074    3
SIG DST   "bin/Count.cgi?"                      5075    3
SIG DST   "../../"                              5076    3
SIG DST   "//////"                              5077    3
SIG DST   "bin/htmlscript?.."                   5078    3
SIG DST   "bin/www-sql/"                        5079    3
SIG DST   "bin/pfdisplay.cgi"                   5080    3
SIG DST   "bin/man.sh?"                         5081    3
SIG DST   ".nfs/?[Oo]pen\0"                     5082    3
SIG DST   "bin/textcounter.pl"                  5083    3
SIG DST   ".asp::$DATA"                         5084    3
SIG DST   "/etc"                                5085    3
SIG DST   "nph-test-cgi"                        5058    4
SIG DST   "bin/webdist.cgi"                     5063    5
SIG DST   "bin/wrap/"                           5065    5
SIG DST   "bin/handler"                         5066    5
SIG DST   "bin/aglimpse/80|IFS"                 5067    5
SIG DST   "bin/websendmail"                     5068    5
SIG DST   "bin/webgais"                         5069    5
SIG DST   "bin/php.cgi"                         5070    5
SIG DST   "bin/test-cgi"                        5071    5
SIG DST   "bin/campas?%0"                       5072    5
SIG DST   "bin/phf"                             5055    5
END
#####################################
#  END WWW Signature Configuration
#####################################
#####################################
#  MailServer Signature Configuration
#####################################
PORTS 25
PTYPE TCP DST SRC
SIG  DST        "[Ee][Xx][Pp][Nn][ ]"  6010    1
SIG  DST        "[Vv][Rr][Ff][Yy][ ]"  6020    1
SIG  DST        [Hh][Ee][Ll][Pp][ ]\n" 6100    2
SIG  DST        from[ ]:[ ]|"   6030    3
SIG  DST        from:\"|"  6050    3
SIG  DST        from[ ]:[ ][‘\"]|"  6040    3
SIG  SRC        8.[65432].[987615432]"  6060    3
SIG  DST        "rcpt to[ ]:[ ]bin"  6080    3
SIG  DST        rcpt to[ ]:[ ]sundiag"    6080
SIG  DST        rcpt to[ ]:[ ]adm"         6080
SIG  DST        rcpt to[ ]:[ ]sync"         6080
SIG  DST        rcpt to[ ]:[ ]operator"  6080
SIG  DST        rcpt to[ ]:[ ]uucp"      6080
SIG  DST        rcpt to[ ]:[ ]daemon"         6080
SIG  DST        rcpt to[ ]:[ ]|"        6090    3
SIG  DST        rcpt to[ ]:[ ][‘\"][ ]|" 6090
SIG  SRC        irectly to pro"   6130    3
SIG  SRCDST        bin:[*!]:"   6135    3
SIG  SRCDST        "0:0:root"   6135    3
SIG  SRCDST        sys:[*!]:"    6135    3
SIG  DST        "[Ww][Ii][Zz][ ]\n"       6110    4
SIG  DST        "[De][Ee][Bb][Uu][Gg][ ]\n" 6120  4
END
########################################
#  END SENDMAIL Signature Configuration
########################################
########################################
#  X11 Signature Configuration
########################################
PORTS 6000
PTYPE AUTH TCP SRC DST
# Service ACL
#ALLOW NET 9.2.75.0 255.255.255.0                2002
#NOTIFY  HOST 9.2.13.109                         2005
#NOTIFY  NET  128.119.0.0 255.155.0.0            2006
#
# Let’s look for the same signatures as telnet/rservices
#
SIG SRCDST      "rpcinfo -p"                    4014    1
SIG SRCDST      "tprof -x"                      4015    1
SIG SRCDST      "expn"                          4016    1
SIG SRCDST      "vrfy"                          4017    1
SIG SRCDST      "mailfrom:|"                    4018    1
SIG SRCDST      "mailfrom:’|"                   4018    1
SIG SRCDST      "mailfrom:\"|"                  4018    1
SIG SRCDST      "finger root"                   4021    1
SIG SRCDST      "finger @"                      4022    1
SIG SRCDST      "bin/phf"                       4023    1
SIG SRCDST      "/phf?Qalias"                   4024    1
SIG SRCDST      "/finger"                       4025    1
SIG SRCDST      "/date"                         4026    1
SIG SRCDST      "/nph-test-cgi"                 4027    1
SIG SRCDST      "/test-cgi"                     4028    1
SIG SRCDST      "?%0A"                          4029    1
SIG SRCDST      "?%0a"                          4029    1
SIG SRCDST      "IFS[ ]=[ ]foo"                 4031    1
SIG DST         ".:$PATH"                       4045    1
SIG SRCDST      "ission [Dd]enie"               4057    1
SIG DST         "dig axfr"                      4040    2
SIG SRCDST      "chmod 2"                       4042    2
SIG SRCDST      "chmod 4"                       4043    2
SIG SRCDST      "chmod 6"                       4044    2
SIG DST         "lquerypv -h"                   4051    2
SIG DST         "workman -p"                    4053    2
SIG SRCDST      "invalid login"                 4055    3
SIG DST         "rlogin -froot"                 4050    3
END
#########################################################
#########################################################
#                 DNS Signature Configuration
#########################################################
#########################################################
#
# DNS - Service Block begin
#
# ALLOW/NOTIFY SERVICE ACLs:
#
# NOTIFY : Issue alert when traffic is detected from this host
# NOTIFY NET  : Issue alert when traffic is detected from this network
# ALLOW  HOST : Do Not Issue alert when traffic is detected from this host
# ALLOW  NET  : Do Not Issue alert when traffic is detected from this host
#
# ALLOW/NOTIFY SERVICE ACLs:
#
# NOTIFY : Issue alert when traffic is detected from this host
#       NET  : Issue alert when traffic is detected from this network
#       HOST : Issue alert when traffic is detected from this host
#       PEAK : Issue alert when traffic is detected during peak hours
#       OFFPEAK: Issue alert when traffic is detected during offpeak hours
#       ANY:   Issue alert anytime traffic is detected
#       NEVER: Never issue alert
# ALLOW : Do not issue alert when traffic is detected from this host
#       NET  : Do not issue alert when traffic is detected from this network
#       HOST : Do not issue alert when traffic is detected from this host
#       PEAK : Do not issue alert when traffic is detected during peak hours
#      OFFPEAK: Do not issue alert when traffic detected during offpeak hours
#       ANY:   do not issue alert anytime traffic is detected
#       NEVER: Always alert
#
# REQ - DNS Request
#       RESOLVE - Address to Resolve    ( to be added )
#               ACL - Access Control
#       ZONEXFER - Zone transfer Request
#               ACL - Access Control
#
# Example:
#       REQ RESOLVE "bad.site.com" NOTIFY  NET 123.45.54.0 255.255.255.0 123
#
#       Notify when site from internal network attempts to resolve
#       "bad.site.com"
#
#       REQ ZONEXFER ALLOW NET 123.45.54.0 255.255.255.0 123
#
#       Allow DNS zone transfer ONLY from Network 123.45.54.xxx and
#       issue Alert 123 for ZONEXFER request outside network
```

```
#
# REPLY - DNS Response
#       SIG "signature" to look for
#               ACL - Access Control
#
# Example:
#       REPLY SIG "/bin/sh" NOTIFY ANY  123
#       Issue Alert  #123 for any DNS responses containing SIG "/bin/sh"
#
# END - End of block
#########################################################
DNS
#ALLOW NET 123.45.54.0 255.255.255.0            765
REQ   ZONEXFER   NOTIFY ANY                     766    2
REPLY SIG "\n"   NOTIFY ANY                     767    3
REPLY SIG "\r"   NOTIFY ANY                     768    3
REPLY SIG "in/"  NOTIFY ANY                     769    3
END
#########################################################
# END of         DNS Signature Configuration
#########################################################
#########################################################
#########################################################
#                PORTMAPPER Configuration
#########################################################
#########################################################
#
#  RPC - RPC Token - Beginning of RPC service block
#
#  Example: RPC PMAP - RPCportmapper
#               MOUNT - Mountd program
#               NFS - Network File System
#               NIS - Name Information Service
#               STATD - STATD service
#               YPUPDATE - YPUPDATE Program
#               YPSERV   - YPServer Program
#               PCNFSD   - Pcnfsd service
#               100000-200000 - other RPC services
#
#
# ALLOW/NOTIFY SERVICE ACLs:
#
# NOTIFY : Issue alert when traffic is detected from this host
#       NET  : Issue alert when traffic is detected from this network
#       HOST : Issue alert when traffic is detected from this host
#       PEAK : Issue alert when traffic is detected during peak hours
#       OFFPEAK: Issue alert when traffic is detected during offpeak hours
#       ANY:   Issue alert anytime traffic is detected
#       NEVER: Never issue alert
# ALLOW : Do not issue alert when traffic is detected from this host
#       NET  : Do not issue alert when traffic is detected from this network
#       HOST : Do not issue alert when traffic is detected from this host
#       PEAK : Do not issue alert when traffic is detected during peak hours
#      OFFPEAK: Do not issue alert when traffic detected during offpeak hours
#       ANY:   do not issue alert anytime traffic is detected
#       NEVER: Always alert
#
# REQ - RPC Request
#       DUMP - DUMP RPC list..
#               ACL - when to issue report
#
# Example:
#
#       REQ   DUMP      ALLOW   NET 128.119.0.0 255.255.0.0     175
#       Allow RPC dump request from the 128.119.0.0 network.. all
#     others will cause an alert #175 to be issued
#
# REPLY - RPC Reply
#       DUMP - Reply from DUMP command
#
# Example:
#       REPLY DUMP      150001  NOTIFY  ANY                     177
#       If the results of an DUMP request contain the  program number
#     150001 (pcnfs - a vulnerable service) report it
#########################################################
#########################################################
RPC   PMAP
#NOTIFY  NET  128.0.0.0   255.0.0.0                      871    2
#REQ   DUMP      ALLOW   NET 128.119.0.0 255.255.0.0     875    0
#REQ   DUMP              NOTIFY  ANY                     876    0
REPLY DUMP      150001  NOTIFY  ANY                      877    1
REPLY DUMP      100024  NOTIFY  ANY                      879    1
REPLY DUMP      100023  NOTIFY ANY                       879    1
REPLY DUMP      100017  NOTIFY ANY                       880    1
REPLY DUMP      100020  NOTIFY ANY                       881    1
REPLY DUMP      100015  NOTIFY ANY                       882    1
END
#########################################################
#########################################################
#               END  PORTMAPPER Configuration
#########################################################
#########################################################
#########################################################
#########################################################
#                BEGIN MOUNT Configuration
#########################################################
#########################################################
#
#  RPC - RPC Token - Beginning of RPC service block
#
#  Example: RPC
#               MOUNT - Mountd program
#
# ALLOW/NOTIFY SERVICE ACLs:
#
# NOTIFY : Issue alert when traffic is detected from this host
#       NET  : Issue alert when traffic is detected from this network
#       HOST : Issue alert when traffic is detected from this host
#       PEAK : Issue alert when traffic is detected during peak hours
#       OFFPEAK: Issue alert when traffic is detected during offpeak hours
#       ANY:   Issue alert anytime traffic is detected
#       NEVER: Never issue alert
# ALLOW : Do not issue alert when traffic is detected from this host
#       NET  : Do not issue alert when traffic is detected from this network
#       HOST : Do not issue alert when traffic is detected from this host
#       PEAK : Do not issue alert when traffic is detected during peak hours
#      OFFPEAK: Do not issue alert when traffic detected during offpeak hours
#       ANY:   do not issue alert anytime traffic is detected
#       NEVER: Always alert
#
# REQ - Requests to Mount Service
#       UMOUNT    - Request to mount filesystem
#                   ACL - when to issue report
#       MOUNT     - Request to mount filesystem
#                   ACL - when to issue report
#       DUMP      - ACL - when to issue report
#       EXPORT    - ACL - when to issue report
#       EXPORTALL - ACL - when to issue report
# Example:
#
#       REQ   DUMP      ALLOW   NET 128.119.0.0 255.255.0.0     175
#       Allow  DUMP  request from the 128.119.0.0 network.. all others
#       will cause an alert #175 to be issued
#
#       REQ   MOUNT  "/cdrom" ALLOW PEAK                        175
#       Allow RPC mount "/cdrom"  request during PEAK hours.. all
#       others will cause an alert #175 to be issued
#
#       REQ    EXPORT   ALLOW   NET 128.23.23.1 255.255.255.0   184
#       Allow export request only from internal ( 128.23.23.0) network.
#       Report attempts from other networks
#
# REPLY - RPC Reply
#       EXPORT - Reply from EXPORT request
#               DIR - Directory exported
#               HOST - Host exported to
#                       ACL - Access Control
#       DUMP   - Reply from DUMP request
#               DIR - Directory mounted
#               HOST - host mounting directory
#                       ACL - Access Control
# Examples:
#
#       REQ    MOUNT   "/CDROM" NOTIFY ANY                                      180
#       Report any attempt to mount "/CDROM"
#
#       REQ    DUMP     ALLOW   NET 128.23.23.0 255.255.255.0           182
#       Report any ‘showmount -a" dump requests from outside network
#
#       REQ    EXPORT   ALLOW   NET 128.23.23.0 255.255.255.0           184
#       Report any ‘showmount -e" exports requests from outside network
#
#       REPLY  EXPORT   NOTIFY  ANY                                     186
#       Report any world-exportable filesystem
#
#       REPLY  EXPORT   DIR "/users/users1"     NOTIFY    ANY             187
#       Report when "/users/users1" is exported to ANY host
#
#       REPLY  EXPORT   HOST "themis"   NOTIFY    ANY                     188
#       Report when any file system is exported to "HOST"
#
#       REPLY  DUMP     HOST "themis"           NOTIFY    ANY             190
#       Report when hosts has mounting any filesystem
#
#       REPLY  DUMP     DIR "/dir"           NOTIFY    ANY             190
#       Report when "/dir" filesystem is mounted
#
#########################################################
#########################################################
RPC    MOUNT
```

```
#ALLOW  NET     9.242.0.0 255.255.0.0                           100500
"External mount req"

#REQ   MOUNT  "/" ALLOW  NET        9.242.0.0 255.255.0.0  100505  "policy
violation : bad mount /"

#REQ    EXPORT  ALLOW   NET 9.242.0.0  255.255.0.0              100515 3
"showmount -a query from outside network"

