

Quest® InTrust 11.3

User Guide



Copyright 2017 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, and the Quest logo are trademarks and registered trademarks of Quest Software Inc. in the U.S.A. and other countries. For a complete list of Quest Software trademarks, please visit our website at <https://www.quest.com/legal>. All other trademarks, servicemarks, registered trademarks, and registered servicemarks are the property of their respective owners.

Legend

-  **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.
-  **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.
-  **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

InTrust User Guide

Updated - February 2017

Version - 11.3

Contents

First Steps	6
Installing InTrust	6
Participation in the Quest Software Improvement Program	7
Collecting Events in Real Time	7
Introduction to Repositories	8
Common Tasks	8
Managing Collections	8
Analyzing Collections	9
Managing Repositories	10
Setting Up Daily Cleanup	10
Gathering Windows Logs Other than Security, Application and System	10
Applications and Services Logs	10
Forwarded Events	11
Load Balancing	11
Example: Configuring Logon and User Session Auditing	12
Start Gathering	12
Put Auditing to a Test	12
View the Results in Repository Viewer	13
Setting Up Gathering of Syslog Data	14
Common Tasks for Syslog Collections	14
Passing Messages On	15
Analyzing Syslog Collections	15
Message Parsing Specifics	15
RFC 3164 Specifics	16
Treatment of Timestamps (RFC 3164)	16
RFC 5424 Specifics	16
Treatment of Timestamps (RFC 5424)	17
Mapping of Event Fields	18
Treatment of Facility and Severity Information	19
Searching for Events in Repository Viewer	21
Where to Run Repository Viewer	21
Working with Repository Viewer and InTrust Server	21
Working with Repository Viewer without an InTrust Server	24
Getting Started with Repository Viewer	24
Running Searches	26
Predefined Searches	26
Changes to Event Fields	27
Custom Searches	29
Managing Repository Groups	31

Filter Parameters in Repository Viewer	31
Configuring Parameters	32
Custom Logic for Parameters	33
Normalized Event Fields (Who, What, When and Others)	33
Advanced Expression-Based Filters	34
Changing the Business Hours and Non-Business Hours Parameters	35
Examining Event Details	36
Configuring the Result Layout	36
Organizing the Grid	36
Grouping	37
Sorting	37
Hiding and Unhiding Events	37
Saving the Results	38
Using Pie Charts and Column Graphs	38
Case Study: Forensic Analysis of Active Directory Tampering	38
Reporting on Events Using Repository Viewer	40
Tracking Report Progress and Running Manually	41
Reconfiguring and Disabling Reports	42
Integration into SIEM Solutions Through Syslog Forwarding	43
Turning Forwarding On and Off	43
Filtering Specifics	44
Data Conversion Formats	44
Basic Event Forwarding Scenario	45
Advanced Event Forwarding Scenario	45
Example: Set Up Forwarding to SecureWorks	45
Make Sure You Have the Data	46
Configure the Forwarding	46
Example: Set Up Forwarding to Splunk	46
Get Splunk Ready	46
Step 1: Define a Source Type	46
Step 2: Configure a Network Input	47
Step 3 (Conditional): Restart Splunk	47
Configure the Forwarding	47
Further Reading	49
About us	50
Contacting Quest	50
Technical support resources	50
Third-Party Contributions	51
Licenses	52
Apache 1.1	52
flex 2.5.25/27	53

GPL (GNU General Public License) 2.0	54
LGPL (GNU Lesser General Public License) 2.1	60
Net-SNMP	70
OpenSSL 1.0	75
Netscape Public License Version 1.1	78
Stanford SRP	84
zlib 1.2.3	86

First Steps

InTrust is an event-log management solution that provides for collection, correlation, archival, and reporting on the heterogeneous audit data from your enterprise-wide network. InTrust real-time alerting and notification capabilities allow you to stay aware of what is going on in your network and how your business-critical resources are functioning.

Although InTrust is a powerful and comprehensive framework for audit data, deployments can range widely in complexity. The following types of coverage are all possible:

- Basic everyday security auditing with a minimal set of components
- Archival of audit data in compressed repositories for regulations compliance
- Fast search and reporting tools that work with repository data
- Real-time monitoring for critical security events, with alert tracking and automated response actions
- Auditing of multiple platforms and custom logs with advanced reporting based on SQL Server Reporting Services
- Combinations of the above

This guide explains only the use of the basic InTrust deployment. More sophisticated features and workflows are described elsewhere in the InTrust documentation set—for example, in the [Deployment Guide](#).

Installing InTrust

Before you begin installation, confirm that the system requirements are met (see [System Requirements](#)). Also note that the InTrust installer verifies this automatically.

To begin installation, use the Autorun application that comes with your InTrust distribution; click **InTrust Default Suite** on the **Install** tab to begin setup.

i | **NOTE:** If you need custom InTrust capabilities, consider the **InTrust Extended Suite** option, which is not covered in this set of topics. For details, see the [Deployment Guide](#).

Next, complete the remaining steps.

! | **CAUTION:** The default InTrust components require that ports 900, 8340 and 8341 be open for inbound traffic. The InTrust installer knows how to configure these ports automatically in Windows Firewall. However, if you use a hardware or third-party software firewall, make sure these ports are open.

Participation in the Quest Software Improvement Program

One of the setup steps prompts you to select the country where you are performing InTrust installation. This choice affects whether your participation in the Quest Software Improvement Program is enabled automatically.

The Software Improvement Program involves Quest receiving anonymous usage statistics from the Quest software you install. No personal identifying data (such as account names) is included in this feedback. The purpose is to determine which features are most popular and find out how their use can be streamlined.

The following information is transmitted:

- Hardware configuration
- Which product features are used
- External IP addresses

Participation is voluntary. Although it is enabled automatically for some countries, you can change your choice at any time after InTrust setup is complete; for details, see the [Installing the First Server in InTrust Organization](#) topic in the [InTrust Deployment Guide](#).

Collecting Events in Real Time

After you have installed the default components, run the InTrust Deployment Manager console by clicking its entry in the Start menu. This console manages gathering of audit data to InTrust repositories.

In the console, you need to specify the computers you want to audit and specify what kinds of events you need. This is done by setting up collections. Collection settings include the computers to collect from, data sources (definitions of the types of events) and the repository to collect to. Simply put, the point of a collection is to “get this kind of data from these computers to this repository”.

For gathering to work, computers in collections need to have InTrust agents installed. You can install agents on specific computers by selecting them in the right pane while a collection is highlighted and clicking **Install Agents**. Alternatively, enable the **Install agents automatically** option while you are creating or editing a collection to automatically install them on all computers in the collection. If this option is off in a newly-created collection, no gathering occurs. Once you enable it, agents are installed and gathering begins.

! CAUTION: If the Install agents automatically option is enabled for a collection, InTrust will try to keep the agents on all computers in the collection. If you uninstall an agent from a computer in such a collection, it will be reinstalled automatically.

In this situation, to stop gathering from a computer, you need to remove it from the collection.

If the Install agents automatically option is disabled, you need to install and uninstall agents manually using toolbar commands.

When you run InTrust Deployment Manager, you are directed to the home view, where you are briefly introduced to the basics of real-time event collection workflow. This view explains collections (how InTrust organizes computers to collect from) and repositories (stores to collect data to), and it provides quick action links to help you get work done.

If you are starting InTrust Deployment Manager for the first time, take the opportunity to create a collection in the home view: either a Windows collection for gathering from Windows computers or a Syslog collection for capturing Syslog messages from devices and hosts.

Introduction to Repositories

An InTrust repository is a store for audit data collected by InTrust. Its architecture is such that massive amounts of data can be stored efficiently in a compact way and indexed for fast browsing in InTrust Repository Viewer and streamlined access by IT Security Search.

This helps achieve security regulations compliance and provides a ready-made toolset for event analysis. For an in-depth description of InTrust repositories, see the [Understanding InTrust Repositories](#) topic.

For the purposes of this guide, however, it is sufficient to know the following about repositories:

- When you set up InTrust, a default repository is automatically created for you in the InTrust installation folder (by default, installation to Program Files is suggested). Note that the default repository is not recommended for real production use, but only for evaluation and training. When you are confident with the InTrust workflow, create your own repository on a server that has ample disk space and is ready for intensive disk writes.
- You can use the default repository for all your logon and user session auditing needs (unless further scaling is required).
- The folder where you create a repository should be available over the network.
- If necessary, you can have multiple repositories, specialized by the type of data they are supposed to contain, by their location, or by some other characteristic. However, try to keep a manageable number of repositories.
- The toolset described in this document works only with indexed repositories and provides no way to disable repository indexing. Although disabling it is possible in the InTrust Manager component, you should not do it for repositories where indexing is ever going to be enabled again, because indexing will not be able to catch up unless this is a very recently created repository.

To manage repositories, use the **Storage** view in InTrust Deployment Manager.

Common Tasks

The following topics describe how you can manage and adapt InTrust using the InTrust Deployment Manager console.

- [Managing Collections](#)
- [Analyzing Collections](#)
- [Managing Repositories](#)
- [Gathering Windows Logs Other than Security, Application and System](#)
- [Load Balancing](#)

Managing Collections

You can add, delete and edit collections at any time. To work with collections, go to the **Collections** view of InTrust Deployment Manager.

To create a collection, right-click the **Collections** node and select **New Windows Collection** or **New Syslog Collection**. To edit or delete a collection, right-click it and use the corresponding command.

To add computers to a collection

Use any of the following methods:

- In the wizard that opens when you edit a collection, change the computer list on the Specify Computers step.
- Select the computers you need in the **Computers not in a collection** search folder in the navigation pane and click **Add to Collection** (in the toolbar or in the shortcut menu), and then select the collection you need.
- Supply a computer list in a plain-text file. For that, in the wizard that opens when you edit a collection, click **Add Computer** and type the local path to the file containing computer names or IP addresses of computers you want to collect from. Note that this is only a one-off import action. InTrust does not track changes to the file or remember its location.

To delete computers from a collection

1. Right-click the collection and select **Edit Collection**.
2. In the wizard that opens, go to the Specify Computers step.
3. In the list of computers, select the computers you do not need, and click **Remove** (in the toolbar or in the shortcut menu).

To stop gathering from a computer without removing it from a collection

This works only in collections where the **Install agents automatically** option is disabled. In such collections, use the **Install agent** and **Uninstall agent** commands (in the toolbar or in the shortcut menu) to manage gathering without affecting collection membership.

In addition, the following management actions can be done in the wizard:

- Change the account used for connecting to the computers in the collection
Set the credentials on the Specify Computers step.
- Change the list of logs that are gathered
Select the data sources you need on the Data Sources and Repository step.

i **NOTE:** Some of the computers in the collection may not have the logs that the data sources expect. If you do not want InTrust to treat such situations as errors, select the **Suppress errors from non-existent data sources** option on the Data Sources and Repository step. This will make sure that the auditing status of an agent will not be affected if a specified log is not found.

- Change the repository that events are gathered to
Select the repository on the Data Sources and Repository step.

Analyzing Collections

When a collection is selected, the right pane shows a table with information about the collection members. The table supports multi-level grouping of collection computers, so that you can organize the computers in tree-like views using any criteria. For example, you can group computers by status, then by domain, then by type.

To use multi-level grouping, drag table column names from the computer list to the area above the list. The computer list changes accordingly.

i **NOTE:** The difference between the “Not Installed” and “Failed” computer statuses is as follows:

- “Not Installed” means agent installation has never been tried for this computer.
- “Failed” means agent installation has been tried but failed

To hide the computers you are not interested in, you can use view filtering. To configure a view filter, use the controls underneath the table column names: click the operator icon to select the operator, and specify the value to filter by.

The same grouping and view filtering techniques are available in the views with search folder results.

Managing Repositories

You can add, delete and edit repositories at any time. To work with repositories, go to the **Storage** view of InTrust Deployment Manager.

In this view, the left-hand pane lists the available repositories, and the right-hand pane shows the properties of the selected repository.

