



Intelligent Response MIR Administration Guide

MIR 2.3

Please visit our forums at <https://forums.mandiant.com/>

MANDIANT Intelligent Response

Disclaimer

Copyright © 2012 Mandiant Corporation. All Rights Reserved.

This documentation and any accompanying software are released “as is.” Mandiant makes no warranty of any kind, expressed or implied, concerning these materials, including without limitation, any warranties of merchantability or fitness for a particular purpose. In no event will Mandiant be liable for any damages, including any lost profits, lost savings, or other incidental or consequential damages arising out of the use, or inability of use, of documentation or any accompanying software, even if informed in advance of the possibility of such damages.

MANDIANT[®], the  logo, Intelligent Response[®], and MIR[®] are registered trademarks of Mandiant Corporation.

Windows[®], Internet Explorer[®], and Windows Vista[®] are registered trademarks of Microsoft Corporation in the United States and/or other countries.

HP[®] and is a registered trademark of Hewlett-Packard Company.

ArcSight[™] is a trademark of ArcSight, Inc.

Table of Contents

1. Introduction	1
1.1. Who Should Use This Book	1
1.2. How This Book is Organized	1
1.3. How to Get Support	2
1.4. For Further Reading	3
2. Introducing MANDIANT Intelligent Response	4
2.1. High-Level System Overview	4
2.2. System Requirements	5
2.3. Installation	7
2.4. Basic Configuration	7
2.5. Configuring the Initial Trust Domain	16
2.6. End to End Function Check	21
3. Appliance Administration	30
3.1. The Appliance Configuration Menu	30
3.2. Controller Backups	31
3.3. Log File Cleanup	35
3.4. Controller Configuration	35
3.5. Controller Diagnostics	42
3.6. Troubleshooting Packages	42
3.7. Field Patch Tracker	43
3.8. Rebooting the Controller	44
3.9. Statistics	44
3.10. The Appliance Upgrade Menu	45
4. Application Administration	47
4.1. Components	48
4.2. Appliance Configuration Details	50
4.3. The Application Database Menu	53
4.4. Discovery	55
4.5. MCIC	56
4.6. RemoteAuth: Active Directory Authorization	56
4.7. Resources	57
4.8. SSL	62
5. MIR Users and Groups	63
5.1. Managing Users	63
5.2. Managing Groups	65
5.3. User Audit Logging	66
6. Understanding the Trust Domain	69
6.1. The Trust Domain	69
6.2. Administering the Trust Domain	76
7. Agent Deployment	85

7.1. The Agent <i>Discovery</i> Service	85
7.2. Agent Installation	89
A. Agent Command-line Reference	101
A.1. Commands and Flags for Using the Agent	101
B. Error Messages and Troubleshooting	108
B.1. Errors, Issues, and Logs	108
B.2. System Reports	115
C. Credits	118
C.1. Component License Notices	118
Index	140

List of Figures

2.1. MIR Architecture	4
2.2. Controller Back Panel	8
2.3. Controller Back Panel	13
6.1. Establishing the TDCA and Issuing a Certificate	72
6.2. Managing the Certificate Revocation List	73
6.3. Controller/Agent Authentication	74
7.1. ADS Functional Overview	86



Chapter 1

Introduction

The latest version of this document is available at <https://forums.MANDIANT.com/topic/official-mir-documentation/>

1.1. Who Should Use This Book

The MANDIANT Intelligent Response *Administration Guide* is for application administrators and other IT personnel who are charged with installation, operation, and maintenance of the MIR product.

To get the maximum benefit from this user guide you should be familiar with TCP/IP networking and system management concepts, such as backup/restoration functions for common IT applications and operating systems. You should also have a solid understanding of the Microsoft Windows operating systems and management of software deployment on those platforms.

This guide is organized around the steps necessary to install and operate the MIR product in an enterprise environment, starting with overall product architecture, continuing with installation of the Controller appliance, and concluding with installation and management of Console and Agent software. For information about the end-user MIR Console and Agents, please consult the *User Guide*.

1.2. How This Book is Organized

This guide contains the following...