#REPLY  EXPORT  DIR "/" ALLOW NET 9.242.0.0 255.255.0.0 987    100525 3 "FS
exported bad network"

REQ    DUMP     NOTIFY ANY                                      982    3

REPLY  EXPORT   NOTIFY  ANY                                     986    4

END

#########################################################

#               END MOUNT Configuration

#########################################################

#########################################################

#########################################################

#

#                BEGIN NFS Configuration

#

#########################################################

#########################################################

#

#  RPC - RPC Token - Beginning of RPC service block

#

#  Example: RPC

#               NFS - Network File System

#

#

# ALLOW/NOTIFY SERVICE ACLs:

#

# NOTIFY : Issue alert when traffic is detected from this host

#       NET  : Issue alert when traffic is detected from this network

#       HOST : Issue alert when traffic is detected from this host

#       PEAK : Issue alert when traffic is detected during peak hours

#       OFFPEAK: Issue alert when traffic is detected during offpeak hours

#       ANY:   Issue alert anytime traffic is detected

#       NEVER: Never issue alert
# ALLOW : Do not issue alert when traffic is detected from this host
#       NET  : Do not issue alert when traffic is detected from this network
#       HOST : Do not issue alert when traffic is detected from this host
#       PEAK : Do not issue alert when traffic is detected during peak hours
#      OFFPEAK: Do not issue alert when traffic detected during offpeak hours
#       ANY:   do not issue alert anytime traffic is detected
#       NEVER: Always alert
#
# REQ - Requests to NFS Service
#
#       LOOKUP "file" - Request to lookup a filename
#       SETATTR       - Request to set/change the attributes of a file
#               FILE "file" - file name
#               UID  uid    - Owner of file
#               GID  gid    - Group of file
#               MODE mode   - File permission mode
#       READLINK - Read file symbolic link
#       CREATE   - Request to Create a file
#               FILE "file" - file name
#               UID  uid    - Owner of file
#               GID  gid    - Group of file
#               MODE mode   - File permission mode
#       REMOVE "file" - Request to remove/delete file
#       RENAME "file1 "file2" - Request to rename a file
#       WRITE - Request to Write to file
#               FILE "file" -   File Name
#               DATA "data" -   File Data
#
#  Examples:
#
#       REQ     LOOKUP  "../."  NOTIFY ANY                              202
#       Report lookup requests containing "../."
#
#       REQ     SETATTR FILE "foop" UID 0 GID 0  MODE 644    NOTIFY ANY   209
#       Report request to change the file "foo" attributes
#         (uid/gid/mode) useful in detecting the creation of suid/root files
```

```
#
#       REQ     CREATE FILE ".rhosts"          NOTIFY ANY   214
#       Report the CREATE request for file .rhosts
#
#       REQ     CREATE FILE ".rhosts" UID 0 GID 0          NOTIFY ANY   214
#       Report the CREATE request for file .rhosts owned by user root/root
#
#       REQ     REMOVE FILE "hosts.deny"
#       Report the deletion request for the file "hosts.deny"
#
#       REQ     WRITE  FILE  ".rhosts"  DATA "+"        NOTIFY ANY      227
#       Report the REQUEST to write "+" in file ".rhosts"
#
#########################################################
#########################################################
RPC     NFS
#ALLOW        NET  9.242.0.0 255.255.0.0    100300    1    "NFS traffic from
outside network"
REQ        LOOKUP  "../.."    NOTIFY ANY   1200      3
REQ        LOOKUP  "passwd"  NOTIFY ANY    1201      3
REQ        LOOKUP  "shadow"  NOTIFY ANY    1202      3
REQ        LOOKUP  ".rhosts" NOTIFY ANY      1203      3
REQ        LOOKUP  "hosts.allow" NOTIFY ANY 1204      3
REQ        LOOKUP  "inetd.conf" NOTIFY ANY  1205      3

REQ     SETATTR FILE "passwd"   NOTIFY ANY 1206     3
REQ     SETATTR FILE "shadow"   NOTIFY ANY 1206     3
REQ     SETATTR FILE ".rhosts"  NOTIFY ANY 1206     3
REQ     SETATTR FILE "net.conf" NOTIFY ANY 1206     3
REQ     SETATTR FILE ".login"   NOTIFY ANY 1206     3
REQ     SETATTR FILE ".cshrc"   NOTIFY ANY 1206     3
REQ     CREATE FILE ".rhosts"   NOTIFY ANY 1207     3
REQ     CREATE UID 0 GID 0      NOTIFY ANY 1208     3

REQ     REMOVE  "passwd"        NOTIFY ANY 1212     3
REQ     REMOVE  ".rhosts"       NOTIFY ANY 1212     3
REQ     REMOVE  "hosts.allow"   NOTIFY ANY 1212     3
REQ     REMOVE  "hosts.deny"    NOTIFY ANY 1212     3
REQ     READ    "passwd"        NOTIFY ANY 1213     3
REQ     READ    ".rhosts"       NOTIFY ANY 1214     3
REQ     READ    "shadow"        NOTIFY ANY 1214     3
REQ     READ    "hosts.equiv"   NOTIFY ANY 1214     3
REQ     READ    "inetd.conf"    NOTIFY ANY 1214     3
REQ     READ    "hosts.allow"   NOTIFY ANY 1214     3
REQ     READ    "hosts.deny"    NOTIFY ANY 1214     3
REQ     READ    "hosts.lpd"     NOTIFY ANY 1214     3
REQ     WRITE  FILE  ".rhosts"  DATA "+" NOTIFY ANY 1215 3
END
#####################################
#  NFS Configuration
#####################################
#####################################
#  RPC STATD Configuration
#####################################
RPC    RSTAT
#ALLOW  NET  9.2.75.0 255.255.255.0 1400    0
END
#####################################
#   RPC STATD  Configuration
#####################################
#####################################
#  YPUPDATE REQUESTS
#####################################
#
# Alert for unusual characters in MAPNAME
RPC YPUPDATED
#NOTIFY  ANY      1551    2
REQ   MAPNAME "|"  NOTIFY  ANY          1550    2
END
#####################################
#  END YPUPDATE REQUESTS
#####################################
```

```
#########################################################
#########################################################
#                        YPSERVE REQUESTS
#########################################################
#########################################################
#
# Alert for specified MAPNAMES
#
RPC YPSERV
#ALLOW  NET  9.2.75.0 255.255.255.0      1601    2
REQ   MAPNAME "passwd"  NOTIFY ANY      1600    1
END
#####################################
#  YPSERVE REQUESTS
#####################################
#########################################################
#########################################################
#       STATUS REQUESTS - Signatures for status exploits
#########################################################
#########################################################
RPC STATUS
#ALLOW NET 123.45.43.0   255.255.255.0                           1700    2
REQ   PATH "../.."  NOTIFY ANY                                  1701    4
REQ   PATH "/bin/"  NOTIFY ANY                                  1702    3
END
#########################################################
#########################################################
#
#                SAMBA/Netbios Configuration
#
#########################################################
#########################################################
#
#  SMB - SAMBA File sharing token
#
# ALLOW/NOTIFY SERVICE ACLs:
```

```
#
# NOTIFY : Issue alert when traffic is detected from this host
#       NET  : Issue alert when traffic is detected from this network
#       HOST : Issue alert when traffic is detected from this host
#       PEAK : Issue alert when traffic is detected during peak hours
#       OFFPEAK: Issue alert when traffic is detected during offpeak hours
#       ANY:   Issue alert anytime traffic is detected
#       NEVER: Never issue alert
# ALLOW : Do not issue alert when traffic is detected from this host
#       NET  : Do not issue alert when traffic is detected from this network
#       HOST : Do not issue alert when traffic is detected from this host
#       PEAK : Do not issue alert when traffic is detected during peak hours
#      OFFPEAK: Do not issue alert when traffic detected during offpeak hours
#       ANY:   do not issue alert anytime traffic is detected
#       NEVER: Always alert
#
# AUTH        : Authentication Entry
#       USER  : look for user
#       ALLOW : Do not issue alert..
#               HOST - For this host
#               NET  - For this Network
#               PEAK - During PEAK times
#               OFFPEAK - during OFFPEAK times
#               ANY  - ANY time
#               NEVER - NEVER
#       NOTIFY: Issue Alert...
#               HOST - For this host
#               NET  - For this Network
#               PEAK - During PEAK times
#               OFFPEAK - during OFFPEAK times
#               ANY  - ANY time
#               NEVER - NEVER
#
# Examples:
#       AUTH USER "test2"      ALLOW NEVER                             107
#       Always report #107  when "test2" establishes a session
```

```
#
#       AUTH PASSWD "\0"      ALLOW NEVER                               107
#      Always report #107  when any session is attempted with a null passwrd
#
# REQ - Requests to NFS Service
#
#       SERVICE - Service offered by host computer
#               ACL -   for Service
#
#  Examples:
#
#       REQ     SERVICE "WINNT$"   NOTIFY ANY                           505
#       Report requests for service "WINNT$"
#
#####################################
#  SAMBA Configuration
#####################################
SMB
AUTH USER "ADMINISTRATOR" PASSWD "\0"            NOTIFY ANY     1500    1
AUTH USER "\0"            PASSWD "\0"            NOTIFY ANY     1501    1
AUTH USER "\0"            PASSWD "ADMINISTRATOR" NOTIFY ANY     1502    1
AUTH USER "\0"                                   NOTIFY ANY     1503    1
AUTH PASSWD "\0"                                 NOTIFY ANY     1504    1
REQ     SERVICE "WINNT$"   NOTIFY ANY                           1505    1
REQ     SERVICE "ROOT"     NOTIFY ANY                           1506    1
END
#####################################
#  SAMBA Configuration
#####################################
#########################################################
#########################################################
#
# ICMP -
#
# ECHO     "MAX PACKET SIZE" ACL
#       Ping of DEATH signature
```

```
#
# REDIRECT                   ACL
#       Alerts to ICMP redirects
#
#########################################################
#########################################################
ICMP
#ALLOW   NET   146.84.32.0 255.255.255.0                         1800    0
ECHO          1024   NOTIFY ANY                                 1801    1
REDIRECT             NOTIFY ANY                                 1802    1
END
RPC PCNFSD
NOTIFY ANY                           10000   3
#REQ PATH "/tmp" NOTIFY ANY          10000   3
#REQ HOST "alce" NOTIFY ANY          10001   3
END
#####################################
#  GOPHER Signature Configuration
#####################################
PORTS 70
PTYPE TCP UDP SRC DST
# Service ACL
#ALLOW  NEVER7000    3    "GOPHER traffic"
#
# Signature List
#
SIG  SRC        [Hh]ost=;/"                7010    2
SIG  SRCDST        "bin:[*!]:"           7020    3
SIG  SRCDST        "0:0:root"          7020        3
SIG  SRCDST        "sys:[*!]:"           7020        3
SIG  SRCDST        "userPassword"        7030    3
END
#####################################
#  END GOPHER Signature Configuration
#####################################
#####################################
```

```
#  IMAP Signature Configuration
#####################################
PORTS 143 220
PTYPE TCP UDP SRC DST
SIG  DST        "passwd"        22002    3
SIG  SRCDST        "bin:[*!]:"        22002    3
SIG  SRCDST        "0:0:root"        22002    3
SIG  SRCDST        "sys:[*!]:"        22002    3
SIG  DST        "bin/"        22000    4
END
#####################################
#  END IMAP Signature Configuration
#####################################
#####################################
#  POP Signature Configuration
#####################################
PORTS 109 110
PTYPE TCP SRC DST
SIG  SRCDST"bin:[*!]:"        11010    3
SIG  SRCDST"0:0:root"        11010    3
SIG  SRCDST"sys:[*!]:"        11010    3
SIG  DST        "bin/"        11000    4
SIG  SRC        "-ERR [Pp]ass"        11005    4
END
#####################################
#  END POP Signature Configuration
#####################################
#####################################
#  IDENT Signature Configuration
#####################################
PORTS 113
PTYPE TCP DST
# Service ACL
#ALLOW  NEVER        11300        3    "IDENT request"
#
# Signature List
```

```
#
SIG  SRCDST        "bin:[*!]:"         11310        3
SIG  SRCDST        "0:0:root"        11310        3
SIG  SRCDST        "sys:[*!]:"         11310       3
SIG  SRCDST     "passwd"           11305        4
SIG  SRCDST     "bin/"                 11305        4
END
#####################################
#  END IDENT Signature Configuration
#####################################
#####################################
#  NNTP Signature Configuration
#####################################
PORTS 119
PTYPE TCP SRC DST
# Service ACL
#ALLOW  NEVER        11900        3    "NNTP request"
#
# Signature List
#
SIG  SRCDST        "bin:[*!]:"        11910        3
SIG  SRCDST        "0:0:root"        11910   3
SIG  SRCDST        "sys:[*!]:"        11910   3
SIG  DST "group ‘/bin/sed:"        11920   4
SIG  DST  "830201540.9220@d"        11905        4
SIG  DST  "bin/"        11905   4
END
#####################################
#  END NNTP Signature Configuration
#####################################
#####################################
#  IRC Signature Configuration
#####################################
PORTS 194 6666 6667
PTYPE TCP UDP SRC DST
# Service ACL
```

```
ALLOW  NEVER        19400        3
SIG  SRCDST        "bin:[*!]:"        19410        3
SIG  SRCDST        "0:0:root"        19410   3
SIG  SRCDST        "sys:[*!]:"        19410   3
END
#####################################
#  END IRC Signature Configuration
#####################################
#####################################
#  LPR Signature Configuration
#####################################
PORTS 515
PTYPE TCP UDP SRC DST
# Service ACL
ALLOW  NEVER        51500        3
SIG  SRCDST        "bin:[*!]:"         51510        3
SIG  SRCDST        "0:0:root"        51510        3
SIG  SRCDST        "sys:[*!]:"         51510       3
END
#####################################
#  END LPR Signature Configuration
#####################################
#####################################
#  TALK Signature Configuration
#####################################
PORTS 517 518
PTYPE TCP SRC DST
# Service ACL
ALLOW  NEVER        51700        3
END
#####################################
#  END TALK Signature Configuration
#####################################
#####################################
#  UUCP Signature Configuration
#####################################
```