To create and delete repositories, use the **New** and **Delete** buttons. To edit the properties of a repository, select it and click the **Edit** link for the group of settings you want.

i **IMPORTANT:** The defining property of a repository is the path to the network share that contains the collected data. When you specify the path, use a UNC name. This makes the repository available to client applications in the network, such as Repository Viewer and IT Security Search. It will also make it easier to integrate the repository into an extended InTrust deployment if you decide to perform it.

You can also create a repository when you create a new collection or edit an existing collection (see [Managing Collections](#)), on the Data Sources and Repository step of the wizard.

Setting Up Daily Cleanup

You can configure a repository to keep only recent data and automatically discard data that is too old. For that, edit the **Daily Cleanup** settings in the repository properties in the **Storage** view. Specify how old data can get before it is considered too old and at what time daily cleanup should start.

Gathering Windows Logs Other than Security, Application and System

Applications and Services Logs

To gather a third-party Windows event log that is available in the Applications and Services Logs subtree in Windows Event Viewer, you need to create a data source for it. This is done in the wizard used for creating and editing collections, on the Data Sources and Repository step.

Proceed to that step, and then do the following:

1. Click **Add**. The New Data Source dialog box opens.
2. Specify a meaningful name for the new data source. Optionally, provide a description.
3. In the text box below, specify the exact log name.

i **NOTE:** If you don't know the name, look it up in Event Viewer, as follows:

1. Run Event Viewer on a computer where the log is available, and locate the log you need.
2. Open the properties of the log. The name is in the **Full Name** text box.
4. Click **OK** to save the new data source, and select the check box next to it in the data source list.
5. Complete the wizard.

Forwarded Events

One of the available Windows log types is Forwarded Events. If subscription-based logging of these events is enabled, InTrust can collect them just like other events. It is possible to configure the gathering using the procedure above; the exact log name in step 3 is **ForwardedEvents** in this case.

However, due to the limitations of this forwarding technology, data in the forwarded events is mostly meaningless. You can gather it to a repository, but you cannot search in it or build reports on it. Therefore, collecting this data is not recommended. Instead, use InTrust to gather the original events from the sender computers.

Load Balancing

The metrics and suggestions in this section are based on tests performed by quality control.

InTrust agents send events to InTrust servers in batches. By default, the event submission rates are as follows:

- On Windows servers, including domain controllers, a batch file is sent every minute.
- On workstations, a batch file is sent every seven minutes.

There are two primary limits to consider when estimating if an InTrust server can cope with its load. On the one hand, an InTrust server can gather from no more than 10,000 computers (servers or workstations) at a time. On the other hand, an InTrust server should not receive more than 60,000 events per second in a steady stream. The rate of events from a computer depends very much on the number of data sources that are processed on that computer.

For example, a collection of about 3000 computers with 5 data sources each, 4 events per second per data source, produces a combined stream of 60,000 events per second. This is a load that a 16-core InTrust server with SSD storage and 16GB of memory should handle without problems.

Tips on avoiding excessive workload on a server:

- Keep track of how many computers there are per InTrust server.
- Add InTrust servers if necessary.
- Assign different servers to different collections.
- Distribute the computers among your collections evenly.

! **CAUTION:** When adding an InTrust server to your existing organization, you should run InTrust setup under an account that can manage the InTrust configuration. The account used for installing the first InTrust server automatically has these privileges. To add InTrust organization administrators, in InTrust Deployment Manager click Manage | Configure Access. Of course, to add organization administrators, you must be an organization administrator yourself.

Example: Configuring Logon and User Session Auditing

InTrust lets you gather two types of data related to users logging on and off computers:

1. **Native Windows Security log events**
These events provide basic logon and logoff information, but contain no indication of the user's presence in the system at any particular time. They only capture the act of logging on and off, and their reliability is limited.
2. **User session events enabled by InTrust**
These advanced events contain enough information to help you track not only logons and logoffs, but also when users are actively using computers. For example, they indicate the exact times and durations of terminal sessions connected to domain controllers. User session events are logged locally on computers that have the InTrust agent installed and are in collections where the "InTrust User Session Tracking" data source is enabled.

For logon and user session tracking to be complete, make sure both the "Windows Security Log" and "InTrust User Session Tracking" data sources are enabled in your collections. For details about enabling data sources, see [Managing Collections](#).

Start Gathering

For the purposes of this topic, configure logon event gathering only from domain controllers. Take the following steps:

1. Right-click **Collections** and select **New Collection**.
2. On the General Properties step, give the collection a name indicating that it contains domain controllers.
3. Proceed to the Specify Computers step of the wizard, and add your domain controllers to the list. Make sure the **Install agents automatically** option is selected.
4. On the Data Sources and Repository step, make sure the "Windows Security Log" and "InTrust User Session Tracking" data sources are enabled.
5. Complete the steps.

After this, agents are installed on the domain controllers, and gathering starts automatically.

If you want to watch other computers in addition to or instead of domain controllers (for example, Exchange or file servers), create a new collection and add all the computers you need to it. Configure the gathering options for this collection likewise.

Put Auditing to a Test

To confirm that auditing is working as intended, deliberately perform some of the activity you are watching for on the computers you are watching. Do any of the following:

- Log on to the computers included in the collection and log off
- Lock and unlock the computers
- Set a low screensaver timeout to cause the screensaver to start
- Switch the user

Next, check that your actions have been captured in the repository.

View the Results in Repository Viewer

The InTrust Repository Viewer application lets you explore and analyze the contents of InTrust repositories. To browse the repository you have been collecting to, run Repository Viewer from the Start menu, and click **File | Open Repository**.

In the dialog box that opens, select the **Production repository** option, and proceed to specify the repository you have been working with.

i | **NOTE:** A *production repository* is a repository that is available in InTrust Deployment Manager or InTrust Manager. For details about production and idle repositories, see [Repository Connections](#).

The left pane of the Repository Viewer console shows:

- A navigation tree that organizes events by domain and log type
- A collection of predefined search folders with preconfigured popular filters for quick event analysis

You can select any of the search folder nodes or any of the repository hierarchy nodes, and view the events they contain by clicking the **Go** button. For the purposes of this document, the following predefined searches are useful:

- Searches in the **Logons** subfolders of the topmost search folders
- Searches in the **User sessions** subfolders of the topmost search folders

Select one of these searches and click **Go**. If events about your activity are displayed in the right pane, then auditing has been set up correctly.

For detailed Repository Viewer documentation, see [Searching for Events in Repository Viewer](#).

Setting Up Gathering of Syslog Data

You can use the InTrust Deployment Manager console to collect and manage Syslog data that is received by InTrust Server. To enable Syslog data capture, you need to set up a Syslog collection, as follows:

- Specify the InTrust server that should listen for Syslog messages
- Specify the devices you want to audit
- Specify the repository where you want to store the collected Syslog data

You can add, delete and edit collections at any time.

The first time you run InTrust Deployment Manager, you are directed to the welcome page, where you are prompted to create a collection. Take the opportunity to create your Syslog collection.

You can create more collections at any time. For that, right-click **Collections** and select **New Syslog Collection**, and then follow the wizard steps.

Common Tasks for Syslog Collections

To add a Syslog collection

1. In the InTrust Deployment Manager console, go to the **Collections** view.
2. Right-click **Collections** and select **New Syslog Collection**.
3. In the **New Syslog Collection** wizard, specify a name and a description for the collection.
4. On the **Set Up Collection** step, specify the InTrust server from which you want to get Syslog audit data and repository. You can collect Syslog data from all devices that send Syslog messages to the InTrust server or specify certain devices by selecting one of the following options:
 - a. **All Syslog data received by InTrust server**
 - b. **Syslog data only from devices you specify on the next step**
5. If you select the **Syslog data only from devices you specify on the next step** option, add the devices you want on the next **Specify Syslog Devices** step. For that click the **Add** button and select **Devices**. In the **Specify Syslog Devices** dialog box, you can add devices from the list or specify the IP address (DNS name) of the certain device.
Also you can upload a text file that contain a list of device IPs, for that click **Add** and select the **Import from file** option. A list file uses the plain text format. Each IP address must be a separate line in the file.

To add devices to a collection

Use any of the following methods:

- In the wizard that opens when you edit a Syslog collection, change the devices list on the **Specify Syslog Devices** step as described in the previous procedure.
- Select the devices you need in the **Syslog devices not in a collection** search folder in the navigation pane and click **Add to collection**, and then select the collection you need.

! **CAUTION:** You cannot add a device from the **Syslog devices not in a collection** search folder to a collection if this collection and this device are related to different InTrust servers.

To delete Syslog devices from a collection

1. Right-click the Syslog collection and select **Edit Collection**.
2. In the wizard that opens, go to the **Specify Syslog Devices** step.
3. In the list of devices, select the devices you do not need, and click **Remove**.

To start a new repository

You can create a repository when you create a new Syslog collection or edit an existing collection, on the **Set Up Collection** step of the wizard. For finer-grained management of repositories, use the Storage view (for details, see [Managing Repositories](#)).

Passing Messages On

If both Syslog listening and forwarding are enabled for a repository at once, then incoming Syslog messages are forwarded unchanged. This happens independently of writing the messages to the repository.

Analyzing Syslog Collections

When a Syslog collection is selected, the right pane shows a table with information about the collection members. The table supports multi-level grouping of collection computers, so that you can organize the computers in tree-like views using any criteria. For example, you can group computers by source status, then by collection, then by timestamp.

The device can have **Not Collecting** or **Collecting** status. If the InTrust server does not receive events from the device for half a week - the device changes the status to **Not Collecting**.

The **Timestamp** field contains the time when the last syslog message was generated by a device or time when the last message was received on the InTrust server (If impossible to determine the time when the message was generated). All syslog devices that are located in the **Syslog devices not in a collection** search folder contains the time when the last event was received on the InTrust server in the **Received** field.

To use multi-level grouping, drag table column names from the devices list to the area above the list. The devices list changes accordingly.

Message Parsing Specifics

InTrust parses the Syslog messages it captures to store a useful representation of them in the repository. Only UDP v4 is used for receiving messages, and they can use either ASCII or UTF-8.

Messages are expected to conform to either [RFC 3164](#) or [RFC 5424](#). The fields of an event entry in the repository are filled in from the fields of a Syslog message.

A message is parsed until the end or until a mismatch occurs. The parser breaks down the message into as many insertion strings as it can. No matter how many fields InTrust is able to parse successfully—all of them, just the first three or none at all—the entire message text is saved in the Description field. This enables you to find the message in Repository Viewer (by using the **Any field** parameter) or IT Security Search even if the fields are not mapped properly.

RFC 3164 Specifics

The following pattern is defined in RFC 3164:

```
<PRI>TIMESTAMP HOSTNAME TAG: MSG
```

An example of a valid message is as follows:

```
<34>Oct 14 22:14:15 mymachine su: 'su root' failed for lonvick on /dev/pts/8
```

The **PRI** field indicates the facility and severity. For details, see [Treatment of Facility and Severity Information](#).

A message has the following parts:

Field	Details
PRI	Indicates the facility and severity. For details, see Treatment of Facility and Severity Information .
TIMESTAMP	See Treatment of Timestamps (RFC 5424) .
HOSTNAME	The name of the host as returned by the hostname command. If it is unknown, the host puts its own IP address in this field.
TAG	This is a piece of data that can help classify the message. It is often followed by the process ID in square brackets. If the process ID is not used, it is followed by a colon.
MSG	The body of the message.

Treatment of Timestamps (RFC 3164)

If the timestamp cannot be parsed, the **Time** event field stores a part of the time that the event was written to the repository (in the InTrust server's time zone). Note that this field is supposed to contain local times. The GMT timestamp is derived from the parsed value. The message contains no time zone information, so it is important that the Syslog device and the InTrust server should best be located in the same time zone; otherwise, the local and GMT timestamps will be wrong.