Chapter 1, Introduction

This chapter, containing information about the *Administration Guide* and concluding with MANDIANT product support details.

Chapter 2, Introducing MANDIANT Intelligent Response

This describes MIR architecture and provides hardware and software configuration instruction. It concludes with a functionality check, after which you can begin to use your new MANDIANT Intelligent Response system.

Chapter 3, Appliance Administration

The MANDIANT Intelligent Response system includes a web-accessed administration tool that configures and controls many aspects of the system. This chapter provides an overview of administration functions and use.

Chapter 4, Application Administration

This chapter documents a number of common maintenance and management tasks.

Chapter 5, MIR Users and Groups

This chapter documents user and group access control.

Chapter 6, Understanding the Trust Domain

The security of your local network, and especially the security of the MIR network appliance – the keystone to your IR/EED investigations – relies on the use of trust

domains and certificate authorities. In this chapter, you will learn how to secure your MIR installation.

Chapter 7, Agent Deployment

The MIR network appliance relies on host-installed Agents to monitor and collect information from individual systems. Agent installation provides a range of options and capabilities; this chapter provides details on Agent deployment and hosting.

Appendix A, Agent Command-line Reference

MANDIANT Intelligent Response Agents can be installed or run audits from the command line, providing much more flexibility than that available through installer-based use. This appendix provides a full description of command line usage, as well as a number of examples of command line use.

Appendix B, Error Messages and Troubleshooting

This appendix provides guidance on identifying and resolving problems you may encounter during your discovery and analysis workflow.

???

MANDIANT Intelligent Response makes use of several Open Source and Shared Source technologies. For compliance with their various licensing terms, their full license declarations are included in this appendix.

1.3. How to Get Support

MANDIANT provides product support for its users. Telephone support is available via our support line at 877-9MANDIA (877-962-6342). Email support is available from <mirsupport@MANDIANT.com>.

Many customers find our community forums to be a valuable resource. Please join us at <https://forums.MANDIANT.com>.

Before contacting Product Support, please have the following information prepared:

- Your name.
- Your company name.
- Email and telephone contact information.
- Your Controller serial number(s).
- The MTMS version number.
- The Agent version number.
- A detailed description of the problem, including any screenshots, error messages, or issues documents, and a list of steps and conditions that produced the problem.

1.3.1. Finding the Agent Version Number

1. Using Windows Explorer, navigate to the Agent installation directory.
2. Right-click `miragent.exe`, choose **Properties**, and select the **Details** tab.
3. Note the **Product Version** number. This is the value needed by MANDIANT Product Support.

1.3.2. Finding the Console Version Number

1. Start the Console.
2. Select **Help** → **About**.

Two version numbers are displayed on the left, in the form 1.x.x (1.x.xxxx).

3. Note the version number in parentheses. This is the value needed by MANDIANT Product Support.

1.3.3. Finding the MIR Version Number

1. Using a web browser, navigate to the Web Administration **Console**,
`https://[Controller URI or Hostname]/version.xml`

Note the version number. This is the value needed by MANDIANT Technical Support.

1.3.4. Finding the Controller Serial Number

The Controller serial number is located on the left side of the Controller frame when facing the front of the appliance. If the unit is rack-mounted, you may need to remove the Controller from the rack.

1.4. For Further Reading

For a list of books and articles that MANDIANT staff members and others have written about computer security, incident response, and electronic evidence discovery, visit http://www.MANDIANT.com/news_events/books/ and http://www.MANDIANT.com/news_events/articles/.