```
PORTS 540
PTYPE TCP UDP SRC DST
# Service ACL
ALLOW  NEVER        54000        3
END#####################################
#  END UUCP Signature Configuration
#####################################
#####################################
#  Kerberos Signature Configuration
#####################################
PORTS 88 750
PTYPE TCP SRC DST ASCII
# Service ACL
ALLOW  NEVER        75000        3
SIG  DST        "krbtgt"        75005 4
END
#####################################
#  END Kerberos Signature Configuration
#####################################
#####################################
#  WRITESERV Signature Configuration
#####################################
PORTS 2401
PTYPE TCP UDP SRC DST ASCII
# Service ACL
ALLOW  NEVER        240100 3
END
```

```
#####################################
#  END UUCP Signature Configuration
#####################################
#####################################
#  Kerberos Signature Configuration
#####################################
PORTS 88 750
PTYPE TCP SRC DST ASCII
# Service ACL
ALLOW  NEVER        75000        3
SIG  DST  "krbtgt"        75005        4
END
#####################################
#  END Kerberos Signature Configuration
#####################################
#####################################
#  WRITESRV Signature Configuration
#####################################
PORTS 2401
PTYPE TCP UDP SRC DST ASCII
# Service ACL
ALLOW  NEVER        240100        3
SIG  SRCDST        "bin/"        240105      3
END
```

Recommended Security Resources

Following is a list of recommended reading on the subject of network security. Most of these resources provide detailed information about network protocols and how they relate to security. The web sites listed also provide information regarding different attacks and how to protect your network from them.

Books

Albitz, Paul and Cricket Liu. *DNS and BIND in a Nutshell.* Sebastopol: O'Reilly & Associates, Inc., 1992.

Anonymous. *Maximum Security - A Hacker's Guide to Protecting Your Internet Site and Network*. Indianapolis: Sams.net, 1997.

Chappell, Laura A. and Roger L. Spicer. *Multiprotocol Internetworking*. San Jose: Novell Press, 1994.

Comer, Douglas E. *Internetworking With TCP/IP Vol.1: Principals, Protocols, and Architecture*. Third Edition. New Jersey: Prentice Hall, 1995.

Nye, Adrian. *X Protocol Reference Manual*. Sebastopol: O'Reilly & Associates, Inc., 1995.

Roberts, Dave. *Internet Protocols Handbook*. Scottsdale: The Coriolis Group, Inc., 1996.

Stevens, W. Richard. *UNIX Network Programming*. New Jersey: Prentice Hall, 1990.

Requests for Comments (RFCs)

Network Working Group. Request for Comments #1001. *Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods*. Internet Activities Board (End-to-End Services Task Force), March 1987.

Network Working Group. Request for Comments #1002. *Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications*. Internet Activities Board, March 1987.

Network Working Group. Request for Comments # 1014. *XDR: External Data Representation Standard*. Sun Microsystems, Inc., June 1987.

Network Working Group. Request for Comments #1035. *Domain Names - Implementation and Specification*. P. Mockapetris, November 1987.

Network Working Group. Request for Comments # 1057. *RPC: Remote Procedure Call Protocol Specification Version 2*. Sun Microsystems, Inc., June 1988.

Network Working Group. Request for Comments # 1094. *NFS: Network File System Protocol Specification*. Sun Microsystems, Inc., March 1989.

Network Working Group. Request for Comments # 1833. *Binding Protocols for ONC RPC Version 2*. R. Srinivasan, August 1995.

Web Sites

http://astalavista.box.sk

http://www.infowar.com

http://www.l0pht.com

http://www.ntbugtraq.com

Cross-Site ACLs and Roles

This section provides reference information about the Cross-Site roles and Access Control Lists (ACLs). Each ACL description includes the ACL's purpose and the default roles that are assigned to the ACL.

The second part of this section lists the preconfigured roles that are installed with Cross-Site and the ACLs to which each role is assigned.

ACLs and Their Roles

The following list includes all of the Cross-Site ACLs and the roles assigned to them by default. You can confirm this list by clicking on the **Permissions** tab in the **Admin** view.

Auth ACL

The roles assigned to this ACL include: **admin**, **install**, and **user**.
The Auth ACL provides authentication and authorization services for the Cross-Site console and clients. The following is a list of services the Auth ACL provides.

- Create a user
- Import a user
- Create a Deployment or Availability agent
- Create a Security agent
- Create a group
- Add a principal (group, user, agent) to a group
- Add and remove a principal from a group

- Delete a principal
- Define a foreign management server in a local domain
- Delete a foreign management server in a local domain
- Set the password for a principal
- Set a role for a principal
- Add a role for a principal
- Remove roles from principal
- Create a new role in the local domain
- Create a foreign role in the local domain
- Define a role in the local domain
- Delete a foreign role in the local domain

EventDispatcher ACL

The role assigned to this ACL is **any**.

This ACL controls who can subscribe to events, register as a provider of events, and notify the dispatcher of events.

The methods associated with this ACL are as follows:

- Record an event subscription in the subscriber memory
- Remove an event subscription from the subscriber memory
- Record an event registration in the registration memory
- Remove a registered event from the registration memory
- Notify the subscribers that a specified event has occurred

EventAdmin ACL

The role assigned to this ACL is **admin**.

This ACL controls what can start and stop the event log. The methods are as follows:

- Start the event log
- Stop the event log

EventHandler ACL

The role assigned to this ACL is **any**.

This ACL forwards notification to the Event Handler’s subject.

Launcher ACL

The role assigned to this ACL is **admin**.

The Launcher ACL controls access to the Cross-Site “launch service” methods. The following is a list of Launcher methods:

- Determine whether a task is currently active
- Reschedule a task
- Launch a task and reschedule it
- Return the completion status of a task

Scheduler ACL

The role assigned to this ACL is **admin**.

The Scheduler ACL is responsible for scheduling internal and external Cross-Site tasks.

The Scheduler ACL is divided into two components. The first is a set of administrative methods: start, stop, and isActive. These are used to control and monitor the Scheduler service itself. The second component is a set of job methods: addJob, removeJob, and isScheduled. These allow a management server console to manipulate the Scheduler’s job queue. The following is a list of Scheduler methods.

- Start the Scheduler service
- Stop the Scheduler service
- Determine whether or not the management server’s Scheduler service is active
- Add a job
- Add a job with a specified timestamp
- Remove a job from the scheduler service’s job queue
- Determine whether a task is in the queue

- Determine whether a task is being executed
- Stop a task

Agent ACL

The role associated with this ACL is **securityAgent**.

The Agent ACL controls access to the interface between the Security agent and the Cross-Site management server. The following is a list of the Agent methods.

- Fetch the agent’s configuration file from the management server
- Send the management server a heartbeat
- Notify the management server of a detected intrusion attempt

Avail ACL

The role associated with the Avail ACL is **availability**.

The Avail ACL provides availability upload services to the Cross-Site for Availability agents. This ACL uses the post method.

Roles and Their ACL Entries

The following is a list of the default roles and the ACL entries, or methods, they are assigned. Each list of ACLs is organized alphabetically.

admin Role

The ACLs and methods to which this role is assigned are as follows:

- Auth: addPrincipalToGroup
- Auth: addRemovePrincipals
- Auth: addRolesForPrincipal
- Auth: clearMgmtServer
- Auth: clearRole
- Auth: createAgent
- Auth: createGroup
- Auth: createRole

- Auth: createSecurityAgent
- Auth: createUser
- Auth: defineMgmtserver
- Auth: defineRole
- Auth: deletePrincipal
- Auth: deleteRole
- Auth: dropRolesFromPrincipal
- Auth: getCredentials?
- Auth: getForeignCredentials?
- Auth: importUsers
- Auth: removePrincipalFromGroup
- Auth: setPassword
- Auth: setRolesForPrincipal
- Auth: updateAgent
- Auth: updateUser
- Event Admin: startEventLog
- Event Admin: stopEventLog
- Launcher: getCompletionStatus
- Launcher: isActive
- Launcher: launch
- Launcher: reschedule
- Scheduler: addJob
- Scheduler: isActive
- Scheduler: isExecuting
- Scheduler: isScheduled
- Scheduler: removeJob

- Scheduler: start
- Scheduler: stop
- Scheduler: terminate
- License: getAvailabilityEndpoints
- License: getDeploymentEndpoints
- License: getSecurityEndpoints
- License: hasValidKey
- XSITagent: gateway

user Role

The ACLs and methods to which this role is assigned are as follows:

- Auth: createAgent
- XSITagent: gateway

install Role

The ACLs and methods to which this role is assigned are as follows:

- Auth: createSecurityAgent
- Auth: updateAgent

any Role

The ACLs and methods to which this role is assigned are as follows:

- License: isAvailabilityEnabled
- License: isDeploymentEnabled
- License: isSecurityEnabled
- EventDispatcher: notify
- EventDispatcher: register
- EventDispatcher: subscribe
- EventDispatcher: unregister

- EventDispatcher: unsubscribe
- EventHandler: notify
- ChanCopy: gateway
- ChannelCopier: gateway
- IdsRt: gateway
- XsiteCertificateMgrChannel: gateway
- XSiteChanMonitor: gateway
- XSiteChanStarter: gateway
- XSiteConsole: gateway
- XSiteCopyWrapper: gateway
- XSiteCpAgentinstalls: gateway
- XSitePopups: gateway
- SSiteTransmitter: gateway
- XSiteTuner: gateway

securityAgent Role

The ACLs and methods to which this role is assigned are as follows:

- Agent: alert
- Agent: getConfigFile
- Agent: indy
- Agent: notify
- Agent: upload

availability Role

This role is assigned to

- Avail: post

Signatures Detected by Cross-Site for Security

Note: *Not all of the signatures that the Cross-Site for Security agent can detect are documented in this appendix. There was insufficient information about some signatures at time of publication of this user's guide, so those signatures were not included. If you require information about signatures that are not documented, see "Recommended Security Resources" on page 223. That section lists recommended reading on network security and includes several web sites that provide information about different attacks and how to protect your network from them.*

This appendix describes the signatures that the Cross-Site for Security agent can detect while monitoring your network. A *signature* is a string of characters that might represent a potential intrusion attempt. A signature can represent normal network behavior, but it could also represent suspicious activity. Signatures are specified in the *base_dir***/XSITsagt/lib/policyfiles/5/ids.rules** file, where the *base_dir* is the base directory of the Cross-Site server. The **ids.rules** file comprises the core of the policy that you assign to a Security agent. You assign Security policy to an agent using the **Policy** explorer on the Cross-Site console.

Alerts are generated when the policy applied to a Security agent is violated. Policy violations occur when a signature is detected by a Security agent, and the signature is defined in the **ids.rules** file as one that may represent an intrusion attempt. Policy violations are also referred to as *incidents*.

Each time a Security agent detects an incident, it logs an alert in a local log file. If the alert is assigned a priority of 1, it is considered *critical*. You designate which incidents should generate critical alert messages by the severity level you assign to each signature in the **ids.rules** file, and the priority level you specify for each alert message on the **Policy** explorer. Severities and priorities range from 1 to 5, with 1 representing the most serious type of incident. An alert should have a priority of 1 if the signature represents activity that you should not see under normal circumstances. The Security agent *immediately* sends critical alerts to the alert log, which is stored in the Cross-Site management repository. All other alerts are uploaded to the alert log periodically.

To enable you to view detailed information about a particular type of incident, the associated signature must have a severity of 1, so that an event will be generated on the Cross-Site console when this type of incident occurs. On the management server, critical alerts are converted to events, forwarded to the Event service, and assigned an event priority. Critical alerts are treated as high priority events and are assigned an event priority of 5. Events are displayed in two places on the console: in the event log and on the **Events** explorer.

The signatures in the **ids.rules** file are organized into sections, by application or service (such as Telnet/R, FTP, and Gopher). Within many of these sections, there are signature lists that specify signature strings, followed by the alert message numbers associated with each signature, and the severity of the incident each signature represents. For example, the following is a signature for the Telnet/R service in the **ids.rules** file:

```
SIG SRCDST    "/phf?Qalias"      4024         1
```

where:

`SIG` Identifies the entry as a signature.

`SRCDST` Indicates that the Security agent is monitoring both inbound (SRC) and outbound (DST) traffic for this signature string.

`"phf?Qalias"`
Specifies the character string being detected.

`4024` Represents the number associated with the alert message that Cross-Site generates when the signature is detected. This number and the corresponding alert message are displayed on the console’s **Events** explorer when you select the **Messages** item in the explorer tree.

`1` Represents the severity of the incident that the signature represents.

This document follows a similar organization, but provides additional information to help you understand the nature of the attack the signature may represent. The following information is provided for each signature that a Security agent can detect:

Signature The character string that the Security agent detects

Severity The severity of the intrusion attempt that the signature represents. The severity of each signature is specified in the **ids.rules** file. Severities range from 1 to 5, with 1 representing the most serious type of incident. An event is generated on the Cross-Site console for an incident with a severity of 1, because this type of incident requires immediate attention.

Implications What the signature might mean to you

Identifying the Source

How you can attempt to identify the source of the signature

What You Can Do

The preventive measures that you can take to protect yourself against such intrusion attempts

In addition to using Cross-Site for Security to monitor your network for possible intrusions, there are several basic good-business practices that you should adhere to when possible, to minimize system vulnerabilities. Ask yourself the following questions to determine whether you are taking proactive measures to protect yourself against intrusions:

- Are all available patches for your operating systems and applications installed?
- Are there a minimum number of people with root access?
- Is there a “smart card” or one-time password for root access?
- Are passwords changed regularly?
- Are old users deleted?
- Do you know where your sensitive information is and how well it is protected?

You can also protect network servers by using firewalls and routers to create barriers between your networks, and between your intranets and the Internet. This prevents easy access to the exchanges made between servers and clients. Finally, set file and directory permissions to ensure that passwords that might be compromised allow intruders access to a minimum amount of data.

Telnet/R Services (and X11) Signatures

Intruders use most X11 attacks to overwrite a buffer and obtain root access.

Signature	SRCDST “rpcinfo -p”
Severity	1
Implications	Intruders run this command against a set of targets to find NFS servers, which may or may not be protected by a filter that blocks traffic to the portmapper. Intruders can find out if the host is running NIS, and if it is an NIS server or slave. They can also discover whether a diskless workstation is around, and whether it is running NFS, or any information services (such as ruserd and rstatd).

Identifying the Source	If this command is done from a Telnet shell, then the login ID will be identified. Also the source IP address should be identified with this alert and it will likely not be spoofed. Since there will also be a time stamp on this alert, several mechanisms are available for identifying the source of the problem.
What to Do	While this isn't an overtly bad signature to detect, it should be considered suspicious. System administrators use this tool to check to see if certain RPC services are up and running. Monitor the traffic in the future, and, if further suspicious behavior is detected, take corrective administrative actions.

Signature	SRCDST "tprof -x"
Severity	1
Implications	Intruders use this command to exploit a vulnerability in tprof for AIX. By entering **tprof -x /bin/sh** on AIX, intruders can access a root shell.
Identifying the Source	If this command is done from a Telnet shell, then the login ID will be identified. Also the source IP address should be identified with this alert and it will likely not be spoofed. Since there will also be a time stamp on this alert. Several mechanisms are available for identifying the source of the problem.
What to Do	This is significant attack that exploits a hole in the tprof program. Any occurrence of this attack should be considered highly questionable and the context of when and why this command occurred should come into question. Keep in mind that falsing can occur, so the surrounding context of this signature should be examined and considered. If the command occurred in a context where it was obviously a user command, then this is a serious occurrence that should be examined further.

I

Signature	SRCDST "expn"
Severity	1

Implications	Intruders use this command to make CPU usage rise to 99 percent and eat all available memory and disk space. The attacker opens a Telnet connection to port 25, issues **helo**, **mail from:** and **rcpt to:** commands. Instead of data, the command uses **expn *@**. The software goes into an infinite loop. Also, this command can be used to implicitly obtain user IDs from the system.
Identifying the Source	This is difficult because, by default, sendmail accepts a message from any incoming connection. The sender of such a message can, therefor, appear to have originated anywhere. Any claim of identity is accepted, so a message's originator can be easily forged. However, the IP address and a time stamp will be available with the alert. Some degree of analysis can be conducted to find the source of the signature.
What to Do	Disable the **expn** command by editing the **sendmail.cf** file. Add the following line and restart the sendmail service: `O PrivacyOptions=needmailhelo, noexpn`

Signature	SRCDST "vrfy"
Severity	1
Implications	Intruders use this command to perform a Denial of Service (DoS) attack. Intruders use Telnet to connect and issue this commend to set the processes running at 100 percent. Also this command can be used to implicitly obtain user IDs from the system.
Identifying the Source	The IP address and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. Also, sendmail logging can be activated with the addition of the line "mail.info<TAB>/var/log/syslog" to the /etc/syslog.conf file. This will help identify the originator of further e-mail.
What to Do	Upgrade to the latest version of sendmail that doesn't suffer from this vulnerability. Another option would be to just not run sendmail on the machine; instead reroute mail through or to some other machine.

Signature	SRCDST “mailfrom:|” SRCDST “mailfrom:’|” SRCDST “mailfrom:\”|”
Severity	1
Implications	Intruders use these commands to exploit the sendmail pipe bug.
Identifying the Source	The IP address and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. Also sendmail logging can be activated with the addition of the line “mail.info<TAB>/var/log/syslog” to the /etc/syslog.conf file. This will help identify the originator of further e-mail.
What to Do	This is a highly suspect command if the system doesn’t currently have a supported software application installed that uses sendmail to communicate with. If this signature is detected, then take into account the entire signature and try to find out what program was activated. The path name to the program that is being activated or was activated is indicated to the right of the vertical bar in the signature. If the activation of this program is suspect, then take appropriate administrative measures.

Signature	SRCDST “finger root”
Severity	1
Implications	Intruders use this command to learn whether root is currently logged into the machine and also to obtain plan information on root. On some systems this command can be used to obtain root. However, this is not necessarily a blatant attack of the system as much as a suspicious command.
Identifying the Source	The IP address and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	The IP address and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.

Signature	SRCDST “finger @”
Severity	1
Implications	Intruders use this command to perform a Denial of Service (DOS) attack on Solaris 2.51 and 2.6 systems, including Sun x86. By supplying a number of host names (e.g., **finger @host@host@host**), intruders can greatly increase the number of processes running and overload the system.
Identifying the Source	The IP address and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	This is a blatant attack and should be treated as such, except for the situation when the finger command is being forwarded through a firewall or other proxy. However, if all host names listed are ordinary machines, then this is clearly an attack that should be responded to.

Signature	SRCDST “bin/phf”
Severity	1
Implications	Intruders use the following command in an http request to query for the **passwd** file: `"/cgi-bin/phf?Qalias=x%0a/bin/cat%20/etc/passwd"` By exploiting the **cgi-handler** bug, a malicious user could start an x-term from the server machine on his/her own system. The intruder may then be able to execute commands with the privileges of the httpd daemon, compromise the http server, and gain privileged access.
Identifying the Source	The IP address and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	If you have the **phf** program installed on your web server and do not require it, you should disable it.

Signature	SRCDST “/phf?Qalias”

Severity	1
Implications	Intruders could use the following command in an http post command to execute any command, including a query for the **passwd** file: `"/cgi-bin/phf?Qalias=x%0a/bin/cat%20/etc/passwd"` A malicious user could scan for hosts that still have the **phf** bug, which gives **etc/passwd**. The main reason for using **phf** on the system is to exploit this bug to execute commands. It can also be used to obtain information about the person calling the script
Identifying the Source	The IP address and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	If you have the **phf** program installed and do not require it, you should disable it.

Signature	SRCDST "/finger"
Severity	1
Implications	Either someone is executing a custom finger program or could be connecting to a web server from a Telnet session and trying to do an HTTP GET of a "/finger" CGI command. The first option, while suspicious, doesn't necessarily warrant further action. However, the latter scenario does and the context of the signature should indicate whether further investigation is needed. If it looks like the user was simply executing a custom finger command from a prompt, there is no need to investigate. However, if the signature was part of an obvious HTTP GET command, then it would be highly questionable.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	If the signature occurred in the context of an HTTP GET, then take appropriate investigatory action.

Signature	SRCDST "/date"

Severity	1
Implications	Either someone is executing a custom date program or could be connecting to a web server from a Telnet session and trying to do an HTTP GET of a “/date” CGI command. The first option while suspicious, probably doesn’t necessarily warrant further action. However, the latter scenario does and the context of the signature should indicate whether further investigation is needed. If it looks like the user was simply executing a custom date command from a prompt, there is no need to investigate. However, if the signature was part of an obvious HTTP GET command, then it would be highly questionable.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell then, the login ID will be identified.
What to Do	If the signature occurred in the context of an HTTP GET, then take appropriate investigatory action.

Signature	SRCDST “/nph-test-cgi”
Severity	1
Implications	Intruders use this command to obtain a list of files on a remote system. Many web sites use a file named **test-cgi**, which is usually in the **cgi-bin** directory. Many **test-cgi** files contain the following command line, which makes the system vulnerable: `echo QUERY_STRING=$QUERY_STRING` Without the quotes, special characters such as (*) get expanded. Therefore, submitting a query of ‘*’ returns the contents of the current directory, where most likely all of the cgi files are stored. Submitting a query of ‘*/’ lists the root directory, etc.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	Place quotes (**"**) around all variables in the **test-cgi** file. For example, `echo QUERY_STRING="$QUERY_STRING"`

Signature	SRCDST “/test-cgi”
Severity	1
Implications	Intruders use this command to obtain a list of files on a remote system. Since there is no equivalent command for this signature on a typical command line, this command should at all times be considered the result of an attack attempt. Especially if the ports associated with this string are indicating a probing connection to an http server.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so, some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login id will be identified.
What to Do	First, consider the surrounding signature context. If this appears to be part of a probing http connection, then identify the source of the attack and take appropriate administrative monitor or other action.

Signature	SRCDST “?%0A” SRCDST “?%0a”
Severity	1
Implications	Intruders use these commands to exploit a generic newline problem in httpd.
Identifying the Source	The IP address, port, and a time stamp will be available with the alerts, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	If the port connected to is a WWW server port and the context of this string is one where it is apparent that the signature indicates an HTTP GET or POS, then this is clearly a questionable signature. This could also be part of a script that is displayed in an editing session and, hence, a questionable signature. If this signature is seen, continue monitoring the subject for other signatures.

Signature	SRCDST “IFS[ ]=[ ]”
Severity	1
Implications	Intruders use this command to reset the Internal Field Separator (IFS) variable so that the command interpreting shell script (that may have SUID privileges) can be tricked into executing commands other than what was intended. This is a very popular system-hacking technique.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	While there might be some perfectly legitimate reasons for resetting the IFS shell variable in a shell script, it is more likely that this variable is reset for non-legitimate reasons. Therefore, it would be most useful if this signature was taken in its proper context. When in doubt about the context in which this signature could have been detected, then further investigate further.

Signature	DST “dig axfr”
Severity	2
Implications	Intruders use this command to attempt a zone transfer attempt. A zone transfer attempt elicits more information than a normal user would typically require. Therefore whenever a broad, sweeping request for information, such as this, comes from the network, it should be considered with caution and in its proper context.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	If this signature occurs in the context of a normal Telnet or R-service session, along with other signatures, then it will likely require immediate action. However, if this signature is detected in the context of a system administrator session, then it is likely that this is a relatively benign signature.

I

Signature	SRCDST “chmod 2”
Severity	2
Implications	Intruders use this command to change permissions to a SGID (set group id) file. Note, however, that this command doesn’t really reset the SGID bit if there are not three following octal numbers (digits less than 8). If an attacker is successful in executing an SGID script, then he/she can potentially read and write files with the group permissions and, therefore, compromise security to some degree.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	First, look at the signature to determine if this is a legitimate attack. If there are three following octal numbers, then it could possibly be an attack. It could also be execution of a legitimate command. Investigate further to determine whether there was a good reason for the command to be executed.

Signature	SRCDST “chmod 4”
Severity	2
Implications	Intruders use this command to try to change the permissions on an SUID (set user id) file, so they can edit it or read it. Because most SUID programs on Solaris are dynamically linked, intruders can gain root access privileges this way. Note, however, that this command doesn’t reset the SGID bit if there are not three following octal numbers (digits less than 8).
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	First, look at the signature to determine if this is a legitimate attack. If there are three following octal numbers, then it could possibly be an attack. It could also be execution of a legitimate command executed. Investigate further to determine whether there was a good reason for the command to be executed.

Signature	SRCDST “chmod 6”
Severity	2
Implications	Intruders use this command to change permissions on the UID (user id) file or SGID (set group id) file, so they can edit it or read it. As a result, they can find out who has root access and grant themselves access. Note, however, that this command doesn’t really reset the SGID bit if there are not three following octal numbers (digits less than 8).
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	First ,look at the signature to determine if this is a legitimate attack. If there are three following octal numbers, then it could possibly be an attack. It could also be execution of a legitimate command. Investigate further i to determine whether there was a good reason for the command to be executed.

Signature	DST “.:$PATH”
Severity	1
Implications	Intruders use this command to attempt to reset the path, possibly within a script, so that the current directory overrides the default PATH variable. This is a highly risky thing to do because of the risk of a Trojan horse execution, and should not be done under normal circumstances.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	Consider the actual signature and determine if it could have a legitimate explanation. Also consider the login, if available. Perhaps a find script that would search for this signature in a particular directory tree would help determine the conditions where this signature was generated.

Signature	DST “rlogin -froot”
Severity	3
Implications	Intruders use this command to exploit a bug that enables them to overwrite a buffer in **gethostbyname()** on Solaris 2.5.1, which gives them a root shell.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	Following are some commands and options you can use to protect yourself against such intrusion attempts. Refer to the manual pages for more information on a command and its options. **rlogind** use ‘**-l**’ to disable validation using **.rhosts** files **fingerd** use ‘**-l**’ to log all connections use ‘**-S**’ to suppress information about login status, home directory, and shell use ‘**-f msg-file**’ to make it just display that file **rshd** use ‘**-a**’ to verify that all incoming remote host names and addresses match use ‘**-l**’ to disable validation using **.rhosts** files use ‘**-L**’ to log all access attempts to syslog

Signature	DST “lquerypv -h”
Severity	2
Implications	Intruders use this command to read any file on an AIX system v4.1 and v4.2. When invoked with the "-h" option, "lquerypv" does not adequately enforce the read permissions on files when it is run by regular (non-"root") users. This can allow users to obtain access to the contents of files that they are not authorized to read.

Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	A vendor patch is available at **http://www.rootshell.com/aix_patch.html**

Signature	DST “workman -p”
Severity	2
Implications	Intruders use this command to exploit a vulnerability in Linux or SysVR4 workman that allows any user to create and write to files owned by the user who is running workman.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	Remove this application from your system unless it is absolutely necessary. If it is, then run a wrapper script around this application and don’t allow world execute directly on this application.

Signature	SRCDST “invalid login”
Severity	3

<table>
<tr>
<td><b>Implications</b></td>
<td>
<p>Intruders generally try to access a system by logging in to several different accounts. They usually start with the default accounts that come with the operating system.</p>
<p>Some of the default accounts are as follows:</p>
<table>
<tr><th>Account</th><th>Password</th></tr>
<tr><td>root</td><td>root</td></tr>
<tr><td>sys</td><td>sys / system / bin</td></tr>
<tr><td>bin</td><td>sys / bin</td></tr>
<tr><td>mountfsys</td><td>mountfsys</td></tr>
<tr><td>adm</td><td>adm</td></tr>
<tr><td>uucp</td><td>uucp</td></tr>
<tr><td>nuucp</td><td>anon</td></tr>
<tr><td>anon</td><td>anon</td></tr>
<tr><td>user</td><td>user</td></tr>
<tr><td>games</td><td>games</td></tr>
<tr><td>install</td><td>install</td></tr>
<tr><td>demo</td><td>demo</td></tr>
<tr><td>umountfsys</td><td>umountfsys</td></tr>
<tr><td>sync</td><td>sync</td></tr>
<tr><td>admin</td><td>admin</td></tr>
<tr><td>guest</td><td>guest</td></tr>
<tr><td>daemon</td><td>daemon</td></tr>
</table>
<p>The root, mountfsys, umountfsys, install, and, sometimes, sync accounts are root-level accounts, which means they grant total power over the operating system. Other logins are user-level account logins, which means they have power over only the files or processes that they own. Command logins do not let you into the operating system, but execute the program associated with the login. For example, the <b>rwho</b> and <b>finger</b> commands display the account names of users that are currently online. This shows intruders several existing accounts.</p>
</td>
</tr>
</table>

Identifying the Source	When someone logs in incorrectly, the error log on the system is updated. This lets the system administrator know that something suspicious could be happening. The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	If a number of these alerts occur in a small time frame on the same IP address and login, then this is probably an attack and proper security measures should be taken. However, if fewer than three are seen, then this could be associated with a user merely forgetting his/her password.

Signature	SRCDST “ission [Dd]enie”
Severity	1
Implications	This is a permission warning. It generally means that the user attempted something and was told by some system that she/he was not allowed to do the attempted operation. This is a good general purpose signature to look for that might shed light on illegal activity.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. If this command is done from a Telnet shell, then the login ID will be identified.
What to Do	Consider the source and destination IP addresses. If either is a sensitive address, then consider under what context this string could have been received. For example, if one of the IP addresses is the address for a router, then perhaps the intruder is attempting to reconfigure the router. Consider the context of this intrusion alert and take appropriate measures.

FTP Signatures

The File Transfer Protocol (FTP) transfers files to and from a remote network site. A vulnerability in ftp client software enables a malicious ftp server to trick the ftp client into executing arbitrary commands. Patches are available from Sun’s web site at: http://sunsolve.sun.com/sunsolve/pubpatches/patches.html.

Signature	DST “[Ee] [Xx] [Ee] [Cc]”
Severity	2
Implications	Intruders exploit a vulnerability in the SITE EXEC command that allows any remote and local user to run commands as root on the target host. The SITE EXEC vulnerability affects ftpd only if the SITE EXEC command feature has been explicitly activated at your site. Support for anonymous ftp is not required to exploit this vulnerability.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	This problem can be eliminated by upgrading to the latest version of ftpd. Otherwise, disable ftpd (comment out in /etc/inetd.conf) until the problem is corrected.

Signature	SRC “uest login”
Severity	3
Implications	Some systems grant limited account status if you log on as guest or anonymous, as well as some of the basic default account names such as demo, games, account, admin, user, basic, and main or maint (for maintenance). By default, a guest user or IUSR_WWW has read access to all files on an NT disk. These files can be read, executed, or downloaded by guests. Under normal conditions, if a system has been designated as a guest-login-enabled machine, then this should be a significant issue. However, if no designation has been given and the given machine isn’t set up for this kind of use, then this should be considered significant and worthy of further investigation. As well, everyone has remote access to an NT system’s registry by default.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.

What to Do	Check for the existence of the following registry key: `<HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurePipeServers\Winreg>` If this registry key does not exist, remote access is not restricted; only the underlying security on individual keys controls access. Versions of NT that are affected include NT 3.5, 3.51, and 4.0.

Signature	DST “CWD ~root”
Severity	3
Implications	Intruders use this attack to access the root directory of a file system. From that directory, intruders can delete log files, as well as the application that is supposed to detect intruders.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	This kind of traffic is rarely explainable within normal FTP traffic and should be considered an attack under most cases. Unless a logical explanation as to why a user would try to switch to this directory is evident ,then immediately begin to investigate and take proper administrative action.

Signature	DST “SITE”
Severity	3
Implications	A vulnerability in the SITE EXEC command feature of ftpd allows any remote or local user to obtain root access. There is also a vulnerability due to a race condition in these implementations. Support for anonymous FTP is not required to exploit this vulnerability.
Identifying the Source	The IP address, port, and a time stamp will be available with the alerts, so some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login id will be of questionable use.

What to Do	Upgrade the latest version of ftpd to ensure that this bug is not an opportunity for exploitation on the given system.

Signature	DST “passwd”
Severity	2
Implications	This could indicate an attempt by an attacker to obtain the password file so he/she can attempt to decode the passwords contained within.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	Ensure that the ftp daemon is unable to read the /etc directory or ensure that the user who generated this signature is not a security risk.

Signature	DST “MKD”
Severity	2
Implications	Intruders use the **MKD** command to create a directory on the system, into which they can then copy malicious scripts and Trojan horse programs.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login id will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	Make sure that this command is not enabled for anonymous users. If this is a signature that is found frequently, then the contents of these directories should be examined.

Signature	DST “[Pp] [Aa] [Ss] [Vv]”
Severity	2

Implications	The FTP service permits passive connections to be established based on the port address given by the client. When multiple passive connections are made to a single NT FTP server using the PASV FTP command, it is possible to use up all available system threads. Requests for additional connections will fail and continue to do so until a client thread is made available. If the system isn't an NT machine, then this signature is probably benign.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	The registry contains an entry in **<System\CurrentControlSet\Services\MSFTPSVC\Parameters>** where the value can be enabled for **<EnablePortAttack: REG_DWORD:>**. Verify that this value is **0** and not **1**.

Signature	DST ".rhosts"
Severity	2
Implications	Intruders use the get command to access the sensitive **.rhosts** system file, so they can determine what systems can log on to the system without authentication. Intruders gain access to a host by exploiting a **+** in the **.rhosts** file. Any occurrence of the **+** character in a .rhosts file is an opportunity for an intrusion if all machines that are granted this access are compromisable.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	If this intrusion signature is detected, it is highly unlikely that there is an appropriate explanation for it. If a reasonable explanation cannot be offered to explain the signature, then take appropriate investigatory or administrative actions.

Signature	DST "hosts.equiv"

Severity	2
Implications	Intruders gain access to a host by exploiting a **+** in the **/etc/hosts.equiv** file. Like the .rhosts file above this file is just as sensitive. See the above signature.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	If this intrusion signature is detected, it would be highly questionable that there is an appropriate explanation for it. If a reasonable explanation cannot be offered to would explain the signature, then take appropriate investigatory or administrative actions.

Signature	DST “../..”
Severity	2
Implications	Intruders use this command to skip past a designated directory point that should be the only directory that is viewable for anonymous ftp. Normally this is ineffective and, under certain circumstances, this signature can have a perfectly legitimate explanation.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature. The login ID will be identified; however, if this is an anonymous ftp site, the login ID will be of questionable use.
What to Do	Investigate whether there could be some reasonable explanation for this signature, such as a user that has descended two or three levels deep in the directory hierarchy needing to cd two levels up.

FTP Data Signatures

The data channel of the ftp protocol is responsible for moving files from one location to another. It is valuable to be able to monitor the contents of these files for possibly sensitive data.

Signature	SRCDST “bin:[*!]:” SRCDST “0:0:root” SRCDST “sys:[*!]:”
Severity	5
Implications	The above signatures are associated with a /etc/passwd file common to most UNIX systems. A common tactic of hackers is to find some way to get the password file so they can run some sort of cracker program on it. One of the ways they may get this file is via FTP. Therefore, if the agent is monitoring for signatures associated with a password file coming over FTP, then this could be a useful alert.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Investigate further to determine if there was some good reason for receiving this signature. Otherwise, try to identify who might have transferred the file.

Finger Signatures

Finger is an old protocol used to allow remote users to find out what users may be on a particular machine and to get details about particular users on a machine.

Signature	DST “bin/”
Severity	4
Implications	This signature could be an attempted buffer overflow attack against a fingered program or a back-door to run a program on the machine. In any event, this signature shouldn’t be in the client traffic to a box that is being fingered.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Consider the entire signature that is reported and attempt to identify whether a reasonable explanation exists as to why this signature would be detected.

Signature	DST “planner”
Severity	5
Implications	This is a possible rootkit signature that might indicate that a system has been compromised with an off-the-shelf rootkit.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If no reasonable explanation as to why a finger signature of planner is being attempted (such has the host name is “planner” or a user named “planner” uses this machine), then take immediate administrative action to isolate the machine and evaluate whether the machine has been compromised.

TFTP Signatures

The Trivial File Transfer Protocol (TFTP) provides remote access to a host’s files without prompting users for a password. While this feature is useful when booting diskless workstations, it is also a security risk. Because TFTP is not authenticated, attackers can more easily gain access to sensitive files over the network.

This vulnerability is inherent in the design of the service. Therefore, it is best to disable TFTP completely by placing a # at the beginning of the **tftp** line in **/etc/inetd.conf**. Otherwise, ensure that only a limited portion of the file system is available to TFTP users by changing the root directory when the tftp daemon is executed. See your TFTP documentation for details.

Signature	DST “passwd” DST “shadow”
Severity	3
Implications	Intruders use this command to gain access to the password file. Once they have access to this file, they can then run several cracking programs on the encrypted passwords.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	See “TFTP Signatures” on page 256 for instructions.

Signature	DST “inetd.conf” DST “master” DST “.rhosts” DST “config” DST “hosts.equiv” DST “../..”
Severity	3
Implications	Intruders use these commands to access sensitive system files, such as password files. If they gain access to these files, intruders can use this information to obtain unauthorized access to other sensitive information: everything from private e-mail to databases.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Comment out (remove) the tftp daemon from the inetd.conf file.

Signature	SRCDST “bin:[*!]:” SRCDST “0:0:root” SRCDST “sys:[*!]:”
Severity	5
Implications	See “FTP Signatures” on page 249.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	See “TFTP Signatures” on page 256 for instructions.

Signature	SRCDST “userPassword”
Severity	3

Implications	Intruders use this command to gain access to the router password file. Once they have access to this file, they can use this information to reconfigure the router to allow traffic from additional ports or hosts, or otherwise modify its functionality, increasing their ability to access the network and, therefore, your vulnerability. If an intruder forces the router to reboot, she/he can also deny service to legitimate users. Routers provide system administrators a facility to manage network connections between computers. As part of a firewall system, the router is generally configured to allow traffic only from specific ports on the host.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert. Hence, some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove TFTP access on the router in question. Also identify the source of the intrusion and conduct a proper investigation.

WWW Signatures

Signature	DST “bin/phf”
Severity	5
Implications	Intruders use the following command to query for the **passwd** file: `set pwQuery "/cgi-bin/phf?Qalias=x%0a/bin/cat%20/etc/passwd"` By exploiting the **cgi-handler** bug, a malicious user could start an x-term from the server machine on his/her own system. Intruders may then be able to execute commands with the privileges of the httpd daemon. They may compromise the http server and gain privileged access.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If you have the **phf** program installed and do not require it, disable it.

Signature	DST “bin/finger”
Severity	1
Implications	Intruders use this command to create a list of users who are currently online. While this attack isn’t very malicious, it is rarely a command used as part of a web page and, therefore, should be considered.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert., so some degree of analysis can be conducted to find the source of the signature.
What to Do	Find out if the signature is part of a web page and also find out if the finger command is on your server and available from the web server. If so, then this may not be a problem; however, if either the command is not available or not part of a web page, then this should be considered a serious attack attempt.

Signature	DST “nph-test-cgi”
Severity	4
Implications	Intruders use this command to obtain a list of files on a remote system. Many web sites use a file named **test-cgi**, which is usually in the **cgi-bin** directory. Many **test-cgi** files contain the following command line, which makes the system vulnerable: `echo QUERY_STRING=$QUERY_STRING` Without the quotes, special characters such as (*****) get expanded. Therefore, submitting a query of ‘*****’ returns the contents of the current directory, where most likely all of the cgi files are stored. Submitting a query of ‘***/**’ lists the root directory, etc.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	As the script is not necessary for normal use of the web server, it should be removed when the installation of the www server is complete.

Signature	DST “bin/perl” DST “bin/tcsh” DST “bin/tcl” DST “-bin/sh%20” DST “bin/csh”
Severity	3
Implications	Intruders use these commands to find out what shells and interpreters are available on the system. A shell processes commands presented in ASCII form and sends them to the operating system. Under normal circumstances, no HTTP POST or GET should interact with a shell script. If this is a normal operating mechanism, then it is a highly risky method of operation. Otherwise, if this isn’t a normal mechanism of interacting with the web server, then this should be considered a probe/attack.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove any interpretive shells from your web server and continue monitoring for these probes. Attempt to identify the possible intruder via the IP address.

Signature	DST “bin/wedist.cgi”
Severity	5

Implications	Several programs provided within the IRIX Mindshare Outbox Environment subsystem have been found to be insecure. These are the cgi-bin programs webdist.cgi, handler and wrap available for IRIX 5.x and 6.x. Each of these programs can be manipulated to execute arbitrary commands with potentially elevated privileges. The webdist.cgi program allows users to install software over a network via an HTML form, webdist.html, installed in the default document root directories for both the Netsite and Out Box servers. webdist.cgi calls webdist, but due to insufficient checking of the arguments passed, it may be possible to execute arbitrary commands with the privileges of the httpd daemon. When installed, webdist.cgi is accessible by anyone who can connect to the httpd daemon. Because of this, the vulnerability may be exploited by remote users as well as local users. Even if a site’s web server is behind a firewall, it may still be vulnerable.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove this service from your web server immediately.

Signature	DST “bin/wrap/”
Severity	5
Implications	Several programs provided within the IRIX Mindshare Outbox Environment subsystem have been found to be insecure. These are the cgi-bin programs webdist.cgi, handler, and wrap available for IRIX 5.x and 6.x. Each of these programs can be manipulated to execute arbitrary commands with potentially elevated privileges. By exploiting this vulnerability, both local and remote users may be able to execute arbitrary commands with the privileges of the httpd daemon. This may be used to compromise the http server and, under certain configurations, gain privileged access.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove this application from the web server application set.