RFC 5424 Specifics

The following pattern is defined in RFC 5424 (the header is **bolded**):

```
<PRI>VERSION TIMESTAMP HOSTNAME APP-NAME PROCID MSGID STRUCTURED-DATA MSG
```

A message has the following parts:

Field	Details
PRI	Indicates the facility and severity. For details, see Treatment of Facility and Severity Information .
VERSION	Syslog version. The presence of a digit after the PRI field is how InTrust can tell this is an RFC 5424-compliant message. However, it doesn't matter which digit it is.
TIMESTAMP	See Treatment of Timestamps (RFC 3164) .
HOSTNAME	This can be an FQDN, IPv4 address, IPv6 address or conventional hostname. It can also be omitted with "-". Examples of valid host names: <ul style="list-style-type: none"> • Machinename • Myhost.domain.com • 10.30.44.135 • fe80::5d3b:41f:38d2:a1b1%13
APP-NAME	This field identifies the application that sent the message. It can be omitted with "-". InTrust does not process this data.
PROCID	This field is often used to provide the process name or process ID associated with a Syslog system. It can be omitted with "-". InTrust does not process this data.
MSGID	This field should identify the type of message. For example, a firewall might use the MSGID "TCPIN" for incoming TCP traffic and the MSGID "TCPOUT" for outgoing TCP traffic. It can be omitted with "-". InTrust does not process this data.
STRUCTURED-DATA	This is a collection of arbitrary key-value pairs. It can be omitted with "-". InTrust does not process this data.

Examples of valid messages:

- `<165>1 2015-05-11T22:14:15.003Z SUPERHOST1 myproc 8710 - - %% It's time to make the do-nuts.`
- `<165>1 2003-10-11T22:14:15.003Z mymachine.domain.com evntslog - ID47 [exampleSDID@32473 iut="3" eventSource="Application" eventID="1011"]`
Message with structured data in the UTC time zone
- `<140>1 2003-10-11T22:14:15.003+3:00 10.30.44.245 evntslog - ID47`
Message with a non-UTC time zone and IP address instead of host name

For an in-depth description of the format, see [Section 6 of RFC 5424](#).

Treatment of Timestamps (RFC 5424)

The timestamp in a message can contain such details as the time zone and milliseconds. Millisecond information is lost when a message is converted to an event entry. It is also possible that the timestamp is omitted altogether, replaced by "-".

The following are examples of valid timestamps:

```
2015-05-12T19:20:50.52-04:00
2015-05-11T22:14:15.003Z
2015-05-24T05:14:15.000003-07:00
-
```

The following timestamps are malformed:

2015-08-24T05:14:15.000000003-07:00

Too many decimal places (there should be no more than six).

08-24-2015T05:14:15-07:00

The order of units in the date is wrong.

2015/08/24T05:14:15-07:00

You cannot use separators other than “-” and “:” in the time part.

If the timestamp cannot be parsed or it is omitted, InTrust substitutes the current time during event generation (in the InTrust server's time zone). The parsed (or substituted) timestamp goes to the **Date** and **Time** fields of the event. Note that messages are supposed to contain local times.

The GMT timestamp is derived from the resulting value, as follows:

- If the time zone is specified, it is used for offsetting the GMT timestamp.
- A message must have either time zone information or local offset information; if neither is available, the timestamp cannot be parsed.

Mapping of Event Fields

When InTrust generates an event entry based on a Syslog message, it uses the rules outlined in the table below. It shows what happens to the following example message:

- RFC 3164-compliant format
`<34>Oct 14 22:14:15 mymachine su: 'su root' failed for lonvick on /dev/pts/8`
- RFC 5424-compliant format
`<34>1 2014-10-14T22:14:15+03:00 mymachine su - ID47 - 'su root' failed for lonvick on /dev/pts/8`

Event field	Value	In the example above
Log	Syslog	Syslog
Event Type	<i>Severity value, derived from the PRI field. For details, see Treatment of Facility and Severity Information. There are more severities than event types, and they are mapped as follows:</i> <ul style="list-style-type: none">• 0–3: error• 4: warning• 5–7: information	error
Source	Syslog Device	Syslog Device
Category	<i>Facility value, derived from the PRI field.</i>	security
Event ID	0	0
Date	<i>The date the event occurred or was put in the repository. For details, see Treatment of Timestamps (RFC 3164)</i>	10/14/2014

Event field	Value	In the example above
	<i>or Treatment of Timestamps (RFC 5424).</i>	
Time	<i>For details, see Treatment of Timestamps (RFC 3164) or Treatment of Timestamps (RFC 5424).</i>	22:14:15
User	<i>Not used.</i>	
Computer	<i>The HOSTNAME field, if it can be parsed. This is the host where the event occurred.</i> <i>If the host name cannot be parsed or is omitted, then InTrust substitutes the IP address of the host that the message came from.</i>	mymachine
Description	<i>The entire message, restored from the insertion strings it was broken down into.</i>	<34>Oct 14 22:14:15 mymachine su: 'su root' failed for lonvick on /dev/pts/8 <34>1 2014-10-14T22:14:15+03:00 mymachine su - ID47 - 'su root' failed for lonvick on /dev/pts/8
Insertion String #1	<i>The host that sent the message; not necessarily the same host that the event occurred on.</i>	mymachine

Treatment of Facility and Severity Information

In both RFC 3164 and RFC 5424, the PRI field indicates the facility and severity. The following table shows how PRI values are interpreted:

Severity →	emergency	alert	critical	error	warning	notice	info	debug
Facility ↓								
kernel	0	1	2	3	4	5	6	7
user	8	9	10	11	12	13	14	15
mail	16	17	18	19	20	21	22	23
system	24	25	26	27	28	29	30	31
security	32	33	34	35	36	37	38	39
syslog	40	41	42	43	44	45	46	47
lpd	48	49	50	51	52	53	54	55
nntp	56	57	58	59	60	61	62	63

Severity →	emergency	alert	critical	error	warning	notice	info	debug
Facility ↓								
uucp	64	65	66	67	68	69	70	71
time	72	73	74	75	76	77	78	79
security	80	81	82	83	84	85	86	87
ftpd	88	89	90	91	92	93	94	95
ntpd	96	97	98	99	100	101	102	103
logaudit	104	105	106	107	108	109	110	111
logalert	112	113	114	115	116	117	118	119
clock	120	121	122	123	124	125	126	127
local0	128	129	130	131	132	133	134	135
local1	136	137	138	139	140	141	142	143
local2	144	145	146	147	148	149	150	151
local3	152	153	154	155	156	157	158	159
local4	160	161	162	163	164	165	166	167
local5	168	169	170	171	172	173	174	175
local6	176	177	178	179	180	181	182	183
local7	184	185	186	187	188	189	190	191

Searching for Events in Repository Viewer

To browse repositories, use the InTrust Repository Viewer application. This console provides tools for event viewing and on-the-spot audit data analysis. Repository Viewer lets you dispense with SSRS-based reporting if your intention is to examine audit data rather than submit formal reports or provide knowledge for regulations compliance.

The primary feature of Repository Viewer is event searching. Searching supports advanced filtering, grouping and sorting. For your searches to work fast, it is recommended that the repository be indexed. (For more information about indexing, see the [Repository Indexing for Advanced Search Capabilities](#) topic.)

You can do the following with the search results:

- Save search criteria as search folders for future use
- Organize the results in a tree using multi-level grouping
- Apply view filters to further refine the scope of data
- Export the results to create an ad-hoc report

In addition, you can schedule a report to be built from an automatic search and have it delivered by email or saved in a network share.

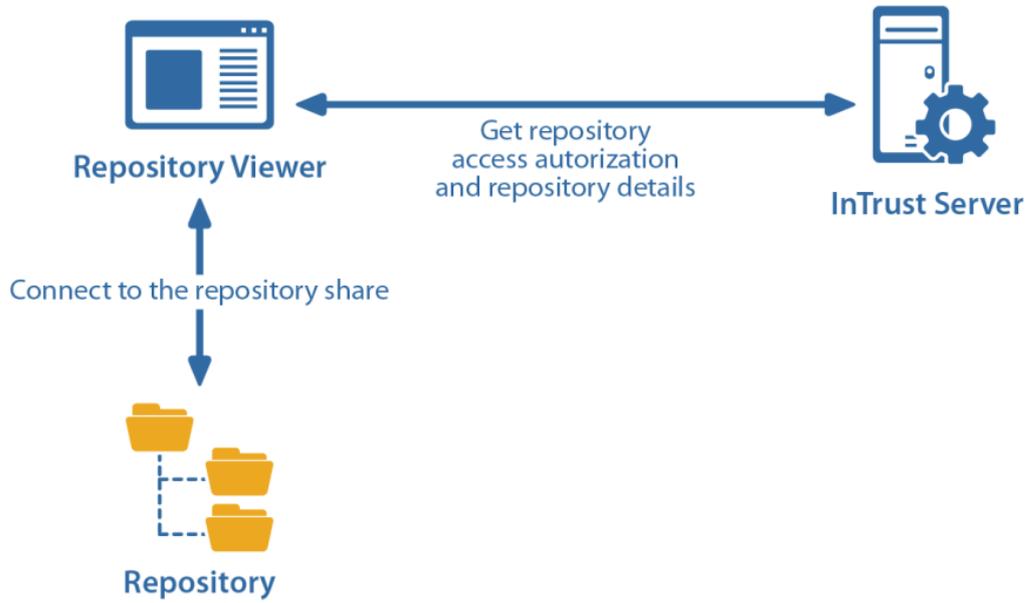
Where to Run Repository Viewer

Repository Viewer does not have complex InTrust component dependencies. However, in the primary use scenario it connects to the repository through an InTrust server, and it matters a lot how far apart the three components are: Repository Viewer, the InTrust server and the repository.

Working with Repository Viewer and InTrust Server

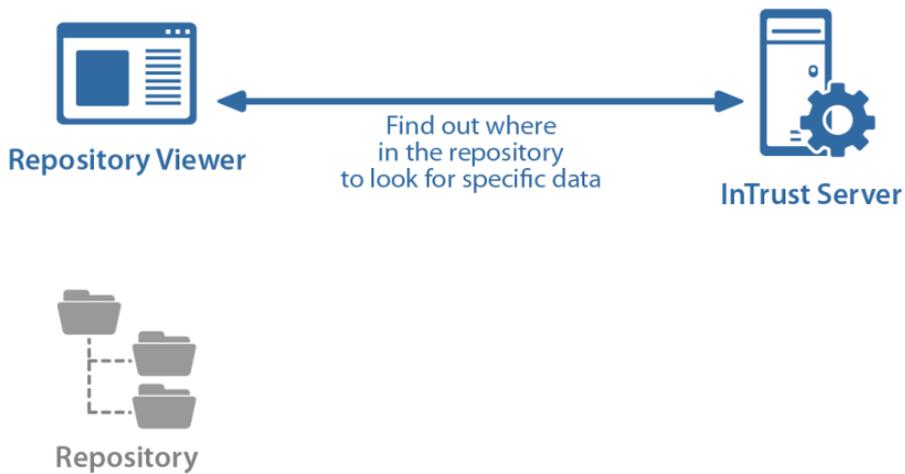
If Repository Viewer opens a repository through an InTrust server, it lets the server manage repository connections.

When Repository Viewer starts working with a repository, it connects to the InTrust server, gets authorization for access to the repository contents, and then connects to the repository.