Chapter 2

Introducing MANDIANT Intelligent Response

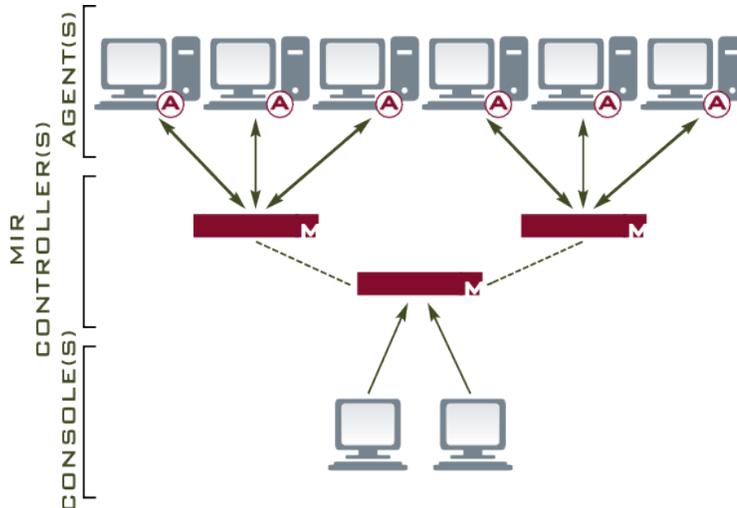
This chapter familiarizes you with the basic product architecture of MIR, then guides you through a basic initial configuration, concluding with a simple function check to ensure its basic services are operating.

2.1. High-Level System Overview

MANDIANT Intelligent Response comprises three main components: Controllers, Agents, and Consoles.

Investigators use the Console application to connect to a Controller appliance. In turn, Controllers connect to Agents. Agents are software components installed on end-user computer systems that enable the Controller to gather information about the system. All data is passed through and tracked by the Controller, ensuring all Consoles are accessing the same information.

Figure 2.1. MIR Architecture



The Controller

The Controller appliance is the heart of MANDIANT Intelligent Response: a combination of specialized hardware and MANDIANT Controller software. Analysis of acquired data is conducted by the Controller, enabling investigators to collaborate on the same set of information.

Future versions of the Controller appliance will allow you to cluster multiple units, enabling you to monitor vast numbers of Agents while still providing a unified view of data, analysis, and reports.

The Agent

The Agent is software installed on a computer system that you want to monitor. Agents allow analysts to gather information about any aspect of a system, from hardware inventory to retrieval of memory, registry, and hard drive data.

Consoles

Consoles provides an interface for using the system, displaying acquired data, and conducting analysis on that data. All data is sourced from the Controller, ensuring every Console has access to the same information and preventing Console users from creating divergent records.

There are two Consoles: MCIC, a simple interface for configuring and running sweeps; and MIR, a robust management tool for the creation and maintenance of host configurations, indicators of compromise, and detailed analysis of results.

2.2. System Requirements

The MANDIANT Intelligent Response system is comprised of hardware and software components. The requirements for these components are as follows:

Network Appliance Requirements:

- Standard 2U, 19" rack-mountable chassis
3.375×19×21 inches (8.6×48.3×53.3 cm) H×W×D
- 32–122° F (0–50° C) operating temperature.
- 60 lbs (27 kg)
- 100–240 VAC, 8.5 A, 50–60 Hz

Console Requirements:

- Microsoft Windows 7, Windows Vista, and Windows XP SP2 or higher.
- 1 GB of RAM.
- Microsoft .NET 3.5 SP1 software installed.
- Microsoft Internet Explorer 6 or higher.
- 75 MB of free disk space.
- Administrator-level privileges for installation.

Agent Requirements:

- 32 bit versions of Windows 7, Microsoft Vista, Windows 2003 SP2, Windows 2000 SP4, and Windows XP SP2 or higher.
- 64 bit versions of Windows 7, Windows 2008 R2, and Microsoft Windows 2003 SP2.
- 512 MB of RAM.
- 12 MB of free disk space.
- By default, the Controller must be able to initiate TCP connections to the Agent on port 22201.
- To test Agent *Discovery Service*, the Agent must be able to initiate TCP connections to the Controller on port 8077.
- Administrator-level privileges for installation.



A long-standing bug in the Windows operating system causes any program using the `windows\temp` directory to fail when the directory has reached its 64K space allotment for files or folders.

In such case, the following message is written to the Agent log file (see *Appendix B, Error Messages and Troubleshooting*):

```
WARNING [Discovery CheckForUpdates]- Agent was not able to
create a temporary download folder. Agent will not update.
(13)
```

The solution is to manually remove some or all of the files in the `windows\temp` directory before attempting to install the Agent again.

2.2.1. Considerations

Before installing the Controller appliance you should consider the following factors regarding its physical requirements and network settings to ensure a smooth deployment.