Signature	DST “bin/handler”
Severity	5
Implications	Several programs provided within the IRIX Mindshare Outbox Environment subsystem have been found to be insecure. These are the cgi-bin programs webdist.cgi, handler ,and wrap available for IRIX 5.x and 6.x. Each of these programs can be manipulated to execute arbitrary commands with potentially elevated privileges. By exploiting this vulnerability, both local and remote users may be able to execute arbitrary commands with the privileges of the httpd daemon. This may be used to compromise the http server and, under certain configurations, gain privileged access.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove this program from the web server or shut down this web server.

Signature	DST “bin/aglimpse/80 | IFS”
Severity	5
Implications	The vulnerabilities exist in the GlimpseHTTP and WebGlimpse packages. Both of these packages provide a web interface that allows you to use Glimpse, an indexing and query system, to provide a search facility for your web site. The cgi-bin programs in these packages perform insufficient argument checking. Due to this, intruders may be able to execute arbitrary commands with the privileges of the httpd process. Glimpse HTTP is an interface to the Glimpse search tool. It is written in PERL.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove the WebGlimpse package from your web server unless the documentation explicitly indicates that this defect is addressed.

Signature	DST “bin/websendmail”
Severity	5
Implications	Intruders use this command to locate the **websendmail** cgi program included in the WebGais package. WebGais is an interface to the Global Area Intelligent Search (GAIS) search tool. It installs a few gateway programs in **/cgi-bin**. It reads a query from a user form and runs the GAIS search engine for that query. Even though the program translates normal www encodings (e.g., **+** to spaces, and **%xx** to real characters), it does not perform any checks for special characters with ambiguous meanings.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove this application if on the web server.

Signature	DST “bin/webgais”
Severity	5
Implications	Intruders use this command to locate the **webgais** cgi program. WebGais is an interface to the Global Area Intelligent Search (GAIS) search tool. It installs a few gateway programs in **/cgi-bin**. The main utility is called “webgais” and is the interface between the search engine and the web. It reads a query from a user form and runs the GAIS search engine for that query. Unfortunately, user input is passed to a system command without making the appropriate checks. The input is enclosed within single quotes, but there is no filtering to remove stray single quotes from the user input.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove or update this application if on the web server.

Signature	DST “bin/php.cgi”
Severity	5

Implications	There have been two attacks connected to PHP, a buffer overflow attack and a failure to check file permissions. The former may lead to arbitrary commands being executed; all operating systems are vulnerable. To determine if a vulnerable PHP is running, connect to the site with any www browser, and type the URL: **http://hostname/cgi-bin/php.cgi**. If the answer is PHP/FI Version 2.0b10, or smaller, the program is vulnerable for the buffer overflow exploit. The latter fault allows any user to view the content of any file readable by the http daemon.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	There are patches for both vulnerabilities, but all users are encouraged to download the latest version of the program from the official site.

Signature	DST “bin/test-cgi”
Severity	5
Implications	Intruders use this command to exploit a vulnerability that allows them to inventory the files on a remote system.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	This script is not necessary for normal operation of the web server. Therefore, it should be removed.

Signature	DST “bin/campas?%0”
Severity	5
Implications	Intruders use this command to exploit the campas cgi script security bug which permits them to execute commands through the campas cgi.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If you do not use the campas cgi script, update it, or delete it from your system.

Signature	DST “filename\”;”
Severity	3
Implications	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
Identifying the Source	
What to Do	

Signature	DST “/uploader.exe”
Severity	3
Implications	The **uploader.exe** cgi program checks and logs access attempts. Intruders use this command to try to access this log to remove their access attempts from the log, destroy the log, and mail the log file to themselves. The log file contains user login information that they can use to gain access to the system.
Identifying the Source	
What to Do	

I

Signature	DST “bin/Count.cgi?”
Severity	3
Implications	Intruders can use this command to spawn an xterm on your $DISPLAY or override the command line. By doing this, intruders can execute commands remotely. The **Count.cgi** program is used to record and display the number of times a web page is accessed. Due to insufficient bounds checking on user-supplied arguments, intruders can overwrite the internal stack space of the **Count.cgi** program while it is executing. They may be able to then execute commands with httpd privileges, which may be used to compromise the http server and gain privileged access.

Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove the exec permissions (chmod -x) from the **count.cgi** executable, located in you httpd's **cgi-bin/** directory.

Signature	DST "../../"
Severity	3
Implications	Intruders use this command to attempt to break out of the current directory and descend into the server's file system. The number of ".." will depend on the location of the file or program they are trying to access.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Verify that your web server is up to date and doesn't allow for this attack to gain access to otherwise secure files.

Signature	DST "//////"
Severity	3
Implications	Intruders use this command to deny access to a particular file for which a user has read access. This type of attack effectively locks files, preventing users from logging in, running programs, and so on, depending on the type of service being denied.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Visit **http://www.ntbugtraq.com/ntfixes.asp** This web page provides a complete list of URLs for patches, including the current security "Hot Fixes" that reflect Denial of Service (DoS) attacks. You can specify your processor, language, and so on, so that the list of URLs reflects your system.

Signature	DST “bin/htmlscript?..”
Severity	3
Implications	Intruders use this command to access system files, which includes any file the web server user can access.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	

Signature	DST “bin/www-sql/”
Severity	3
Implications	Intruders use this command to access a msql database via an http server, and create pages from a query result. This program acts as a filter, using the PATH_TRANSLATED file to access html files on your server. It translates **<!sql...>** tags into html-viewable text, letting other parts of the html file remain unchanged. Unfortunately, www-sql does not verify whether a user can access the intended PATH_TRANSLATED file. For example, let’s say your document tree is **/home/htdocs/** and you have a **/home/htdocs/protected/** directory for which you restrict access using an **.htaccess** file. Enter **http://*your.server*/protected/file.html** in your browser. You are prompted for a username and password. Now, enter **http://*your.server*/cgi-bin/www-sql/protected /file.html**. No prompt is displayed and you get the file you requested.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	This is a common characteristic of other “cgi-wrapper” programs as well, including **w3-msql** and **php.cgi**. The latter addresses this by providing the option to set PATTERN_RESTRICT at compile time. This way, the program will only load files with unusual endings, such as **.phtml**. You can also compile as an Apache module.

I

Signature	DST “bin/pfdisplay.cgi”
Severity	3
Implications	Intruders use this command to locate the **pfdisplay.cgi** program. They can then exploit a vulnerability in **pfdisplay.cgi** that allows any user to view files on the system with user “nobody” privileges.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If you are using SGI cgi scripts, consider limiting access to your domain.

Signature	DST “bin/man.sh”
Severity	3
Implications	
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	

Signature	DST “.nfs/?[Oo]pen\0”
Severity	3
Implications	By appending domcfg**.nsf/?open** to the base URL of a target site, intruders can determine whether that site’s database configuration can be edited by outside users. If, after trying this, the intruder is not prompted for a password, it is likely that those files are both readable and writable. Armed with such access, an intruder could easily redirect the entire web site to any other domain just by filling out a simple form.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.

What to Do	This vulnerability stems from the Access Control Lists (ACLs) used to set the Lotus Domino server's security permissions. Domino's default ACLs allow any web user to have read and write access to the database. There is no way to verify the security of the server configuration databases other than manually verifying the ACLs associated with each database. Visit **http://www.lotus.com** for instructions on how to fix this security flaw in Lotus Domino.

Signature	DST "bin/textcounter.pl"
Severity	3
Implications	Intruders use this command to access the **textcounter.pl** cgi script. They can then exploit a vulnerability in this script that enables everyone to execute commands on the system with the same rights are the httpd daemon.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Remove this script from the web server.

Signature	DST ".asp::$DATA"
Severity	3
Implications	By appending **::$DATA** to **.asp** URLs, intruders can download the active server page (ASP) source code from Microsoft IIS web servers. They can even execute code through the URL.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	A vendor patch is available at **http://www.microsoft.com/security/bulletins/ms98-003/htm**

Signature	"/etc"
Severity	3

Implications	This is more of a general purpose signature that should alert when an attacker tries to grab at some files on the system. This is not a real attack signature per se, however, merely a section of text that will likely be found in other attacks that grab at files on a UNIX system.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If this signature is found, then it could indicate legitimate traffic. Consider the entire signature in context and decide if it could be a legitimate HTTP request or not. If not, then take appropriate actions.

Mail Server (SENDMAIL) Signatures

Sendmail, or the Simple Mail Transfer Protocol (SMTP), is used to send and receive mail messages, and is a standard part of the UNIX operating system. The sendmail program, which is normally only invoked in daemon mode by the root user, can be bypassed to give unprivileged local users the ability to gain root privileges. Cross-Site for Security agents detect attempts to exploit known weaknesses in sendmail, such as mail alias expansion and address verification.

Signature	DST “[Ee][Xx][Pp][Nn]”
Severity	1
Implications	The expn command is a sensitive, sendmail command. A client sends this command to the server in an effort to expand an alias or address out to its canonical form. Also, with certain versions of sendmail, this can be an exploited command to perform a Denial of Service (DoS) attack or to obtain a list of all users on the system. Another issue to consider is that this signature is short enough that when large e-mails are processed that contain MIME enclosures, there is a chance that this signature may be falsely generated. An indication of a false alert is that the signature contains a long string of characters containing the expn signature versus the string being the first set of characters in the signature followed by some white space and an e-mail address.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.

What to Do	Disable the **expn** command by editing the **sendmail.cf** file. Add the following line and restart the sendmail service: `O PrivacyOptions=needmailhelo, noexpn`

Signature	DST “[Vv][Rr][Ff][Yy]”
Severity	1
Implications	This is a sensitive command that has legitimate uses. However, it can also be abused and is an opportunity for an attack. Therefore, any occurrence of this signature should be considered in the true context of the signature. Another issue to consider: this signature is short enough that when large e-mails are processed containing contain MIME enclosures, there if a chance that this signature may be falsely generated. An indication of a false alert id that the signature contains a long string of characters containing the expn signature versus the string being the first set of characters in the signature followed by some white space and an e-mail address.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Consider upgrading to the latest version of sendmail on the system in question. This should assure that any direct attacks using VRFY will fail.

<table>
<tr><td><b>Signature</b></td><td>DST “from[ ]:[ ] |”
DST “from:\”|”
DST “from[ ]:[ ][‘\”]|”</td></tr>
<tr><td><b>Severity</b></td><td>3</td></tr>
<tr><td><b>Implications</b></td><td>Intruders use this command to exploit a PIPE bug.</td></tr>
<tr><td><b>Identifying the Source</b></td><td>The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.</td></tr>
<tr><td><b>What to Do</b></td><td></td></tr>
</table>

Signature	SRC “8.[65432].[98765432]”
Severity	3
Implications	Actually this isn’t an attack signature as much as an alert to the system administrator that an old version of sendmail might be running on the network. This identified version is one in which there are known problems/issues that could indicated that a system is compromisable.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Investigate further based on the IP addresses of the systems involved. A Telnet to the mail port number 25 on the server machine should give you a banner message that indicates what version of sendmail it is running. By either identifying the server or eliminating this as a possible weak sendmail version, one can identify where the signature came from.

Signature	DST “[Hh] [Ee] [Ll] [Pp] [ ] \n”
Severity	2
Implications	Intruders use this command to invoke the SGI help system. They exploit a vulnerability in the help system that enables users to get unauthorized root access if they can log into an account on the system or get physical access to the system console. Another issue to consider is that this signature is short enough that when large e-mails containing MIME enclosures, there is a chance that this signature may be falsely generated. An indication of a false alert is that the signature contains a long string of characters containing the help signature versus the string being the first set of characters in the signature. Also note that any normal e-mail with the string HELP in it could cause a false positive for this signature. Perhaps the best advice is to look for this signature only if the suspect SGI system is somewhere in the connection.
Identifying the Source	The IP address, port ,and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If you see this signature on an SGI system, investigate further.

Signature	DST “[Ww] [Ii] [Zz] [ ] \n”
Severity	4
Implications	
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	This is a signature that has the potential to be falsely generated. If all that is seen is the signature with no surrounding words, then this could be an attack. Most likely, however, this alert will be triggered as part of the contents of an existing e-mail message and can be ignored.

Signature	DST “[Dd] [Ee] [Bb] [Uu] [Gg] [ ] \ n”
Severity	4
Implications	Intruders use the DEBUG command to access and use a more informative debug level, which enables them to look at the actual instructions being executed by the computer. This type of information can then be used to figure out what type of system is running, and, therefore, what areas may be vulnerable.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	This is a signature that has the potential to be falsely generated. If all that is seen is the signature with no surrounding words, then this could be an attack. Most likely, however, this alert will be triggered as part of the contents of an existing e-mail message and can be ignored.

Signature	SRC “irectly to pro”
Severity	3
Implications	This signature represents part of a “cannot mail directly to programs” error message. This signature can be used to identify intruders who are trying to mail a file directly to a program. The program then launches the file, causing any number of problems.

Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature
What to Do	Like many of the other signatures, this signature can be falsely generated; however, if the signature actually spells out the string “cannot mail directly to programs”, then it is highly suspect and should be investigated further.

Signature	SRCDST “bin: [* !] :” SRCDST “0:0:root” SRCDST “sys: [* !] :”
Severity	3
Implications	These signatures indicate that a password file is being mailed. This is a common practice for attackers. Intruders can then use this information to decode a password and obtain some access to the network.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Examine the syslog of the mail relay and see if the e-mail origin and destination were logged somewhere, then take appropriate action.

Signature	SRCDST “& {IFS}”
Severity	4
Implications	Intruders use this command to gain access to euid mail. This exploits part of an rmail IFS vulnerability in AIX 3.2.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Upgrade the sendmail on the AIX system in question or disable e-mail altogether.

DNS Signatures

DNS is an abbreviation for Domain Name System. This system was developed to allow for an hierarchical naming architecture that would allow individuals to refer to machines on the Internet by name rather than by IP address. Most machines that have access to the Internet use the DNS to resolve host names to their address. This system, however common, is capable of being attacked. Some of the more common attack signatures are documented here.

Signature	REQ ZONEXFER
Severity	3
Implications	Since a DNS system must be available at all times, a secondary DNS system is sometimes employed to allow for the primary DNS to come down for maintenance. The secondary DNS periodically polls the primary DNS i to find out if there have been any changes in the zone. If so, it will initiate a zone transfer. This zone transfer is a highly sensitive operation and releases a large quantity of information.
Identifying the Source	The IP address, port ,and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Compare the IP address in the alert to those IP addresses of your known secondary DNS systems. If there isn’t a match then this could likely be a preliminary probe of the domain. Record and monitor which IP address originated this request.

Signature	REPLY SIG “\n” REPLY SIG “\r”
Severity	3
Implications	This signature represents a newline or carriage return in a DNS response. Under normal conditions, these characters are never in a DNS packet. However, if an attacker uses some of the various off-the-shelf attacks, then they would likely contain one of these characters.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.