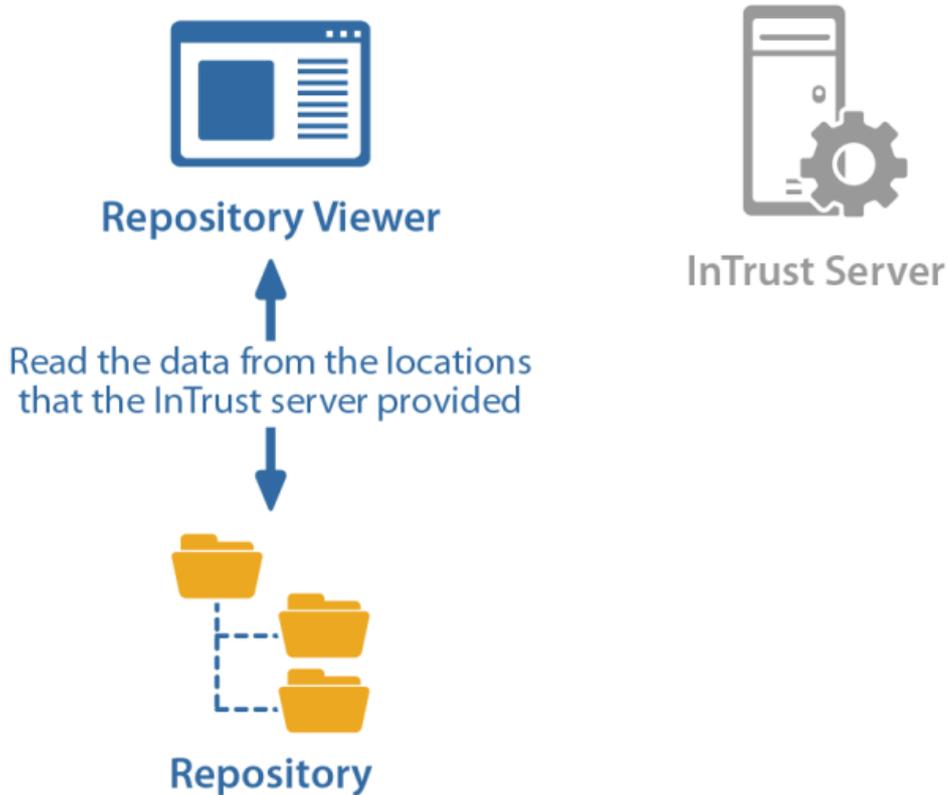


As Repository Viewer continues to work with the repository, it repeats the following steps:

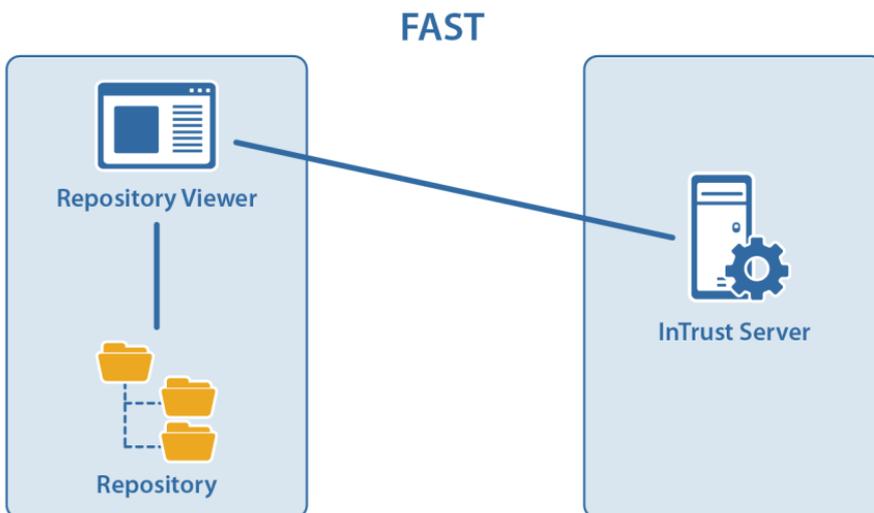
1. Ask the InTrust server for the exact locations of the requested data in the repository structure.

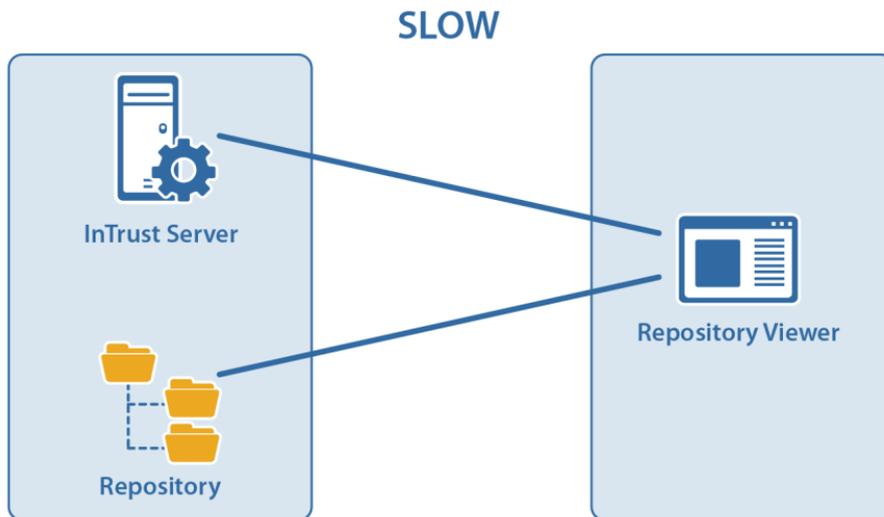


2. After this negotiation, read the data directly from the repository, using the information from the server.



The actual reading of repository data is the most traffic-intensive part of the process. Therefore, you should try to run Repository Viewer as close as possible to the repository share, especially in geographically-dispersed networks. Ideally, they should be on the same computer, but if that is not possible, you should run Repository Viewer on a computer or virtual machine located in the part of the network that is nearest the repository share location. How close Repository Viewer is to the InTrust server is far less important, because the amount of data they exchange is insignificant.





Working with Repository Viewer without an InTrust Server

You can use this option to analyze data from an idle repository; for example, a backup copy of a production repository with historical data.

Getting Started with Repository Viewer

i NOTES:

- When you launch Repository Viewer for the first time, the console asks you to specify the repository to look in.
- Repository Viewer remembers the most recently used repository and opens it automatically on startup.

To open a repository, click **Repositories | Open** in the main menu. You are prompted to select what kind of repository to connect to: idle repository or production repository. These options mean the following:

- **Production repository**
This ensures that the InTrust server you specify handles the communication between Repository Viewer and the repository. Always use this option if the repository you need is managed by an InTrust server and is available for gathering, consolidation and other operations. This method does not lock down the repository index, and multiple instances of Repository Viewer can use its index simultaneously.
- **Idle repository**
This makes Repository Viewer read data directly from the repository without any intermediary components. Use this option only when the repository you need is not attached to any InTrust server. For example, it can be a backup copy of a production repository or an idle repository with historical data. Note that using direct connection locks the index of a repository so that only the first-connected instance of Repository Viewer can use the advantages of indexing.

Production repositories can be grouped together to form *repository groups*. A repository group acts as a single unit: you can run searches on it and create reports as if it were a regular repository. For details about repository group membership, see [Managing Repository Groups](#).

i **NOTE:** Repository Viewer works with repository groups concurrently, but multi-repository searching is not completely overhead-free.

Repository groups are stored in InTrust configuration, and they are available to every instance of Repository Viewer connected to the InTrust organization.

To open a production repository or repository group

1. Select whether you want to connect by specifying an InTrust organization or a specific InTrust server.
2. Select the organization or server.
3. Select whether you want to open individual repositories or a repository group.

The following happens for individual repositories:

- If you select a single repository, it will open in a temporary group. If there is still one repository in the group by the time you finish the session, the group is not saved in InTrust configuration.
- If you select multiple individual repositories or repository groups, a new group will be created for them, and the group will open. It will include all members in your selection.

If you select a repository group, that group will open.

You should always use the index if it is available and up to date. The index makes Repository Viewer operation interactive.

i **NOTES:**

- For access to a production repository, Repository Viewer must be running under any of the following:
 - An account that is listed as an organization administrator.
 - An account which has at least **Read** permissions on the repository and index and is a member of the computer local **AMS Readers** group on the InTrust server that manages the repository (or repositories) and on the InTrust server that Repository Viewer connects to (these may be two different servers).
- Make sure all InTrust servers in the organization have the agent communication port (900 by default) and InTrust Server management port (8340 by default) open for inbound traffic.
- If Repository Viewer connects from a remote computer, inbound TCP ports 1024 to 65535 must be open on that computer for communication with the InTrust server.
- After you have opened a repository or repository group from some InTrust organization, there is a quick way to open other repositories from the same organization. For that, click **Repositories | Change**.

To open an idle repository

1. Specify the local or network path to the repository root folder.
2. To use the index of the repository, in the **Index location** group of options select **Repository folder** or supply a path in the **This location** text box. To continue without an index, select **No index**.

i **NOTE:** For access to an idle repository, Repository Viewer must be running under an account with at least **Read** permissions on the share that contains the repository.

Once you have opened a repository or repository group, the left pane shows the following:

- A navigation tree with the repository structure
The tree represents the repository structure using multiple levels, such as environment (Microsoft Windows or Unix), domain (for Windows only) and computer, down to the event batch file.
- Predefined search folders with search condition presets
These are essentially built-in interactive reports. For details, see [Predefined Searches](#).
- Custom search folders
These are search folders that you create yourself, either based on existing ones or from scratch. For details, see [Custom Searches](#).

The right pane contains search tools.

i **NOTE:** Any tab can be detached and docked freely in the right pane. To detach a tab, drag it away from where it is docked. To dock a pane, drag it onto any of the areas of the view compass that appears. To make it a tab again, right-click its caption and select **Tabbed Document**.

Running Searches

To run a search, click **Go**. The context of your search depends on the following:

1. Where in the navigation tree you are
Selecting a node in the navigation tree means that your searches will include only the events available at that node's level. For example, to look in the entire repository group, select the repository group node; to get events only from a particular repository or computer, select that repository or computer's node.
2. Whether you are using any parameter filters
Running the search without any parameters will show you all events at your current navigation tree level. If you add any filters, they are applied during the search. If you have selected any search folder in the left pane, you are already using the filter set configured for that search folder.

By default, the number of search results that can be displayed at once is capped at 5000. If you reach this limit, consider specifying better filtering conditions. You can also change the search result limit on the **Search Filter** tab.

i **NOTES:**

- The higher the search limit, the more memory is used by Repository Viewer. If you want to increase the search result limit beyond 5000, do it with caution.
- Use filtering by date whenever the date range is known. This speeds up searches considerably.

Predefined Searches

Repository Viewer provides an extensive set of preconfigured search folders out of the box. They will likely cover most of your event analysis needs; consider trying these searches before you begin creating your own. To view and use the searches included by default, expand the **Predefined Search Folders** node. Predefined search folders are available only when you are working with production repositories.

i | **NOTES:**

- Predefined searches are stored in the InTrust configuration database.
- Predefined searches are updated from one InTrust version to another. This can cause upgraded Repository Viewer to fail to find events that the old version was able to find. For details, see [Changes to Event Fields](#). Events gathered to the repository after the InTrust upgrade are fully compatible with the updated predefined searches.

You can freely modify these search folders in the **Search Filter** tab (see [Filter Parameters](#) for details). However, any changes you make are applied only for the current session. The next time you open Repository Viewer, predefined search folders will be in their default state. If you want to save your changes permanently, make a copy of the modified search folder using the **Copy To** button in the toolbar of the **Search Filter** tab. A predefined search folder can be a convenient starting point for creating your own search folder.

i | **NOTE:** The **Copy To** button is available only when an existing search folder is selected. When the filter parameters are configured from scratch, the button is labeled **Save As**.

In addition to the search filter configuration, the saved search folder includes the event list layout. If you have configured grouping and sorting for the search (see [Configuring the Result Layout](#) for details), these settings are preserved.

After you have saved your own search folder, all subsequent changes to it are applied immediately and permanently. See also the [Custom Searches](#) topic.

Changes to Event Fields

The set of fields in events stored in the InTrust repository has been expanded from version to version. Predefined searches in Repository Viewer have kept up with those changes and incorporated the newly-added fields. As a result, predefined searches may not always work as expected on event data that was collected by older versions of InTrust. This topic lists the added fields by InTrust version.

If your search unexpectedly turns up too little old data, you may want to modify the search to exclude recently implemented fields.