2.2.1.1. Operating Requirements

The Controller appliance is a ruggedized enterprise-class device, designed to ensure robust operation. Nonetheless, providing an optimal operating environment will help ensure its longevity.

- 32–122° F (0–50° C) operating temperature
- 60 lbs (27 kg)
- 100–240 VAC, 8.5 A, 50–60 Hz, conditioned and protected

2.2.1.2. Network Requirements

The Controller has two 10/100/1000 Mbps Ethernet interfaces (RJ-45 connector). Both interfaces may be used for Console traffic to the Controller and data collections from Agents. If both interfaces are going to be used, we recommend placing them on different subnets. We also recommend assigning static IP addresses and hostnames to these interfaces either directly or via persistent DHCP leases.

Please have the following information available before beginning the configuration process:

- One or two IP addresses if you will directly configure static addresses on the Controller instead of using persistent DHCP leases (on separate subnets if you want to activate both interfaces on the Controller appliance).
- The network mask associated with each address.
- The DNS server IP address, if you will not be using DHCP leases.
- The gateway IP address for each interface, if you will not be using DHCP leases.

2.2.1.3. Firewall Requirements

The Console, Controller, and Agent must be able to communicate with one another. The following provides a summary of connectivity requirements:

Console ⇒ Controller TCP 443

The Console opens multiple SSL connections to the Controller when accessing data. The default listening port on the Controller is TCP 443.

Controller ⇒ Agent TCP 22201

The Controller initiates connections to deployed Agents when performing Audits. The default listening port on the Agent is TCP 22201. Data collections are performed across these connections; the speed of data acquisition is directly related to the bandwidth available between Controller and Agent.

Agent ⇒ Controller TCP 8077

OPTIONAL: Agents initiate connections to the Controller to register themselves using the Agent *Discovery Service*. If the Agent *Discovery Service* is not being used, this connection is not required.

2.3. Installation

Instructions for installation of the MIR Controller appliance are provided with the hardware. Please contact MANDIANT support if you did not receive installation instructions or need additional information regarding physical installation.

The Controller appliance is housed in a 2U rack-mountable chassis. MANDIANT recommends rack-mounting the Controller to ensure proper cooling of the device. The LED panel of the Controller is the front; the panel containing network ports, interface ports, and power cord is the back.



Do not initially connect network cables to the Controller interfaces. You must complete initial configuration of the device before attaching production network connections to the system, as described below.

2.4. Basic Configuration

Once the Controller is physically installed and powered-on, it must be configured before being connected to the local network. This section provides the minimum set of steps necessary to bring the Controller online for normal user operations.

Basic configuration consists of the following steps:

1. *Section 2.4.1, “Configuring the New Controller”.*
2. *Section 2.4.2, “Connecting the Controller to the Local Network”*
3. *Section 2.4.3, “Setting up Controller Accounts and Admin Credentials”.*
4. *Section 2.4.4, “Adding a “Test” User Account”.*
5. *Section 2.4.5, “Configuring Network Time Protocol”.*



The Administration Console supports Firefox 4 or greater. MSIE is not supported.

2.4.1. Configuring the New Controller

The Controller must be configured before it can be successfully attached to your production network. First, you need to connect a network-capable workstation or laptop to the Controller. The workstation/laptop must have Firefox 4 or greater installed in order to access the Administration Console.

2.4.1.1. Connecting to an Unconfigured Controller

1. Ensure that the Controller is powered-on with `eth0` disconnected from the network.
2. Log into the Controller administration account. Using `sudo ifconfig`, find the `eth0` IP address the Controller has assigned itself. This will be a link local address in the range `169.254.1.0` through `169.254.254.255`, or the automatic fallback address `192.168.31.41`.
3. On the laptop or workstation, configure its ethernet port with a static IP address in the same range (excluding the Controllers' address) with a netmask of `255.255.255.0`.

DNS and Gateway settings are not required.

4. Connect the Controller `eth0` network port to the laptop or workstation using a standard Ethernet cable. The Controller supports automatic MDI/MDI-X sensing, eliminating the need for a cross-over cable. The connection must be a direct connection from the Controller to the laptop or workstation, or a connection through the same layer 2 switch will work. Do not use a router.