What to Do	This signature could be falsed by a TXT DNS field; however, this is unlikely. This signature should be taken seriously. Further monitoring and investigation should be taken to identify who or what could have caused this signature.

Signature	REPLY SIG “ in/ ”
Severity	3
Implications	This signature represents a **in/** in a DNS response. Under no circumstance should a “/” appear in the normal encoding of a DNS packet. However, many buffer overflow attacks against a DNS server would likely contain this character.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Again, this is a serious attack that merits further investigation to find if this signature has been falsely generated or is a legitimate attack.

Gopher Signatures

The Gopher protocol is used for distributing documents and performing searches of documents, as defined in RFC 1436. A Gopher client presents the user with a hierarchy of items and directories like an ordinary file system. These documents can be dispersed across a network, and are not limited to one host only.

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3
Implications	These signatures indicate that a password file came over the Gopher connection. This is a highly suspect signature to receive and should be investigated further to determine if there is a reasonable explanation as to why this alert would be generated.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.

What to Do	Shut down whatever Gopher server is up and running. Gopher servers are notorious for having security holes in them. It is highly recommended that alternate solutions to requirements be entertained.

IMAP Signatures

The Internet Message Access Protocol (IMAP) is a protocol that addresses the need of individuals who want to have their mail storage and folders kept in a common location. IMAP is a protocol that has been adopted by many mail client programs such as Netscape, MS Outlook, and Eudora.

Signature	SRCDST ”bin/”
Severity	4
Implications	Various IMAP programs running on servers have had a number of problems with buffer overrun attacks. Typically, the buffer overrun will contain code to execute a command on the remote server machine. A typical signature found in this attack likely executes a program with this string signature contained within it. This is a common signature, however, and numerous falsing could occur for legitimate reasons.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Examine the entire signature. If there appears to be a significant amount of binary data around the signature, then this is may be a buffer overflow attack. However, if the signature appears to be normal text, then this may be a falsely generated attack alert.

Signature	DST ”passwd”
Severity	2
Implications	This signature can also be associated with a buffer-overflow attack. The first thing intruders tend to do is to use the **passwd** command to set up a password for themselves.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.

What to Do	Again, this is a signature that is highly likely to generate some false positives; the context of the entire signature should be considered at all times. If a significant amount of binary data is found preceding the signature, then this could be an attack.

l

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3
Implications	These are signatures that are commonly found in a UNIX password file. A common strategy of attackers is to mail password files to themselves so that the encrypted files can be cracked. Therefore, it is a good idea to search all possible e-mail traffic for this common practice of e-mailing a password file.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If the signature seems to be an actual line from a UNIX password file investigate further. This might require the examination of the contents of the e-mail files and logs directly.

IPOP Signatures

The Post Office Protocol (POP) is an older protocol used for the storage of Internet e-mail. It is primarily used for the temporary storage of e-mail until a personal computer can request the e-mail.

Signature	DST ”bin/”
Severity	4

Implications	Please also refer to the discussion of this same signature in the IMAP section above. A buffer overflow vulnerability has been identified in some POP servers based on QUALCOMM's qpopper. Qpopper is a POP server used for downloading Internet e-mail. Versions of QUALCOMM qpopper prior to 2.5 and some Santa Cruz Operation, Inc. systems are vulnerable.
Identifying the Source	The IP address, port and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If this looks like a buffer-overflow signature with its characteristic binary data in front of the signature, then further investigation is merited. Otherwise, this signature could be part of an e-mail, and has falsely generating this alert. If the signature appears within a sentence or some other context that would be characteristic of an e-mail, then disregard this alert.

Signature	SRC ”-ERR [Pp]ass”
Severity	4
Implications	This signature represents a POP login failure. This typically means that an intruder may be attempting to guess at a password. This is a typical strategy for attackers since the POP program might not use the standard logging mechanisms to report invalid login attempts versus a normal Telnet password guess.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If several (more than five) of these alerts are generated from the same IP address, then investigate further. If fewer than five of these alerts are found, then disregard.

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3

Implications	See the same set of signatures in the previous IMAP protocol.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	If the signature looks like a legitimate line of data from a signature file, then investigate further.

IDENT Signatures

The IDENT protocol was specified in RFC 1413 as a simple protocol that allows a server to verify and identify the owner of a connection from a given client machine. However, many of the IDENT servers are susceptible to buffer overrun errors and can also be a significant source of information leak on a network. The signatures contained in this section are primarily associated with detecting buffer overrun attacks.

Signature	SRCDST ”passwd” SRCDST ”bin/”
Severity	4
Implications	Refer to the discussions in the IMAP section for implications about these signatures.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Refer to the discussions in the IMAP section for implications about these signatures.

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3
Implications	Refer to the discussions in the IMAP section for implications about these signatures.

Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Refer to the IMAP section for implications about these signatures.

NNTP Signatures

The Network News Transport Protocol (NNTP) is the protocol used to post messages to the various message boards on the Internet.

Signature	DST ”group ‘ /bin/sed:”
Severity	4
Implications	This signature represents an intruder issuing a shell command in a new group message.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	

Signature	SRCDST ”bin/”
Severity	4
Implications	Please refer to the IMAP description of this exploit.
Identifying the Source	
What to Do	

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3

Implications	Please refer to the IMAP description of this exploit.
Identifying the Source	
What to Do	

IRC Signatures

Internet Relay Chat (IRC) is a protocol used to support the concept of virtual named rooms and the ability of a user to enter the room and listen to all communication on between members within the room.

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3
Implications	Please refer to the IMAP description of these signatures.
Identifying the Source	
What to Do	

LPR Signatures

The Line Printer (LPR) protocol is associated with streaming data in a flow-controlled manner to a printer. This protocol isn’t typically found on the Internet but is typically found on LANs.

Signature	SRCDST ”bin : [*!] :” SRCDST ”0:0:root” SRCDST ”sys : [*!] :”
Severity	3
Implications	Please refer to these signature explanations found in the IMAP protocol section.
Identifying the Source	
What to Do	

Kerberos Signatures

Kerberos is an authentication protocol that has been integrated into many systems. It is a service that is very enticing to hackers looking for ways to attack or compromise it.

Signature	DST ”krbtgt”
Severity	4
Implications	Kerberos 4 Ticket Granting Tickets (TGTs) are susceptible to dictionary attacks because they contain a constant string that can by used for compares: **krbgt**. It is, therefore, possible to query a Kerberos server, hand in a valid principle (user and kerberos realm), receive a TGT, and decrypt the DES ticket using dictionary words for the key. If the string “krbgt” exists in the decrypted packet, the attacker has the correct Kerberos for password authentication.
Identifying the Source	The IP address, port, and a time stamp will be available with the alert, so some degree of analysis can be conducted to find the source of the signature.
What to Do	Intruders require a user list and Kerberos realm name in order to mount this attack, therefore, ensure that services such as **finger** and **rusers** are turned off to the external users. Intruders can, however, still sniff incoming and outgoing e-mail traffic for usernames. This data leak does not exist in Kerberos5, nor does it appear to be in Kerberos 5 running in Kerberos 4 compatibility mode. Unfortunately, many companies still use Kerberos 4 because of legacy systems and the incompatibility between the two versions. OpenBSD OS distribution has fixed this problem.
	A Related Vulnerability Upon receiving a malformed UDP packet, the Kerberos 4 server returns a packet containing an error string and the principle from some “unsanitized” data structure (e.g., a UDP packet containing a null pointer). Because it is expecting more data than what is being delivered, the pointer to the reused structure references the unpurged principle information. This unsanitized data contains the name of the last user to request a TGT, and the kerberos realm name: the information required to request a TGT for that user and perform a dictionary attack on the response.

WRITESRV Signatures

The **writesrv** daemon allows users to send messages to and receive responses from users on a remote system. Messages are sent with the **write** command. The **writesrv** utility receives incoming requests from a write command and creates a server process to handle the request. This server process communicates with the client process (**write**) and provides whatever services are requested.

To perform this service, the **writesrv** daemon creates a socket that is attached to the port defined in the **/etc/services** file. All requests for service are sent as messages to this socket.

Signature	SRCDST ”bin/”
Severity	3
Implications	Please refer to these signature explanations found in the IMAP protocol section.
Identifying the Source	
What to Do	

Glossary

If you cannot find a term in this glossary, refer to the IBM Dictionary of Computing, which is located at the following URL:

http://www.networking.ibm.com/nsg/nsgmain.htm

It defines technical terms used in the documentation for many IBM products. It also includes IBM product terminology as well as selected terms and definitions from various industry sources.

A

access control
In computer security, the process of ensuring that the resources of a computer system can only be accessed by authorized users in authorized ways. In Cross-Site, access control restricts access to services and channels on a Cross-Site management server.

access control list (ACL)
A list associated with an object that identifies all the subjects that can access the object and their access rights; for example, a list associated with a file that identifies users who can access the file and identifies their access rights to that file. In Cross-Site, there are also ACL entries. These are the individual Cross-Site methods to which roles are assigned.

admin role See *role*.

administrator
A system or web administrator who is authorized to perform management tasks on the Cross-Site console. See also *end user* and *user resource*.

agent resource
A Cross-Site resource that represents the client on which a Cross-Site agent has been installed. When a Cross-Site agent is installed, it registers itself with its Cross-Site management server. A representation of the agent is then created in the management server's repository; this is the agent resource.

agent role See *role*.

alert A message generated when a Cross-Site for Security agent detects an intrusion attempt. An alert contains information about the incident and the severity of the intrusion.

any role See *role*.

application channel
See *channel*.

attack Any attempt by an unauthorized person to compromise the functionality of a networked system. See also *intrusion attempt.*

authentication
Verification of the identity of a user or the user's eligibility to access an object.

authorization
The process of granting a user either complete or restricted access to an object, resource, or function.

availability role
See *role*.

B

BAROC file (.baroc file)
In the event server of the Tivoli Enterprise Console, the internal representation of the defined event classes. Cross-Site provides a BAROC file for event integration with the TEC.

C

CA See *certificate authority*.

Castanet A suite of Java-based applications provided by Marimba Inc. that automatically distribute and maintain software applications and content within a company or across the Internet. Cross-Site uses the

Castanet technology to implement it's Cross-Site for Deployment and auto-update features.

certificate A digital document obtained from a registered certification authority (CA) that contains the identity and public key for a user or system component. In Cross-Site, certificates are used for authenticating and signing channels (using signing certificates), and for securing the Cross-Site management server (using root certificates). See also *foreign certificate*.

certificate authority (CA)
An organization that issues and signs certificates. A CA authenticates the certificate owner's identity, the services that the owner is authorized to use, issues new certificates, renews existing certificates, and revokes certificates belonging to users who are no longer authorized to use them.

certification authority
See *certification authority*.

channel A container that resides on the Cross-Site server and provides restricted access to applications and information. A channel's content, which is one or more content bundles, can be downloaded, installed, and executed by agents. Cross-Site provides two types of channels:

data channel, which enables you to bundle any information with a channel that an end user can then download and copy to their local machine.

application channel, which enables you to bundle Java applications, applets, Windows applications, VisualBasic applications, or any native applications with a channel. The channel then runs on the end users machine, thereby running the contained application. This type of channel is supported through Tivoli Services only.

collaborative management
A cooperative relationship between Internet commerce partners and Internet service providers (ISPs) to ensure the successful completion of business transactions.

collection A Cross-Site resource that groups other resources, thus providing users with a single view of related resources. The Cross-Site console uses folders to represent collections in trees.

console See *Cross-Site console*.

content bundle

A grouping of application or data files that define the content associated with a Deployment channel.

coordinated universal time (UTC)

The time scale, based on the Systeme International (SI) second, as defined and recommended by the Comite Consultatif International de la Radio (CCIR) and maintained (using an atomic clock) by the Bureau International des Poids et Mesures (BIPM). The Systeme International is based on three fundamental units of measure (the meter, the kilogram, and the second) and is sometimes called the “MKS system” because of these units. For most practical purposes, UTC is equivalent to the mean solar time at the prime meridian (0 degrees longitude) of Greenwich, England, which is known as Greenwich mean time.

core services

Services provided by the Cross-Site management server that support all Cross-Site applications. These services include the event service, task manager, policy manager, data manager, channel manager, framework security, and auto-updating.

Cross-Site agent

The client software that can include the Cross-Site for Availability, Cross-Site for Deployment, and Cross-Site for Security components.

Cross-Site console

Cross-Site’s desktop. The console is the user interface where basic operations to manage the Cross-Site environment are performed. The console displays and organizes the resources, policies, tasks, events, and reports of the different Cross-Site applications.

Cross-Site for Availability

A policy-based application that gathers Internet and intranet information. Availability agents monitor and collect data about the performance of actual Internet and TCP/IP transactions from client machines. Availability also scans web sites to determine the integrity of each page.

Cross-Site for Deployment

A policy-based Cross-Site application that enables companies to securely distribute applications to end users over an intranet or the Internet. Cross-Site for Deployment leverages Marimba’s Castanet technology to provide its distribution and update capabilities.

Cross-Site for Security

A policy-based intrusion detection system that relies on the Security agent to monitor networks to detect external and internal attacks.

Cross-Site management repository
A collection of tables and data that support Cross-Site functions in a relational database management system (RDBMS). The management repository is created in your RDBMS during Cross-Site installation.

Cross-Site management server
In the Cross-Site environment, the server hosts all of the resources and provides services to all of the Cross-Site applications. The Cross-Site management server works in conjunction with the web server, database server, and IBM WebSphere Application Server to provide services to the Cross-Site applications and console. Also known as the management server or Cross-Site server.

D

database A collection of interrelated data organized according to a database schema to serve one or more applications.

data channel
See *channel*.

default policy
A set of resource property values that are assigned to a resource when the resource is created. When you create a Cross-Site resource, it is assigned the default policy for its type. See also *policy types*.

demilitarized zone (DMZ)
The area between two firewalls that insulates your internal network from the Internet.

deployment role
See *role*.

DHCP See *Dynamic Host Configuration Protocol*.

digital certificate
See *certificate*.

digital signature
Data that is appended to or cryptographically transformed by a digital certificate. This data, or digital signature, enables the recipient of the data to verify the source and integrity of the data and to recognize potential forgery. In Cross-Site, some channels must be signed with a digital signature from a trusted certification authority.

DMZ See *demilitarized zone*.

DNS See *Domain Name System*.

domain In general, the part of a computer network in which the data processing resources are under common control. In Cross-Site, all elements of Cross-Site that enable the core services and applications to function and that are managed by a Cross-Site management server. This can also be referred to as an administrative domain. See also *foreign domain*.