Added Between 10.5 and 10.6

Field Name	Field Display Name
Affected_Group	Affected Group

Added Between 10.6 and 10.7

Field Name	Field Display Name
Filer	Filer
New_path	New path
Scope	Scope
Number_of_results	Number of results
Query_filter	Query filter

Field Name	Field Display Name
Attribute_name	Attribute name
Elapsed	Elapsed
Query_type	Query type
TPAM_Operation	Operation
TPAM_Role	Role
TPAM_Target	Target
TPAM_Failed	Failed
UNIX_Result	Result
UNIX_OS	OS
QPMU_Service	Service
QPMU_Master_host	Master host
QPMU_Submit_host	Submit host
QPMU_Submit_user	Submit user
QPMU_Run_host	Run host
QPMU_Run_user	Run user
QPMU_Command_line	Command line
Permissions_Changed	Permissions Changed
Original_Owner	Original Owner
New_Owner	New Owner
Data_Written	Data Written
Permission_level_name	Permission level name
Permission_level_allow_mask	Permission level allow mask
Permission_level_deny_mask	Permission level deny mask
Site_URL	Site URL
List_URL	List URL
List_relative_URL	List relative URL
User_Logon_Name	User Logon Name
Applied_to	Applied to
Inherited_from	Inherited from

Field Name	Field Display Name
Version	Version
Grantee_user_name	Grantee user name
Grantee_group_name	Grantee group name
Field_Name	Field Name
Old_value	Old value
New_value	New value
Attachment_file_name	Attachment file name

Added Between 10.7 and 11.0

Field Name	Field Display Name
UNIX_AUDIT_NAME	Audit Event
UNIX_AUDIT_CLASS	Audit Class
UNIX_AUDIT_CALL	Audit Call
UNIX_AUDIT_TRAIL	Audit Trail
UNIX_AUDIT_COMMAND	Audit Command

Added Between 11.0 and 11.0.5

None.

Added Between 11.0.5 and 11.1

Field Name	Field Display Name
Facility	Facility
Object_New_DN	Object New DN
Object_Old_DN	Object Old DN
Severity	Severity

Custom Searches

If the predefined Repository Viewer search folders do not cover your specific needs, use custom searches: either based on the predefined ones or created from scratch.

i IMPORTANT:

- Starting with InTrust 11.0, custom searches are stored in the InTrust configuration database in a revised format. If you have upgraded InTrust, you have the option of migrating your existing custom searches.
- To create custom searches, you need to make sure your account is an InTrust organization administrator. To view and edit the list of organization administrators, do one of the following:
 - In InTrust Deployment Manager, click **Manage | Configure Access**.
 - In InTrust Manager, open the properties of the root node.

The default organization administrators are the accounts used for installing InTrust and for running InTrust services.

Ad-Hoc Searches

To run an ad-hoc search with parameters, use the **Search Filter** tab, which is under the event list in the default layout. The **Add or Remove Parameters** button lets you customize your search, as follows:

1. Click **Add or Remove Parameters**.
2. In the Select Filter Parameters tool bar that opens, select the parameters that you want to define for the filter. See [Filter Parameters](#) for details.
3. When you have added the necessary parameters, close the Select Filter Parameters tool bar, and specify the values you want to filter by and the operators to use for value matching.
4. Click **Go**.

If you expect to use the same set of parameters in the future, you can save it as a custom search folder. For details, see [Custom Search Folders](#) below.

Custom Search Folders

Any search filter configuration can be saved as a search folder. You can make custom search folders:

1. By modifying predefined search folders and saving your changes, as described in the Predefined Searches topic. This method can save you a great deal of time and effort.
2. By building a set of filters from scratch when only a node in the navigation tree is selected, and saving this.

To create a search folder based on your current filter configuration and place in the navigation tree, click **Save As** in the **Search Filter** tab when it shows your filter settings, and specify the name of the new search folder in the dialog box that appears.

i **NOTE:** The **Save As** button is available only when the filter parameters are configured from scratch. When an existing search folder is selected, the button is labeled **Copy To**.

Mind that the node currently selected in the navigation tree can affect the set of parameters defined for the search folder. For example, if a particular computer is selected, an additional parameter will be automatically added to show events only from this computer. If you want to avoid this, create search folders while the root folder of the repository is selected.

i **NOTE:** Each user's custom search folders are saved in the InTrust configuration database. They are available to all InTrust organization administrators (for reading and writing) and members of the **AMS Readers** local group on the repository-managing InTrust server (for reading).

Organizing Search Folders

To logically nest search folders, organize them into containers:

- To create containers for your search folders in advance, right-click **Custom Search Folders** in the left pane and select **Create Container**.
- To create a container while saving the search folder, click the folder icon in the **Save As** dialog box.

Best Practice: Search Across Event Fields

If you want to find specific information no matter which event field it is in, use the **Any Field** parameter for your search term. This is especially helpful if you are not familiar with the information layout in the events you are working with.

To find this parameter in the Select Filter Parameters dialog box, select the **Primary** option in the drop-down list. **Any Field** is the first item in the list.

Generally, this is a good starting point for refining a search: it let you exclude the fields where you don't want the term to occur instead of trying to include all the fields where it might occur.

Managing Repository Groups

After you have loaded a repository group into Repository Viewer, you can manage its membership as follows:

- Using the **Remove** command in the member repository shortcut menu
- Using the **Add Repository** command in the group shortcut menu

The shortcut menu for a repository group also contains the **Rename Repository Group** and **Delete Repository Group** commands. The **Delete Repository Group** command erases the group from InTrust configuration. The other place where you can delete a repository group is in the Open Repositories wizard; all existing repository groups in the InTrust organization are listed there.

i | **IMPORTANT:** Whenever a repository is added to a group or removed from it, the change is immediately applied in all instances of Repository Viewer connected to an InTrust organization. These changes should be made responsibly.

Filter Parameters in Repository Viewer

Repository Viewer provides a variety of fields to look in. To list all of them, select **All** in the drop-down list in the Select Filter Parameters toolbar. By default, only the normalized fields (such as *Who*, *When* or *What*) are shown.

The parameters include:

1. Regular event fields (available in the Primary set and under All)
2. Additional parameters:
 - The **Insertion strings** set
These are the unnamed insertion strings that events use for storing various information. You can use these fields if you know precisely what they are used for in the events you are working with.
 - The **Resolved insertion strings** set
These are regular insertion strings that have been processed to resolve any GUIDs and SIDs that occur in them. Note that the resolution works only for events that were gathered using InTrust agents.
 - The **Named insertion strings** set
These are friendly labels for regular insertion strings. Note that different types of events use identically-numbered insertion strings for different kinds of data, so you should make sure the meaning is right if you use a named insertion string in your search. Named insertion strings are intended for improving presentation, especially if you are preparing custom searches for someone else to use.
 - The **Normalized event fields** set
See [Normalized Event Fields \(Who, What, When and Others\)](#) for details.
 - The **Any Field** parameter
See the *Best Practice: Search Across Event Fields* section in [Custom Searches](#) for details.
 - The **Custom** parameter
See [Advanced Expression-Based Filters](#) for details.

Configuring Parameters

When you have added a parameter to the **Search Filter** tab, specify the following:

1. The operator to apply
Use the leftmost button in the operator block. The operators are "Equals", "Contains", "Ends with" and so on.
2. The parameter value
This is a combo box where in addition to an explicit value, this can be one of the following options:
 - **Blanks**
Matches if the field is empty.
 - **NonBlanks**
Matches if the field is not empty.
 - **Custom**
Lets you build a logical condition tree that works within this particular parameter; see below for details.

i **NOTE:** In the current version of Repository Viewer, the following issues are known to exist in search filters:

- The value used for the Any Field parameter matches only the beginnings of words.
- The "Contains" operator matches only the beginnings of words.

All the parameters you include in the filter are combined using logical AND—they must all match for the filter as a whole to match. For details about using OR operations, see [Advanced Expression-Based Filters](#).

! **CAUTION:** For some search filter operators, there is no search speedup if the repository is indexed. The following operators cannot take advantage of the index:

- **Not equals**
- **Does not contain**
- **Not like (wildcards)**
- **Does not start with**
- **Ends with**
- **Does not end with**

Custom Logic for Parameters

Selecting **Custom** in the parameter value combo box opens a dialog box that lets you set up multiple matching conditions and manage their flow with the AND and OR operators.

- To change the list, use the **Add Condition** and **Remove Condition(s)** buttons.
- To select conditions, use the leftmost column: you can Ctrl-click, Shift-click and drag-select items.
- To apply the AND operator (meaning, match all of them) to selected conditions, click the 'And' Group button. The grouping will be visualized as a blue line that spans the operators.
- To apply the OR operator (meaning, match any of them) to selected conditions, click the 'Or' Group button. The grouping will be visualized as an orange line that spans the operators.
- To change a group's operator from OR to AND or the other way around, click the line that marks the grouping, or select a member of the group and click the Toggle button.
- To remove one or more conditions from a group, select them and click the **Ungroup** button.

Note that this logic is processed for values of a single parameter. If you want to analyze multiple parameters, see [Advanced Expression-Based Filters](#) for details.

Normalized Event Fields (Who, What, When and Others)

These fields are not present in the original events; they are filled in by InTrust based on knowledge about the contents of regular fields in various types of events. Normalized fields make it easier to retrieve the most important information from the event; you do not have to know which particular original fields contain which kind of information.

The current set of supported normalized fields is as follows:

FIELD	MEANING
What	<p>A brief description of what the event is about. It is related to such fields as Description and Category.</p> <p>Example: For all events that have to do with logging on, the What field says Logon, regardless of the event category, platform where it occurred, or nature of the logon.</p>

FIELD	MEANING
When	When the event was generated. The time is automatically converted to the local time on the computer where Repository Viewer is running.
Where	The computer where the event happened (had effect).
Where From	The name or IP address of the computer from which the activity (such as a logon, or a configuration change) was performed. This is not necessarily the same computer as the one where the activity had effect.
Who	Plain user name of the account that caused the event. Example: Using this field helps you track user activity across platforms: Windows, Unix, VMware and so on.
WhoDomain	The Active Directory domain of the account that caused the event, where applicable.
Whom	The user account that was affected by the event, where applicable. Example: In password change events, this field shows whose password was changed.

i **NOTE:** Use Event-o-Pedia (<http://eventopedia.cloudapp.net/>) to learn more about the events you can audit. This Web site is a knowledge base that helps you find out the meaning, structure and importance of the events you encounter.

Advanced Expression-Based Filters

The **Custom** filter parameter lets you specify expressions for very specific filtering needs that cannot be covered by the built-in options (for example, complex time ranges). The parameter accepts expressions in the REL expression language, which is used for event analysis throughout InTrust. The language is described in the [InTrust Customization Kit](#) document.

The immediate and intuitive advantage of custom expressions is the ability to use logical OR across multiple fields to branch your matching conditions. Effectively, this lets you combine multiple searches.

Examples of expression-based filters:

What you want to find	Expression
Events where the Computer field is "SRV01" or the User Name field is "DOMAIN1\jdoe", but not necessarily both at once.	<code>(Computer = "SRV01") or (UserName = "DOMAIN1\\jdoe")</code>
Events where the Who field is an account that is a member of the Domain Admins group.	<code>member_of(Who, 'Domain Admins', true)</code> Important: This expression works only for global and universal groups, not for domain local groups. It is suitable in this case, because Domain Admins is a global group.

For more advanced REL techniques, refer to the [InTrust Customization Kit](#).

Changing the Business Hours and Non-Business Hours Parameters

The **Business Hours** and **Non-Business Hours** parameters define fixed time patterns, and no user interface is provided for editing these patterns. If you need to adjust the hours for a particular search folder, you can do so using native SQL Server tools, as follows:

1. Run an SQL query on the InTrust configuration database to find the search folder you need. For example:

```
select [Guid], [Query] from [dbo].[SearchItem] where [name] = '<search_folder_name>'
```

This returns the GUID of the search folder and the search query that it uses. Here is a sample search query:

```
<SearchQuery>
  <SimpleCriteria>
    <SimpleCriteria>
      <name>When</name>
      <condition>
        <GroupOperator>And</GroupOperator>
        <Items>
          <DateTimeComparisonCondition_BusinessHours>
            <start_time>8</start_time>
            <end_time>19</end_time>
            <start_dow>1</start_dow>
            <end_dow>5</end_dow>
          </DateTimeComparisonCondition_BusinessHours>
        </Items>
      </condition>
    </SimpleCriteria>
  </SimpleCriteria>
  <FullTextSearchCriteriaItem>
    <FTS/>
  </FullTextSearchCriteriaItem>
</SearchQuery>
```

2. Edit the search query so that it meets your requirements. You need to make changes to the contents of the **DateTimeComparisonCondition_BusinessHours** or **DateTimeComparisonCondition_NonBusinessHours** node. In particular, you need to modify the integer values of the following:

- **start_time**
What time the business or non-business hours start
- **end_time**
What time the business or non-business hours end
- **start_dow**
The first work day in the case of business hours; the first day off in the case of non-business hours (0 through 6 is Sunday through Saturday)
- **end_dow**
The last work day in the case of business hours; the last day off in the case of non-business hours (0 through 6 is Sunday through Saturday)

NOTE: It is assumed that the times you specify are in the time zones of the computers where the events were logged. If you want these original timestamps to appear in Repository Viewer and scheduled reports, make sure the **Local Time** column is displayed in the grid. This column is hidden by default. For details about changing the grid, see [Configuring the Result Layout](#).