Figure 2.2. Controller Back Panel

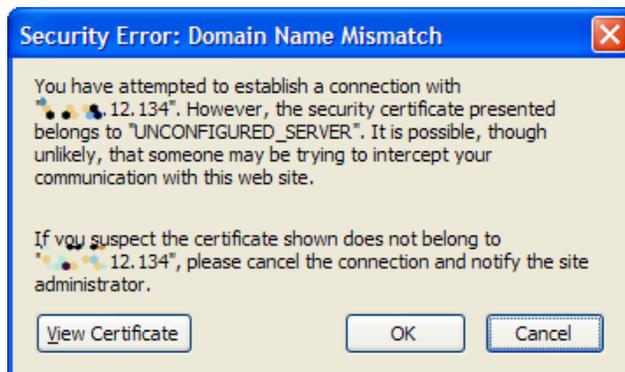


5. Launch a web browser and navigate to the IP address obtained from the Controller in the first step: `https://169.254.x.x/administration/` (or if the controller is configured with the automatic fallback address, `https://192.168.31.41/administration/`) Your browser will generate a certificate error, as shown below.

If you are using Firefox, select **Accept this certificate temporarily for this session** and click **OK**.



6. You will be presented with another error dialog warning you about a domain name mismatch. Click **OK**.



7. You will be prompted for login credentials. Use the following:

```
username: admin  
password: [provided by Sales]
```

1. After logging-in, you will be forwarded to the processes screen for the Controller:

```
MIR_Processes
View the state of currently running MIR processes on the Controller.
Page refresh in - 6
```

Component	Process Status	Availability	# Running
mir_lcd	stopped	dead	1
mir_web_admin	started	alive	1
mir_web	started	alive	1
mir_data	started	alive	1
mir_discovery_server	started	alive	1
mir_mbus	started	alive	1
mir_discovery_proxy	started	alive	1

(Some systems may show `mir_lcd` as started and alive.)

Next, configure the Controller with a static or dynamic IP address, as described in the following two sections, *Section 2.4.1.2, "Configuring the Controller with Static Addresses"* or *Section 2.4.1.3, "Configuring the Controller with DHCP-Persistent Leases"*.

2.4.1.2. Configuring the Controller with Static Addresses

1. In the Administration Console, select **Appliance** → **Config** → **Network**.
2. Provide a **Hostname** for the Controller.
3. For **Enable eth0**, select **Yes**.
4. For **Use DHCP for eth0?**, select **No**. This allows you to use manual configuration options for the `eth0` interface.

Network Configuration

Configure appliance network interfaces. Changing the network configuration will cause the network interfaces to be restarted, along with all MIR Controller services. **WARNING: an invalid network configuration will make the appliance and this administration interface unreachable.** Carefully validate settings before proceeding. Refer to the Administrator Guide for additional information.

* Hostname:

Enable eth0?: Yes No

Use DHCP for eth0?: Yes No

MAC Addr for eth0: 08:00:27:fc:16:d2

Domain Name for eth0:

* IP for eth0:

* Netmask for eth0:

Gateway for eth0:

Broadcast for eth0:

MTU for eth0:

DNS Server Search Order for eth0:

DNS Server 0 for eth0:

DNS Server 1 for eth0:

Enable eth1?: Yes No

5. Set the network configuration appropriate to your MIR environment.



Double-check your entries: an incorrect network configuration may render the Administration Console unavailable.

The fields are as follows

Domain Name (Required)

The domain name suffix for the Controller. If the Controller is `controller.MANDIANT.com` this field would be `MANDIANT.com`.

IP (Required)

The IP address for the interface. Example: `192.168.0.3`.

Netmask (Required)

The netmask for the interface in dotted-quad format. Example: `255.255.255.0`.

Gateway (Required)

The IP address of the default gateway for the interface. Example: `192.168.0.1`.

Broadcast (Optional)

The broadcast address for the interface. If not specified, it defaults to setting all bits within the host portion of the network mask to all 1s. Example:
192.168.0.255.