domain name

On the Web, a domain name is the part of the URL that tells a server using Domain Namfce System (DNS) where to forward a request for a web page. A domain name consists of a sequence of sub-domain names separated by a delineator character. For example, if the fully qualified domain name (FQDN) of a host system is ralvm7.vnet.ibm.com, each of the following is a domain name: ralvm7.vnet.ibm.com, vnet.ibm.com, and ibm.com.

Domain Name System (DNS)

The way that Internet domain names are located and translated into Internet protocol (IP) addresses. Because maintaining a central list of domain-name or IP-address correspondences would be impractical, the lists of domain names and IP addresses are distributed throughout the Internet in a hierarchy of authority.

Dynamic Host Configuration Protocol (DHCP)

A protocol defined by the Internet Engineering Task Force (IETF) that is used for dynamically assigning IP addresses to computers in a network.

E

encryption A method of encoding messages to provide privacy for communications as they move over intranets or the Internet. Some methods of encrypting, such as 128-bit encryption, are so difficult to break that US export laws permit them to be used only within the United States. The Cross-Site suite is available in two levels of encryption: 56-bit encryption, for non-financial institutes outside of the US and Canada, and 128-bit encryption, for all other customers.

end user Any computer user who relies on or requires the services of the Cross-Site agent. In particular, an end user interacts with the agent on a client machine and does not have access to the Cross-Site console or administrative functions. See also *administrator* and *user resource*.

ethernet A 10-Mbps baseband local area network that allows multiple stations to access the transmission medium at will without prior coordination, avoids contention by using carrier sense and deference, and resolves contention by using collision detection and delayed retransmission. See also *token ring*.

event In Cross-Site, an alert, error, or status message generated by Cross-Site server and application operations and displayed on the console. An event informs you when a service or application is functioning and notifies you when policy has been violated.

event adapter
In a Tivoli environment, software that converts events into a format that the Tivoli Enterprise Console can use and forwards the events to the event server.

explorer A panel on the Cross-Site console that enables you to review and set properties for resources, policy, tasks, and events in a view. Explorers separate the elements that are managed. For example, if you are working in the Admin view, the Roles, Certificates, and Permissions explorers are displayed.

extranet A private, virtual network that uses access control and security features to restrict the usage of one or more intranets attached to the Internet to selected subscribers (such as personnel from a sponsoring company and its business partners). See also *intranet*.

F

firewall A gateway device that protects and controls the connection of one network to other networks. The firewall (a) prevents unwanted or unauthorized communication traffic from entering the protected network and (b) allows only selected communication traffic to leave the protected network.

flooding In computer security, “classic” Denial of Service (DOS) attacks that do not directly harm the target system but make it difficult for other systems to use the target system. Flooding is used by hackers to determine where the holes on a network exist. This is done by scanning a network for the systems (IP addresses) and ports that respond to connection requests.

foreign certificate
A certificate that is generated by a certificate authority that Cross-Site does not support or recognize, such as one generated by your company.

foreign certificate authority (CA)

A certificate authority that Cross-Site does not support or recognize.

foreign domain

A domain associated with a Cross-Site management server other than the local management server. (Each management server represents a domain.)

framework services

See *core services*.

fully qualified domain name

In the Internet suite of protocols, the name of a host system that includes all of the subnames of the domain name. An example of a fully qualified domain name is ralvm7.vnet.ibm.com.

G

Greenwich mean time (GMT)

The mean solar time at the prime meridian (0 degrees longitude) of Greenwich, England. Greenwich mean time is sometimes called Z time or Zulu time. Although Greenwich mean time and coordinated universal time are sometimes used interchangeably, they are not synonyms. Greenwich mean time is an approximate time. Because the second is no longer defined in terms of astronomical phenomena, the preferred name for this time scale is coordinated universal time (UTC).

H

heartbeat A signal that the Cross-Site for Security agent sends to the Cross-Site management server to convey that it is still active.

host name See *fully qualified domain name*.

HTML See *Hypertext Markup Language*.

HTTP See *Hypertext Transfer Protocol*.

HTTPS See *Hypertext Transfer Protocol Secure*.

HTTP proxy server

An HTTP server that receives requests intended for another server and that acts on the client’s behalf (as the client’s proxy) to obtain the requested service.

Hypertext Markup Language (HTML)
A markup language that is specified by an SGML document type definition (DTD) and is understood by all web servers.

Hypertext Transfer Protocol (HTTP)
The protocol that is used to transfer and display hypertext documents. HTTP is the standard web browser protocol.

Hypertext Transfer Protocol Secure (HTTPS)
The standard secure web protocol. This protocol uses Secure Sockets Layer (SSL) encryption.

I

IBM WebSphere Application Server (IBMWebAS or WebSphere)
IBM’s Java servlet-based application environment for building, displaying, and managing web applications. In Cross-Site, WebSphere is installed behind the web server.

incident A signature detected by the Cross-Site for Security agent that indicates questionable network activity.

install role See *role*.

Internet service provider (ISP)
An organization that provides access to the Internet.

intranet A private network that integrates Internet standards and applications (such as web browsers) with an organization’s existing computer networking infrastructure. See also *extranet*.

intrusion attempt
An attempt by an unauthorized person to access or destroy a network resource.

Internet protocol (IP) address
The unique 32-bit address that specifies the location of each device or workstation on the Internet. For example, 9.67.97.103 is an IP address.

ISP See *Internet service provider*.

J

Java Database Connectivity (JDBC)
An application programming interface (API) that has the same characteristics as Open Database Connectivity (ODBC) but is

specifically designed for use by Java database applications. Also, for databases that do not have a JDBC driver, JDBC includes a JDBC to ODBC bridge, which is a mechanism for converting JDBC to ODBC; it presents the JDBC API to Java database applications and converts this to ODBC. JDBC was developed by Sun Microsystems, Inc. and various partners and vendors.

Java Development Kit (JDK)
A software package that can be used to write compile, debug, and run Java applets and applications.

Java Runtime Environment (JRE)
A subset of the Java Development Kit (JDK) that contains the core executables and files that constitute the standard Java platform. The JRE includes the Java Virtual Machine (JVM), core classes, and supporting files.

Java servlet The server-side analog of a Java applet. Servlets are dynamically loaded in response to HTTP requests. Their effects might range from simple HTML generation to complex DBMS queries. In Cross-Site, all management server facilities are implemented through Java servlets.

Java Virtual Machine (JVM)
A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications).

JDBC See *Java Database Connectivity*.

JDK See *Java Development Kit*.

JRE See *Java Runtime Environment*.

JVM See *Java Virtual Machine*.

K

KeyRing A special control file that governs access to secure web servers and Cross-Site management servers. The KeyRing contains the public keys, private keys, trusted roots, certificates, names of sites, and CAs that are trusted and, therefore, authorized for access.

key file password
The password for the web server’s key file that secures a server. This password is chosen when a certificate is purchased from a CA and is needed when you start the (secure) web server.

key ring See *KeyRing*.

M

management repository
See *Cross-Site management repository*.

management server
See *Cross-Site management server*.

mgmtserver role
See *role*.

N

NetView See *Tivoli NetView*.

Network File System (NFS)
A protocol developed by Sun Microsystems Incorporated, that allows any host in a network to mount another host's file directories. Once mounted, the file directory appears to reside on the local host.

O

Open Systems Interconnection (OSI)
A standard architecture or model for how messages should be transmitted between any two points in a telecommunication network. The reference model defines seven layers of functions that take place at communication endpoints. This standard is a guide for product developers that ensures that their products will consistently work with other products.

P

PDF See *portable document format*.

performance monitoring
The Cross-Site for Availability function that monitors TCP/IP services on a client system and logs data to a local file. Availability agents upload this data to the Cross-Site management server.

policy A set of rules applied to Cross-Site resources that control the behavior of those resources.

policy types The predefined set of policies from which you can create policy for a Cross-Site resource. Policy can be created in Cross-Site using the following pre-defined policy types:

Intrusion Detection policy is a Cross-Site for Security policy that assigns Security-related rules to an agent resource to define security and intrusion detection behavior.

Availability Agent policy is a Cross-Site for Availability policy that assigns monitoring rules to an agent.

Site Scan policy is a Cross-Site for Availability policy that controls the behavior of the scan that collects information about a particular web site.

Channel policy is a Cross-Site for Deployment policy that controls how often each agent checks the channel for updates; the interval between each check for channel updates; and the time period during which checks for channel updates should occur.

Portable Document Format (PDF)

A standard specified by Adobe Systems, Incorporated, for the electronic distribution of documents. PDF files are compact; can be distributed globally via e-mail, the Web, intranets, or CD-ROM; and can be viewed with the Acrobat Reader, which is software from Adobe Systems that can be downloaded at no cost from the Adobe Systems home page, which is located at the following URL:

http://www.adobe.com

PostScript (PS)

A standard specified by Adobe Systems, Incorporated, that defines how text and graphics are presented on printers and display devices.

proxy server

A server that receives requests intended for another server and acts on the client’s behalf to access the requested server. A proxy server is often used when a client and the server are incompatible for direct connection. For example, when the client is unable to meet the security authentication requirements of a server but should be permitted some service, a proxy with the appropriate access is used. See also *HTTP proxy server* and *SOCKS server*.

principal A Cross-Site entity to which security roles can be assigned in order to grant permissions.

PS See *PostScript*.

publish The process of making a channel available on a Cross-Site server, enabling the channel to be downloaded by subscribed agents.

R

RDBMS See relational database management system.

realm A name used by a browser in correlation with a URL to save the password information you enter so that it can authenticate automatically on the next challenge.

relational database management system (RDBMS)
A collection of hardware and software that organizes and provides access to a relational database.

remote procedure call (RPC)
1) A facility that a client uses to request the execution of a procedure call from a server. This facility includes a library of procedures and an external data representation.
(2) A client request to a service provider located in another node.

resource In Cross-Site, any entity that can have policy applied to it and, therefore, can be managed. A resource is represented in the management repository by a database object and, on the Cross-Site console, by an icon in the explorer's tree. Resources include agents, sites, channels, collection, and users.

role A set of actions and responsibilities associated with a particular activity. In Cross-Site, roles are used to grant privileges to principals and control access to services and Deployment channels. The following roles are provided by Cross-Site:

admin, which is for administrators who need full access to all services and channels. Assign this role to a person who is responsible for making organizational and content changes to the management server. This role is equivalent to root on UNIX, or Administrator on NT. agent, which is assigned only to Cross-Site agents for Cross-Site operations.

any, which is only used in ACL entries. It provides read-only access to non-critical information for any authenticated principal.

availability, which is for end users of Cross-Site agents who must install and use the agent component of the Cross-Site for Availability application.

deployment, which is for end users of Cross-Site agents who must install and use the agent component of Cross-Site for Deployment.

install, which is for users who are responsible for installing Cross-Site for Security agents.

mgmtserver, which is automatically assigned to Cross-Site servers only; this role is for inter-domain operations. This role is not for users.

securityAgent, which is automatically assigned to all Cross-Site for Security agents. This role is not for users.

user, which is for end users who run agents an need read-only access to information an executable files on the Cross-Site server. This role should be assigned to end users who do not administer the product.

RPC See *remote procedure call*.

S

Secure Sockets Layer (SSL)
A secure protocol that allows data to be encrypted and enables clients to authenticate a server in client-server communication.

securityAgent role
See *role*.

Security agent
See *Cross-Site agent*.

servlet See *Java servlet*.

signature In Cross-Site for Security, a sequence of IP packets that characterizes a security threat.

signing certificate
See *certificate.*

signing password
In general, the password required to use a digital certificate. In Cross-Site, the password that must be entered by the publisher of the channel to sign a channel with the particular signing certificate.

Simple Mail Transfer Protocol (SMTP)
In the Internet suite of protocols, an application protocol for transferring mail among users in the Internet environment. SMTP specifies the mail exchange sequences and message format. It assumes that the Transmission Control Protocol (TCP) is the underlying protocol.

Simple Network Management Protocol (SNMP)
In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an

application layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

site resource
A Cross-Site resource that represents a web site that you wish to monitor. You create a site resource to designate it as the target of an Availability task, or site scan.

site scan An Availability task that is used to "crawl" web sites, scanning them for availability information such as the presence of broken links.

SMIT See *System Management Interface Tool.*

SMTP See *Simple Mail Transfer Protocol*.

SNMP See *Simple Network Management Protocol*.

SOCKS server
A circuit-level gateway that provides a secure one-way connection through a firewall to server applications in network that is not secure.

SSL See *Secure Sockets Layer.*

subscribe An action that enables an interested end user to use his or her Deployment agent to download, install, launch and view a channel's files onto your local drive.

System Management Interface Tool (SMIT)
An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

T

tablespace In relational database management systems, an abstraction of a collection of containers in which database objects are stored.

task An action that can be performed on one or more resources in a Cross-Site domain. A task defines the application or class to be invoked when the task is executed. It also defines the schedule under which the task will execute.

TCP/IP See *Transmission Control Protocol/Internet Protocol*.

TEC See *Tivoli Enterprise Console*.

timeout A time interval allotted for certain operations to complete; for example, the period of time allotted for a response before a system operation is interrupted and must be restarted.

time stamp The value on an object that is an indication of the system time at some critical point in the history of the object.

Tivoli Enterprise Console (TEC)
A Tivoli Enterprise product that collects, processes, and automatically initiates corrective actions for system, application, network, and database events; it is the central control point for events from all sources. The Tivoli Enterprise Console provides a centralized, global view of the network computing environment; it uses distributed event monitors to collect information, a central event server to process information, and distributed event consoles to present information to system administrators.

Tivoli NetView
A Tivoli Enterprise product that enables distributed network management across multiple operating systems and protocols.

Tivoli Software Distribution
A Tivoli Enterprise product that automates software distribution to clients and servers in a network computing environment. An organization can use this product to install and update applications and software in a coordinated, consistent manner across a network. Tivoli Software Distribution creates file packages and distributes them to predefined subscribers.

token ring (1) According to IEEE 802.5, network technology that controls media access by passing a token (special packet or frame) between media-attached stations.
(2) A FDDI or IEEE 802.5 network with a ring topology that passes tokens from one attaching ring station (node) to another.
See also *token ring*.

Transmission Control Protocol/Internet Protocol (TCP/IP)
(1) The Transmission Control Protocol and the Internet Protocol, which together provide reliable end-to-end connections between applications over interconnected networks of different types.
(2) The suite of transport and application protocols that run over the Internet Protocol.

trusted root In the Secure Sockets Layer (SSL), the public key and associated distinguished name of a certificate authority (CA).

U

user resource
A Cross-Site principal that provides unique access to different types of Cross-Site administrators and end users. You, acting as a Cross-Site administrator, create a user resource that the agent’s installation program can use to access the Cross-Site server.

user role See *role*.

UTC See *coordinated universal time*.

V

view On the Cross-Site console, a collection of panels that enables you to view and manipulate Cross-Site objects. Cross-Site offers six views: Global, Availability, Security, Deployment, Admin, and Help. The views that are available to you depend on the particular applications that you have installed.

W

WebSphere See *IBM WebSphere Application Server*.

Index