3. Overwrite the original search query with your modified version in the configuration database, using the previously extracted GUID to identify the search folder. For that, use an SQL query like the following:

```
update [dbo].[SearchItem] set [Query] = '<modified_search_query_string>' where [Guid] = '<search_folder_GUID>'
```

Examining Event Details

To view the details of a selected event, use the **Event Details** tab. Double-click the event to open this tab. To copy the details, right-click anywhere in the tab and click **Copy to Clipboard**.

Configuring the Result Layout

You can set up event display in the right pane exactly the way you want your search results to be presented, using sorting and grouping.

Organizing the Grid

The default event view in Repository Viewer is a grid, and the default grid layout is a table, where the columns are named after event fields.

You can snap the column names together like building blocks, vertically as well as horizontally, to make compact layouts instead of using a plain table. The grid will use your block layout for every event it displays.

group by area Drag a field here to group by that field

Log	Event ID	When	Who	WhoDomain	What	Where	Where From	Whom
Security	4634	10/31/2014 12:52:51 PM	FRAGGLE\$	FRAGGLEROCKK	Logoff	Fraggle.FraggleRock.local		
Security	4624	10/31/2014 12:52:51 PM	FRAGGLE\$	FRAGGLEROCKK	Logon	Fraggle.FraggleRock.local	127.0.0.1	
Security	4672	10/31/2014 12:52:51 PM	FRAGGLE\$	FRAGGLEROCKK	Special Privileges assigned	Fraggle.FraggleRock.local		
Security	4728	10/31/2014 12:52:49 PM	adm1	FRAGGLEROCKK	Group Member Added	Fraggle.FraggleRock.local		FRAGGLEROCKK\Admin
Security	4737	10/31/2014 12:52:49 PM	adm1	FRAGGLEROCKK	Group Changed	Fraggle.FraggleRock.local		TestGroupGS
Security	4634	10/31/2014 12:52:49 PM	RapidFire	IT	Logoff	Fraggle.FraggleRock.local		
Security	4624	10/31/2014 12:52:49 PM	RapidFire	IT	Logon	Fraggle.FraggleRock.local	10.30.38.196	
Security	4751	10/31/2014 12:52:47 PM	adm1	FRAGGLEROCKK	Group Member Added	Fraggle.FraggleRock.local		FRAGGLEROCKK\Admin
Security	4750	10/31/2014 12:52:47 PM	adm1	FRAGGLEROCKK	Group Changed	Fraggle.FraggleRock.local		TestGroupGD
Security	4746	10/31/2014 12:52:45 PM	adm1	FRAGGLEROCKK	Group Member Added	Fraggle.FraggleRock.local		FRAGGLEROCKK\Admin
Security	4745	10/31/2014 12:52:45 PM	adm1	FRAGGLEROCKK	Group Changed	Fraggle.FraggleRock.local		TestGroupDD
Security	4759	10/31/2014 12:52:42 PM	adm1	FRAGGLEROCKK	Group Created	Fraggle.FraggleRock.local		TestGroupUD
Security	4727	10/31/2014 12:52:42 PM	adm1	FRAGGLEROCKK	Group Created	Fraggle.FraggleRock.local		TestGroupGS
Security	4749	10/31/2014 12:52:41 PM	adm1	FRAGGLEROCKK	Group Created	Fraggle.FraggleRock.local		TestGroupGD
Security	4744	10/31/2014 12:52:40 PM	adm1	FRAGGLEROCKK	Group Created	Fraggle.FraggleRock.local		TestGroupDD
Security	4738	10/31/2014 12:52:40 PM	adm1	FRAGGLEROCKK	User Account Changed	Fraggle.FraggleRock.local		Admin

You may want to hide the fields you do not need or display the blocks that you want to work with. For that, click the icon next to the leftmost block name and change the selection in the **Field Chooser** toolbar. The following fields are available:

- Normalized event fields
- Regular event fields (available in the **Primary** set and under **All**)
- Insertion strings
- Named insertion strings
- Resolved insertion strings

For details about the fields, see [Filter Parameters](#).

Grouping

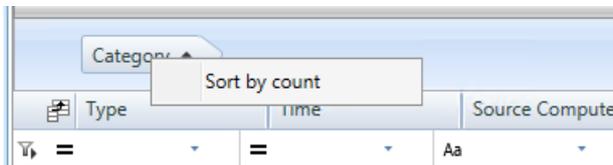
Repository Viewer supports multi-level grouping of events, so that you can organize the results in tree-like views using any criteria. For example, you can group events by log, then by event ID, and then by user.

To use multi-level grouping, in the **Events** pane, drag column names from the event list to the area above the event grid. The event list changes accordingly.

Sorting

To sort the results by a particular field, click that field's block in the grid. Clicking a block repeatedly toggles between ascending and descending order.

Items are sorted by name. However, for groups you also have the option of sorting items by count. To enable it, right-click the block you need in the grouping area and select **Sort by count**.



This option is set independently for each grouping level.

Hiding and Unhiding Events

Hiding and unhiding events is useful when you need to repeatedly locate specific events in the same pool of audit data. This does not change your list of search results, but only specifies which parts of it are shown. It is quicker than redefining search filters and redoing searches every time, and if you are using a search folder, it helps you avoid modifying it.

To configure a view filter, use the controls underneath the column names in the event view: click the operator icon to select the operator, and specify the value to filter by. For details about operators, see the [Filter Parameters](#) topic.

Saving the Results

In any event view, you can export the currently displayed events to a file. For that, click **Report | Save Report** button. In the save dialog box, you have the options of saving the current grid "as is" or running a fresh search without an item limit and possibly with more recent results.

The **Report** drop-down menu also contains scheduling options. For details about scheduled reporting, see [Reporting on Events Using Repository Viewer](#).

Using Pie Charts and Column Graphs

Pie charts and column graphs are graphical alternatives to the grid-based textual representation of search results.

An event list can use multi-level grouping, but pie charts and column graphs work only if single-level grouping is used. In addition, the charts are most informative when they have only a few elements to display. Otherwise, the visual clutter can make them useless.

To switch to a non-default event view, select the **Pie Chart** or **Column Graph** tab.

i | NOTE:

A pie chart or column graph cannot be saved to a file. You can only view it in Repository Viewer.

Case Study: Forensic Analysis of Active Directory Tampering

This example is based on an actual investigation, but the details have been changed. In the example environment, a business-critical server named **acc05** hosts the payroll in a network share. Access to the share is controlled through share permissions. Only the members of the **Finance and Accounting** Active Directory group have read and write access. Jake, the investigator, has grounds to suspect that some of the payroll files have been tampered with, and needs to perform forensic analysis. Here is how he does it using the InTrust-collected audit data from the Security log and Change Auditor File Access Audit event log:

Jake, the investigator, has grounds to suspect that some of the payroll files have been tampered with, and needs to perform forensic analysis. Here is how he does it using InTrust-collected audit data from the Security log and Change Auditor File Access Audit event log:

1. The starting point is the **acc05** computer where the breach supposedly happened. In Repository Viewer, Jake runs a search that shows events from the computer for the past 24 hours. The filter parameters are as follows:
 - Computer: "acc05"
 - When: "Last 1 day"
2. He groups the results by **Who** then by **What**, and checks who accessed the share. In the Where From field, he spots an IP address that needs looking into.
3. Jake checks what else was done from the same IP address. For that, he runs a new search with the **Any Field** parameter set to the suspicious IP address. He finds out that a logon to a domain controller from the suspicious IP address occurred under an administrator account. Jake notes the time of the logon.

4. He then finds out what the administrator account did after the logon to the domain controller. For that, he runs a new search with the **Who** parameter set to the administrator account name and the **When** parameter set to after the logon.
5. It turns out that the impersonator cleared the Security log in an attempt to cover the tracks. However, the events from the log were not lost, because the InTrust agent on the domain controllers was running with log backup enabled.
6. It also turns out that user **david_shore** was added to the **Finance and Accounting** Active Directory group and removed from it shortly afterwards. This is the apparent intruder. To confirm it, Jake checks if this account did anything to the payroll files.
7. He runs a new search with the following parameters:
 - Computer: "acc05"
 - When: "Last 1 day"
 - Who: "david_shore"
8. The search turns up changes to the payroll files. This is clear evidence of a breach, and the perpetrator is now known.

Reporting on Events Using Repository Viewer

In addition to interactive work, Repository Viewer can be used for running scheduled reports based on repository contents. Scheduled reports are essentially searches that have been configured in advance and run automatically at regular intervals.

Therefore, the same considerations mostly apply to scheduled reports as to regular searches (see [Searching for Events in Repository Viewer](#) for details).

Scheduled reports work only on production repositories, not on idle repositories. For details about the difference between them, see [Repository Connections](#).

! CAUTION:

At this time, to set a schedule, you need to start with a custom search. You cannot schedule a report based on repository tree browsing results or a predefined search. If you want a scheduled report based on a predefined search or your current event view configuration, first save the settings using the Copy To or Save As button in the results pane. For details, see the [Predefined Searches](#) topic.

To configure a scheduled report

1. Select the custom search you need in the treeview and click **Report | Schedule Report** in the results pane.
2. On the first step of the wizard that opens, specify the time range that the report should cover. If you are setting up a report with the most recent events going back a specific time period, it is important to pick the right keyword:
 - Use **Last** to specify a period that starts with the beginning of a complete time unit and ends with the end of a complete time unit. For example, a report with a "last 7 days" time range that runs early on Monday will contain events from 12:00 AM the previous Monday to 23:59 PM on Sunday. The **Last** keyword is recommended for reports that run regularly, because it helps make sure that no results are lost between consecutive report runs.
 - Use **This** to specify a period up to the time the report generation begins. This is the same behavior as during searches, but the option should be used with caution for reports. There is no telling when exactly the report will really start building—it may be minutes after the time specified by the schedule. Therefore, the **This** keyword is not recommended for reports that run regularly, because there might be gaps in event continuity from report to report.
 - The **Before** and **Between** keywords do not make much sense in a report. They are best used for one-off searches that you can export to files.
3. Specify the type of report that you want. Depending on the layout in the event grid, you may have a choice of presentation. The **Table** type is always available and has the same layout as the grid. The **Pie**

Chart and **Column Graph** types are available only if data in the grid is grouped by exactly one field.

i | **NOTE:** The **Pie Chart** or **Column Graph** choice is selected automatically if you have clicked **Report | Schedule Report** on the **Pie Chart** or **Column Graph** tab in the results pane, respectively.

4. Specify the desired file format and the delivery method for the report: whether to send it by email, save it in a network share, or both.

i | **NOTES:**

- Not all formats are available for all report types. For example, a pie chart cannot be saved to CSV.
- If you select the CSV format, grouping and sorting settings will be ignored. Your report will contain a plain table of events sorted by time in descending order.
- CSV is the only format without limits on the number of included events. For all other available formats, the number of entries is capped at one million.

If you select to send the reports as email attachments, note the following:

- You can change SMTP settings for email delivery under the list of recipients.
 - Consider setting the maximum attachment size to avoid putting unnecessary load on the SMTP server.
 - If you expect the report files to be large, consider using both delivery methods instead of just email.
 - In the event of email delivery failure due to an exceeded maximum attachment size, the specified recipients get a notification message about this.
 - The report recipients you specify will get not only the resulting reports, but also any messages about possible reporting failures.
5. Set the schedule for the regular report runs. You can specify a very precise pattern. Note that if a lot of reports happen to be scheduled for the same time, they are queued to run one after another, and the start of your report run may be delayed.
If necessary, change the InTrust server that will run the report in the **Server** drop-down list. Report generation is resource-intensive, and this option can help distribute the server load evenly.
 6. Review the resulting configuration and complete the wizard. On the final step, you can select to run the report immediately after you click **OK**.