MTU (Optional)

The MTU size in bytes for the interface. If not specified, defaults to 1500.

DNS Server Search Order (Optional)

A space-delimited list of up to six domain names (256 characters) to be used when a hostname but no domain name is provided for lookup. Example: `company.com MANDIANT.com`.

DNS Server 0 (Optional)

The IP address of a DNS server for this interface. If one is not specified the Controller will not be able to resolve any host names. Example: 192.168.0.2.

DNS Server 1 (Optional)

The IP address of a DNS server for this interface. Example: 192.168.2.2.

6. Click **Update Network Configuration** when you have completed configuring the interface.

The Administration Console will display a message indicating that network changes are being applied. If the IP address has been changed from its prior setting, an error will be displayed when the Administration Console page attempts to refresh (typically within four minutes). In that instance you may manually point your browser to the address you set for the Controller as follows:

```
https://[new ip address]/administration/
```

If you wish to configure `eth1` as well, repeat this task; the same configuration options will be presented for its interface.

The Controller is now ready to be connected to your production network. See *Section 2.4.2, "Connecting the Controller to the Local Network"*.

2.4.1.3. Configuring the Controller with DHCP-Persistent Leases

1. First, obtain the hardware address for the **eth0** interface:
 - a. In the Administration Console, select **Appliance** → **Statistics** → **Network Stats** in the left navigation pane.
 - b. Select **Interfaces** on the right. Details for each active interface will be displayed.

In the **Interfaces** information, locate and note the **HWaddr** value for the `eth0` interface; if `eth1` has also been activated, make note of its value as well.

```

Network Status
Interfaces
eth0  Link encap:Ethernet  HWaddr 08:00:27: : :d2
      inet addr:192.168.56.103 Bcast:192.168.56.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:1311 errors:1 dropped:0 overruns:0 frame:0
      TX packets:1339 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:329008 (321.2 KiB) TX bytes:781174 (762.8 KiB)
      Interrupt:10 Base address:0xd020

lo    Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      UP LOOPBACK RUNNING MTU:16406 Metric:1
  
```

2. Next, configure your network DHCP server to reserve a lease for the `eth0` (and, optionally, `eth1`) interface:
 - a. In the Administration Console, select **Appliance** → **Config** → **Network**.
 - b. Provide a **Hostname** for the Controller.
 - c. For **Enable eth0**, select **Yes**.
 - d. For **Use DHCP for eth0?**, select **Yes**. The Controller will now obtain its IP address from the DHCP server.

3. Click **Update Network Configuration** when you have completed configuring the interface.

If you wish to configure `eth1` as well, enable it similarly.

The Controller is now ready to be connected to your local (production) network.

2.4.2. Connecting the Controller to the Local Network

1. Connect the Controller to your production network using one of the Controller Ethernet interfaces, as appropriate for the configuration of the Controller and network.

Figure 2.3. Controller Back Panel



2. Test the connection by navigating to the Administration Console interface, using a web browser on the network.

2.4.3. Setting up Controller Accounts and Admin Credentials

Once the Controller is properly networked you should change the administrator account credentials from their default settings to settings that comply with your local security policies. You should also set up a test account so that you can complete end-to-end testing.

1. Log into the Administration Console using this URI:

`https://[Controller URI or Hostname]/administration/`



You will still receive the certificate errors as described in ???. Log in using the default administrator credentials.

After setting up accounts, the next section, *Section 2.5, "Configuring the Initial Trust Domain"*, will resolve the certificate errors.

2. Select **Users** in the left navigation bar. A set of controls will appear under it.

From **Current Users** select admin and then click **Edit**.

3. On the right, you will be presented with a user editing form. In **New Password**, type the new password for the admin account. Repeat for **Confirm New Password**, then click **Update User Record**.

If the passwords match, you will be prompted to login to the Administration Console with the new password. Log in again using `admin` for the user name, and the new password. You will be returned to the password change form.