Tracking Report Progress and Running Manually

Searches that are not scheduled have magnifying glass icons in the left-pane treeview. Searches that are scheduled have the icons of their respective report types.

To view a list of the reports that are currently scheduled, select the **Scheduled Report Status** node in the navigation treeview. For each report, the last known status is shown in the details pane; it can be one of the following:

- **Scheduled**
The report has never been run.
- **Running**

- Succeeded
- Failed

You can run any idle report in this list by right-clicking the report and selecting **Run**. To stop a running report, right-click it in the list and select **Stop**.

i | **NOTE:** An InTrust server can run no more than two reports at once. If more reports have overlapping schedules, they are queued.

Just like elsewhere in Repository Viewer, you can group the results in the details pane by dragging table column names to the grouping area.

Reconfiguring and Disabling Reports

Each custom search has individual reporting settings. If you want to make changes to your reports, consider the following:

- To change the result-related options, such as the report layout or filter settings, modify the search itself. This will affect all subsequent scheduled runs.
- To change the schedule, select the search you want and click **Report | Schedule Report** in the results pane. This will affect the schedule of that particular search.

To disable a schedule, select the search you want and click **Report | Remove Schedule** in the results pane. If you do this while a report is being generated, it will be completed, and the subsequent runs will be canceled.

Integration into SIEM Solutions Through Syslog Forwarding

Events that arrive in a repository can be passed on to SIEM systems that know how to receive, store and index them for analysis. This is known as audit data forwarding and is configured on a per-repository basis.

- [Turning Forwarding On and Off](#)
- [Data Conversion Formats](#)
- [Basic Event Forwarding Scenario](#)
- [Advanced Event Forwarding Scenario](#)
- [Example: Set Up Forwarding to SecureWorks](#)
- [Example: Set Up Forwarding to Splunk](#)

Turning Forwarding On and Off

Forwarding has a dedicated group of settings in the properties of a repository. Use the **Enable forwarding** option to turn it on and off for the repository you are working with.

For details about repository options, see [Managing Repositories](#).

! **CAUTION:** Do not forward events to an InTrust server that listens for Syslog messages, because the messages will arrive with incorrect timestamps.

The following options control how forwarding is performed:

- **Destination host**
The host that listens for forwarded messages.
- **Port**
The port that the destination host uses for listening.
- **Message encoding**
By default, Western European is used.
- **Message filtering**
If you need only a subset of the repository data, you can specify one of the available filters. These filters are really Repository Viewer search folders. If you want to add or modify a filter, open Repository Viewer and make your changes. Your filter will be available the next time you configure forwarding. For details about working with search folders, see [Searching for Events in Repository Viewer](#). Using search folders as filters has some important implications; see [Filtering Specifics](#) below for details.
- **Message format**
The format in which data is expected on the receiving end; see [Data Conversion Formats](#) for details. This

setting has no effect on data that arrives from Syslog devices; such data is forwarded unchanged. Only collected Windows event log data is converted to the specified format.

Filtering Specifics

- Repository Viewer search folders support grouping and sorting, but these settings have no meaning for message forwarding and will be ignored.
- If you edit a search folder that is already used as a filter, your changes will affect the filtering. Consider making dedicated search folders for filtering purposes.
- If a filtering search folder is deleted, filtering is turned off for the repository that used it.
- If you use predefined search folders as filters, note that changes made to them in Repository Viewer are not applied.
- Be careful when specifying the time range for the search folders that will be used as filters. If you set the wrong type of range, this can effectively turn off message forwarding. For example, if you set a time range based on the “Last” keyword, no matches will ever occur. You should not specify a time range for a filtering search folder.

Data Conversion Formats

SIEM appliances expect data in a specific format. For forwarding to be useful, InTrust must convert the contents of the repository to that format before passing them on.

The following output formats are supported:

- Dell SecureWorks
For details, see [Example: Set Up Forwarding to SecureWorks](#).
- IBM QRadar
- Tibco LogLogic
- Splunk (JSON)
For details, see [Example: Set Up Forwarding to Splunk](#).

You can add support for other formats by providing custom format definition scripts.

To specify a different format, select the **Custom Format** item in the **Message format** drop-down list, click **Edit**, and use the editor that opens.

Note the following specifics:

1. Your custom formatting code must implement the **Transform()** function. This function will be used as the entry point by the event forwarding engine. It takes an event object and its sequential number as arguments, and it returns a string.
2. The custom message format will be applied only to the repository you are working with, and will not be replicated to other repositories.
3. Switching from the custom format to the predefined format resets the custom format script to its default state. Back up your custom format script in a file.

For more details about formatting custom messages, study the default formatting script provided in the built-in editor. This is a valid script that replicates the functionality of the predefined SecureWorks forwarding component in InTrust. To change the message format, either edit the **Format** variable or write your own custom

script using this default script as an example. In the **Format** string, event field names enclosed in percent signs (%) will be replaced by their values.

For details about event objects and the InTrust object model in general, see [Customization Kit](#).

Basic Event Forwarding Scenario

This scenario applies if both of the following are true for the repository that you want to forward events from:

1. The InTrust server that manages the repository has at least 8 CPU cores.
2. The rate of incoming events is no more than 2,000 per second.

i | **NOTE:** If you use custom script-based format conversion, the rate of outgoing events will be considerably lower than with the predefined format.

In this case, all you need to do is enable event forwarding for your existing repository, as described in the [Turning Forwarding On and Off](#) topic.

InTrust logs its event forwarding activities and gives you errors if the forwarding queue overflows. If this happens, the event rate is too high, and there will be gaps in the continuity of forwarded events. In that case, you should use the recommendations from the [Advanced Event Forwarding Scenario](#) topic.

i | **NOTE:** The retention threshold for the event forwarding queue is 48 hours by default. Events that are older than the threshold value are dropped from the queue and cannot be forwarded.

Advanced Event Forwarding Scenario

In this scenario, you use a dedicated repository for event forwarding. Create a new collection specifically for the events you want to forward, and select to create a new repository for this collection.

As a final step, you can make sure that disk space is not wasted on the repository contents that you are not going to use. Set up automatic cleanup of the repository contents. For that, use the InTrust Manager console from an extended InTrust deployment to do the following:

1. Connect InTrust Manager to the organization where your repository resides.
2. Create a task and schedule it to run periodically; for example, every day.
3. Within the task, create a single repository cleanup job that clears everything from the repository used for the forwarding.
4. Commit your configuration changes.

For details about performing these steps, see the [Auditing Guide](#).

Example: Set Up Forwarding to SecureWorks

Suppose SecureWorks is already in place in your environment and is used for tracking the operation of Syslog-enabled systems. For Windows network auditing, you use InTrust and Change Auditor. You would like to extend the scope of your SecureWorks coverage to include suspicious user activity in the Windows network.

Make Sure You Have the Data

To capture suspicious administrative activity, you would need to look at the following:

- User session events provided by InTrust
These events provide a deep insight into user logons, logoffs and sessions.
- Change Auditor for Active Directory log
This log provides fine-grained information about all changes to Active Directory.

Confirm that these data sources are used by the collections that work with your repository.

Configure the Forwarding

You need to enable forwarding for the repository with the necessary data. Go to the properties of the repository and, on the **Forwarding** tab, select **Enable forwarding** and specify where the messages should go.

After you have completed the collection setup, confirm that the forwarding is really working. Wait a few minutes for the new settings to take effect. After that, log on to some of the computers that InTrust is watching, and try to make Active Directory changes. Then check on the SecureWorks appliance whether it has registered your activity.

Example: Set Up Forwarding to Splunk

Suppose Splunk is deployed in your environment for analyzing Windows security events. You would like to use InTrust as the forwarding mechanism. The data you need goes to a repository that is set aside specifically for forwarding purposes. The repository has only Windows Security log data.

Get Splunk Ready

! CAUTION: For the sake of speed, the Splunk forwarding component of InTrust uses the UDP protocol, so successful delivery of forwarded data is not guaranteed.

You need to perform two procedures in Splunk (and maybe restart it), as described below.

Step 1: Define a Source Type

To make sure that event fields are recognized correctly, make a specialized source type for incoming InTrust data. If you want to use the Splunk UI for this, configure the options as follows (the last three options are set up in the **Advanced** group):

Option	Value
Category	Structured
Indexed extractions	json
NO_BINARY_CHECK	true

Option	Value
SHOULD_LINEMERGE	false
pulldown_type	1

If you want to skip configuration through the Splunk UI, include the following snippet in the **<Splunk_installation_folder>etc\apps\search\local\props.conf** file:

```
[InTrust]
DATETIME_CONFIG =
INDEXED_EXTRactions = json
NO_BINARY_CHECK = true
SHOULD_LINEMERGE = false
category = Structured
pulldown_type = 1
```

Step 2: Configure a Network Input

In Splunk, add a new UDP network input and apply your new source type to it. Configure the network input as necessary, but make sure you set up the following:

1. It must use the UDP protocol.
2. Specify the source type you defined earlier; in this example, it is **InTrust**.

Make a note of the port number where Splunk will listen for forwarded UDP traffic. You are going to need it for InTrust forwarding configuration.

If you want to skip configuration through the Splunk UI, include the following snippet in the **<Splunk_installation_folder>etc\apps\search\local\inputs.conf** file:

```
[udp://514]
connection_host = ip
index = main
sourcetype = InTrust
```

For details about the various ways that you can add network inputs in Splunk, see the "Get data from TCP and UDP ports" article in the documentation of your version of Splunk.

Step 3 (Conditional): Restart Splunk

If you made your changes by editing configuration files, restart Splunk to apply them; use either the **splunk stop** and **splunk start** commands or the **Restart** action in the Splunk UI. For details, see the Splunk documentation.

Configure the Forwarding

To send data to Splunk, enable forwarding for the repository with the necessary data. Go to the properties of the repository and, on the **Forwarding** tab, select **Enable forwarding** and specify where the data should go.

Select **Splunk (JSON)** as the message format, and specify the correct Splunk host name and the UDP port where the forwarded data is expected.

After you have completed the collection setup, confirm that the forwarding is really working. Wait a few minutes for the new settings to take effect. After that, log on to some of the computers that InTrust is watching, and try to make Active Directory changes. Then open Splunk and check whether your activity has registered.

Further Reading

This guide dealt with the default InTrust configuration. If you are interested in other InTrust capabilities and alternative workflows, or if you need in-depth information about the topics covered here, go to the [InTrust online documentation library](#).

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece — you — to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit <https://www.quest.com/company/contact-us.aspx> or call +1 949 754-8000.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product

Third-Party Contributions

InTrust, version 11.3 contains the third-party components listed below. For third-party license information, go to <http://software.dell.com/legal/license-agreements.aspx>. Source code for components marked with an asterisk (*) is available at <http://opensource.dell.com>.

Component	License and/or Acknowledgement
bison 1.28*	GPL (GNU General Public License) 2.0
boost 1.54*	Boost Software License 1.0
CLucene 0.9	Apache version 1.1 This product includes software developed by the Apache Software Foundation (http://www.apache.org .)
expat 1.95.5	MIT
flex 2.5.25	flex 2.5.25/27
flex 2.5.27	flex 2.5.25/27
flex 2.5.4	flex 2.5.25/27
GNU standard C++ class library 3*	GPL (GNU General Public License) 2.0 with the "runtime exception" Copyright (C) 2004 Free Software Foundation, Inc.
libiconv 1.1	LGPL (GNU Lesser General Public License) 2.1
Net-SNMP 5.0.3	Net-SNMP
NLog 2.0	BSD - Kowalski 2011
OpenSSL 1.0.1d	OpenSSL 1.0 Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved. Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/)
SpiderMonkey 1.5	Netscape Public License Version 1.1 ("NPL") 1.1
Stanford SRP 1.7.5	Stanford SRP Copyright (c) 1997-2001 The Stanford SRP Authentication Project

Component	License and/or Acknowledgement
	All Rights Reserved. This product includes software developed by Tom Wu and Eugene Jhong for the SRP Distribution (http://srp.stanford.edu/). This product uses the "Secure Remote Password" cryptographic authentication system developed by Tom Wu (tjw@CS.Stanford.EDU).
ZLib 1.1.4	zlib 1.2.3 Copyright (C) 1995-2002 Jean-loup Gailly and Mark Adler

* A copy of the source code for this component is available at <http://opensource.dell.com>.