The screenshot shows the 'User Administration' interface. At the top, it says 'User Administration' and 'Manage user accounts and user/group affiliations.' Below that, it identifies the user as 'Administrator (admin)'. The editing form contains the following fields: 'User Name' (pre-filled with 'admin'), 'First Name' (pre-filled with 'Administrator'), 'Last Name' (empty), 'New Password' (empty), and 'Confirm New Password' (empty). At the bottom of the form is a button labeled 'Update User Record'.

2.4.4. Adding a “Test” User Account

1. Log into the Administration Console.
2. Select **Users** → **Create User**.
3. On the right, you will be presented with a user editing form.

Create a test user account, as demonstrated below. Note that in **Assigned Groups** (near the bottom right) the new account is assigned to the **Users** group by default. This is sufficient for function testing.

The screenshot shows the 'User Administration' interface. At the top, it says 'User Administration' and 'Create and manage user accounts.' Below this is the 'Create New User' section, which contains a form with the following fields: 'User Name' (testuser), 'First Name' (Test), 'Last Name' (User), 'New Password' (masked with asterisks), and 'Confirm New Password' (masked with asterisks). Below the form is the 'Groups' section, which has two columns: 'Available Groups' and 'Assigned Groups'. The 'Available Groups' column contains a list with 'Administration', 'ReadOnly', and 'SearchOnly'. The 'Assigned Groups' column contains a list with 'Users'. Between the columns are four buttons: '>>>', '>', '<', and '<<<'. At the bottom of the form is a 'Create User' button.

Note that **User Name** is case-sensitive and must be a unique identifier.

4. Assign group memberships.

An account may also be **ReadOnly**, allowing a user to only read items stored on the Controller. They can not search, run Jobs, or modify any data. **SearchOnly** is similar, except the user can run searches.

5. Click **Create User** to complete. The **Users** → **Current Users** selector will now include the new account.
6. Click **Create User** to complete. The **Users** → **Current Users** selection list on the left will now include the new account.

2.4.5. Configuring Network Time Protocol

To ensure the Controller persistently syncs its clock to a remote time source, you need to configure Network Time Protocol (NTP) services:

1. In the Administration Console, select **Appliance** → **Config** → **NTP**.

NTP Configuration

Configure appliance NTP settings.

Single IP NTP Servers

NTP Server #0:

NTP Server #1:

NTP Server #2:

NTP Server #3:

Multiple IP NTP Servers

NTP Servers #0:

2. By default the Controller is set to synchronize with servers maintained by the *NTP Pool Project*. If your Controller can connect to resources on the internet, these settings may be sufficient.

Otherwise, remove those entries and configure the appliance to sync to your preferred time services.

3. When your configuration is complete click **Update NTP Configuration**.

You may also set the time manually; see *Section 3.4.6, "Setting the Time and Date Manually"*.

2.5. Configuring the Initial Trust Domain

The next step in bringing the Controller fully online is to configure its Trust Domain. This largely involves setting up an internal Certificate Authority and associated SSL certificates. For full details about MIR's security infrastructure see *Chapter 6, Understanding the Trust Domain*.

There are five major steps necessary to set up the Trust Domain:

1. *Section 2.5.1, "Updating OpenSSL Settings"*.
2. *Section 2.5.2, "Generating Keys and Certificates"*.
3. *Section 2.5.3, "Generating Server Keys and Certificates"*.
4. *Section 2.5.4, "Generating a Certificate Revocation List"*.
5. *Section 2.5.5, "Restarting All MIR Processes"*.
6. *Section 2.5.6, "Backing Up the TDCA"*.

2.5.1. Updating OpenSSL Settings

1. On the Initial Configuration page, beside 3.1 **Update OpenSSL Settings**, select the **Here** link.

OR

Log into the Administration Console and select **Application** → **SSL** → **CA** in the navigation panel.

2. You will be presented with three configurable options. The selected defaults will allow you to deploy successfully. See *Chapter 6, Understanding the Trust Domain* for full details on CA configuration.

Trust Domain Certificate Authority Configuration

Update configuration options used by the Trust Domain Certificate Authority (TDCA) to create and administer SSL certificates used in the Trust Domain used for MIR operations.