Licenses

Apache 1.1

```

/* =====
* The Apache Software License, Version 1.1
*
* Copyright © 2000 The Apache Software Foundation. All rights
* reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. The end-user documentation included with the redistribution,
* if any, must include the following acknowledgment:
* "This product includes software developed by the
* Apache Software Foundation (http://www.apache.org/)."
* Alternately, this acknowledgment may appear in the software itself,
* if and wherever such third-party acknowledgments normally appear.
*
* 4. The names "Apache" and "Apache Software Foundation" must

```

* not be used to endorse or promote products derived from this
 * software without prior written permission. For written
 * permission, please contact apache@apache.org.
 *

* 5. Products derived from this software may not be called "Apache",
 * nor may "Apache" appear in their name, without prior written
 * permission of the Apache Software Foundation.
 *

* THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED
 * WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES
 * OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
 * DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR
 * ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
 * SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
 * LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF
 * USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
 * ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
 * OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT
 * OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 * =====

* This software consists of voluntary contributions made by many
 * individuals on behalf of the Apache Software Foundation. For more
 * information on the Apache Software Foundation, please see
 * .
 *

* Portions of this software are based upon public domain software
 * originally written at the National Center for Supercomputing Applications,
 * University of Illinois, Urbana-Champaign.

flex 2.5.25/27

Flex carries the copyright used for BSD software, slightly modified
 because it originated at the Lawrence Berkeley (not Livermore!) Laboratory,
 which operates under a contract with the Department of Energy:

Copyright (c) 2001 by W. L. Estes <wlestes@uncg.edu>

Copyright (c) 1990, 1997 The Regents of the University of California.

All rights reserved.

This code is derived from software contributed to Berkeley by

Vern Paxson.

The United States Government has rights in this work pursuant
 to contract no. DE-AC03-76SF00098 between the United States

Department of Energy and the University of California.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

This basically says "do whatever you please with this software except remove this notice or take advantage of the University's (or the flex authors') name".

Note that the "flex.skl" scanner skeleton carries no copyright notice.

You are free to do whatever you please with scanners generated using flex; for them, you are not even bound by the above copyright.

GPL (GNU General Public License) 2.0

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not

price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights.

These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not

covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the

entire whole, and thus to each and every part regardless of who wrote it. Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program

except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made

generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

As a special exception, you may use this file as part of a free software library without restriction. Specifically, if other files instantiate templates or use macros or inline functions from this file, or you compile this file and link it with other files to produce an executable, this file does not by itself cause the resulting executable to be covered by the GNU General Public License. This exception does not however invalidate any other reasons why the executable file might be covered by the GNU General Public License.

END OF TERMS AND CONDITIONS

LGPL (GNU Lesser General Public License) 2.1

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use,

not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the

entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License").

Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that,

in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form

under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License.

Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not.

Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6.

Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the

Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license

restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or

otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances. It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

Copyright (C)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public

License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!

Net-SNMP

License

Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001. An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, Inc. to the project since 2003.

Code has been contributed to this project by many people over the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS

IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) ----

Copyright (c) 2003-2006, Sparta, Inc

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of Sparta, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 6: Cisco/BUPTNIC copyright notice (BSD) ----

Copyright (c) 2004, Cisco, Inc and Information Network

Center of Beijing University of Posts and Telecommunications.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) ----

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003

oss@fabasoft.com

Author: Bernhard Penz

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN

IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

OpenSSL 1.0

License

This is a copy of the current LICENSE file inside the CVS repository.

LICENSE ISSUES

=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License

/* =====

* Copyright (c) 1998-2007 The OpenSSL Project. All rights reserved.

*

* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:

*

* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.

*

* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.

*

* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

*

* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.

*

* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written

* permission of the OpenSSL Project.

*

* 6. Redistributions of any form whatsoever must retain the following

* acknowledgment:

* "This product includes software developed by the OpenSSL Project

* for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

*

* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.

* =====

*

* This product includes cryptographic software written by Eric Young

* (eay@cryptsoft.com). This product includes software written by Tim

* Hudson (tjh@cryptsoft.com).

*

*/

Original SSLeay License

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)

* All rights reserved.

*

* This package is an SSL implementation written

* by Eric Young (eay@cryptsoft.com).

* The implementation was written so as to conform with Netscapes SSL.

*

* This library is free for commercial and non-commercial use as long as

* the following conditions are aheared to. The following conditions

* apply to all code found in this distribution, be it the RC4, RSA,

* lhash, DES, etc., code; not just the SSL code. The SSL documentation

* included with this distribution is covered by the same copyright terms

* except that the holder is Tim Hudson (tjh@cryptsoft.com).

*

* Copyright remains Eric Young's, and as such any Copyright notices in
 * the code are not to be removed.

* If this package is used in a product, Eric Young should be given attribution
 * as the author of the parts of the library used.

* This can be in the form of a textual message at program startup or
 * in documentation (online or textual) provided with the package.
 *

* Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:

* 1. Redistributions of source code must retain the copyright
 * notice, this list of conditions and the following disclaimer.

* 2. Redistributions in binary form must reproduce the above copyright
 * notice, this list of conditions and the following disclaimer in the
 * documentation and/or other materials provided with the distribution.

* 3. All advertising materials mentioning features or use of this software
 * must display the following acknowledgement:
 * "This product includes cryptographic software written by
 * Eric Young (eay@cryptsoft.com)"
 * The word 'cryptographic' can be left out if the routines from the library
 * being used are not cryptographic related :-).

* 4. If you include any Windows specific code (or a derivative thereof) from
 * the apps directory (application code) you must include an acknowledgement:
 * "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
 *

* THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 *

* The licence and distribution terms for any publically available version or
 * derivative of this code cannot be changed. i.e. this code cannot simply be
 * copied and put under another distribution licence
 * [including the GNU Public Licence.]
 */

Netscape Public License Version 1.1

AMENDMENTS

The Netscape Public License Version 1.1 ("NPL") consists of the Mozilla Public License Version 1.1 with the following Amendments, including Exhibit A-Netscape Public License. Files identified with "Exhibit A-Netscape Public License" are governed by the Netscape Public License Version 1.1.

Additional Terms applicable to the Netscape Public License.

I. Effect.

These additional terms described in this Netscape Public License -- Amendments shall apply to the Mozilla Communicator client code and to all Covered Code under this License.

II. "Netscape's Branded Code" means Covered Code that Netscape distributes and/or permits others to distribute under one or more trademark(s) which are controlled by Netscape but which are not licensed for use under this License.

III. Netscape and logo.

This License does not grant any rights to use the trademarks "Netscape", the "Netscape N and horizon" logo or the "Netscape lighthouse" logo, "Netcenter", "Gecko", "Java" or "JavaScript", "Smart Browsing" even if such marks are included in the Original Code or Modifications.

IV. Inability to Comply Due to Contractual Obligation.

Prior to licensing the Original Code under this License, Netscape has licensed third party code for use in Netscape's Branded Code. To the extent that Netscape is limited contractually from making such third party code available under this License, Netscape may choose to reintegrate such code into Covered Code without being required to distribute such code in Source Code form, even if such code would otherwise be considered "Modifications" under this License.

V. Use of Modifications and Covered Code by Initial Developer.

V.1. In General.

The obligations of Section 3 apply to Netscape, except to the extent specified in this Amendment, Section V.2 and V.3.

V.2. Other Products.

Netscape may include Covered Code in products other than the Netscape's Branded Code which are released by Netscape during the two (2) years following the release date of the Original Code, without such additional products becoming subject to the terms of this License, and may license such additional products on different terms from those contained in this License.

V.3. Alternative Licensing.

Netscape may license the Source Code of Netscape's Branded Code, including Modifications incorporated therein, without such Netscape Branded Code becoming subject to the terms of this License, and may license such Netscape Branded Code on different terms from those contained in this License.

VI. Litigation.

Notwithstanding the limitations of Section 11 above, the provisions regarding litigation in Section 11(a), (b) and (c) of the License shall apply to all disputes relating to this License.

EXHIBIT A-Netscape Public License.

"The contents of this file are subject to the Netscape Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/NPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is Mozilla Communicator client code, released March 31, 1998.

The Initial Developer of the Original Code is Netscape Communications Corporation. Portions created by Netscape are Copyright (C) 1998-1999 Netscape Communications Corporation. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[] License"), in which case the provisions of [] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [] License and not to allow others to use your version of this file under the NPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [] License. If you do not delete the provisions above, a recipient may use your version of this file under either the NPL or the [] License."

MOZILLA PUBLIC LICENSE

Version 1.1

1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version,

related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL

PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

``The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____. Portions created by _____ are Copyright (C) _____. All Rights

Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[] License"), in which case the provisions of [] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

Stanford SRP

Licensing

SRP is royalty-free worldwide for commercial and non-commercial use.

The SRP library has been carefully written not to depend on any encumbered algorithms, and it is distributed under a standard BSD-style Open Source license which is shown below. This license covers implementations based on the SRP library as well as independent implementations based on RFC 2945.

The SRP distribution itself contains algorithms and code from various freeware packages; these parts fall under both the SRP Open Source license and the packages' own licenses. Care has been taken to ensure that these licenses are compatible with Open Source distribution, but it is the responsibility of the licensee to comply with the terms of these licenses. This disclaimer also applies to third-party libraries that may be linked into the distribution, since they may contain patented intellectual property. The file "Copyrights" contains a list of the copyrights incorporated by portions of the software.

Broader use of the SRP authentication technology, such as variants incorporating the use of an explicit server secret (SRP-Z), may require a license; please contact the Stanford Office of Technology Licensing (<http://otl.stanford.edu/>) for more information about terms and conditions.

This software is covered under the following copyright:

/*

* Copyright (c) 1997-2001 The Stanford SRP Authentication Project
* All Rights Reserved.

*

* Permission is hereby granted, free of charge, to any person obtaining
* a copy of this software and associated documentation files (the
* "Software"), to deal in the Software without restriction, including
* without limitation the rights to use, copy, modify, merge, publish,
* distribute, sublicense, and/or sell copies of the Software, and to
* permit persons to whom the Software is furnished to do so, subject to
* the following conditions:

*

* The above copyright notice and this permission notice shall be
* included in all copies or substantial portions of the Software.

*

* THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND,
* EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY
* WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

*

* IN NO EVENT SHALL STANFORD BE LIABLE FOR ANY SPECIAL, INCIDENTAL,
* INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER
* RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF
* THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT
* OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

*

* In addition, the following conditions apply:

*

* 1. Any software that incorporates the SRP authentication technology
* is requested to display the following acknowledgment:
* "This product uses the 'Secure Remote Password' cryptographic
* authentication system developed by Tom Wu (tjw@CS.Stanford.EDU)."

*

* 2. Any software that incorporates all or part of the SRP distribution
* itself must display the following acknowledgment:
* "This product includes software developed by Tom Wu and Eugene
* Jhong for the SRP Distribution (<http://srp.stanford.edu/>)."

*

* 3. Redistributions in source or binary form must retain an intact copy
* of this copyright notice and list of conditions.

*/

Address all questions regarding this license to:

Tom Wu

tjw@cs.Stanford.EDU

zlib 1.2.3

License

/* zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.3, July 18th, 2005

Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied
warranty. In no event will the authors be held liable for any damages
arising from the use of this software.

Permission is granted to anyone to use this software for any purpose,
including commercial applications, and to alter it and redistribute it
freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not
claim that you wrote the original software. If you use this software
in a product, an acknowledgment in the product documentation would be
appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be

misrepresented as being the original software.

3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

*/