SSL Key Length:

New Certificates Valid For:

New CRLs Valid For:

SSL Key Length

The length of the public/private key pairs (in bits) associated with all entities in the system. MANDIANT recommends using 2048 bits for maximum security. Other valid values include 1024 and 4096. Note that longer key lengths increase the processing time for network connections.

New Certificates Valid For

The duration (in days) that Controller certificates are valid for. This affects how often you will need to re-install certificates, and is identical in concept to certificate duration for webservers.

New CRLs Valid For

The duration (in days) that Certificate Revocation Lists are valid for. Once a CRL expires, all parties relying upon it must fetch or receive a new copy. MANDIANT recommends the default setting of 180 days.

3. If you changed any settings click the **Update Certificate Authority Configs** button on the page.

2.5.2. Generating Keys and Certificates

1. Certificates On the Initial Configuration page, beside 3.2 **Generate TDCA Keys and Cert**, select the **Here** link.

OR

Log into the Administration Console and select **Application** → **SSL** → **Keys** in the navigation panel.

2. The main page will display the Controller and TDCA key and certificate statuses. Note that Controller (appliance) keys and certs already exist, but that the TDCA keys and certificate have not yet been generated.

Keys

MIR systems can have up to two sets of key pairs: One if the system is acting as the Trust Domain Certificate Authority for this MIR installation, and another for server operations.

WARNING: Deleting the TDCA keys will prevent you from adding additional Controllers to your Trust Domain or issuing updated Certificate Revocation Lists. Regeneration of a TDCA requires re-install of any Agents distributed with the previous TDCA certificate. Refer to the Administrator Guide for additional information.

Appliance Keys & Appliance Certificate	Exists	<input type="button" value="Delete Appliance Keys/Cert"/>
TDCA Keys & TDCA Certificate	Exists	<input type="button" value="Delete TDCA Keys/Cert"/>

3. To create the TDCA key pairs and certificate, click **Create TDCA Keys/Cert**. When complete, the TDCA status will change to **Exists**

2.5.3. Generating Server Keys and Certificates

1. On the Initial Configuration page, beside 3.2 **Generate Server Keys and Signed Cert**, select the **Here** link.

OR

Log into the Administration Console and select **Application** → **SSL** → **Keys** in the navigation panel.

2. The main page will display key and certificate status for the TDCA and the Controller's server keys. Note that although a server key exists, it has not been properly signed by your new TDCA. To create server key pairs and a certificate signed by the TDCA, click **Delete Server Keys/Cert**. You will be asked to confirm the operation: click **OK**.
3. Note the Server Keys and Server Certificate status has changed to **Does Not Exist**.
4. Next, click **Create Server Keys/Cert**. When complete, the **Server Keys and Certificate** status will change to **Exists**.

2.5.4. Generating a Certificate Revocation List

1. On the Initial Configuration page, beside 3.2 **Generate a New CRL**, select the **Here** link.

OR

Log into the Administration Console and select **Application** → **SSL** → **CRL** in the navigation panel.

2. Click on the **(Re)Generate CRL** button on the main page. The CRL will now be displayed in the **TDCA Certificate Revocation List** box on the main page.

Certificate Revocation Lists

Certificate Revocation Lists (CRLs) are used to notify systems within the Trust Domain that a formerly trusted certificate has been revoked and should not be trusted for authentication. Only the TDCA can issue a CRL.

TDCA Certificate Revocation List: Provides the currently active CRL in base64 encoded format.

Revoke Certificate: Revokes the certificate pasted into the input field and adds it to the CRL. Note, this only works if the Controller is configured as the TDCA for the Trust Domain. See the Administrator Guide for more information.

TDCA Certificate Revocation List

```
-----BEGIN X509 CRL-----
MIIBiTBzAgEBMAOGCSqGSIb3DQEBBQUAMDExDzANBgNVBAMUBk1JU19DQTERMA8G
A1UEChMITWFuZG1hbWQxMzA1VTFwOxMDAyMTgyMDQ5NTRaFw0xMTAy
MTgyMDQ5NTRaOAA4wDDAKBgNVHRQEAwIBATANBgkqhkiG9w0BAQUFAAOCAQEAgrzI
iWniAQcjfeZwblxEMIJbJQmXWe2c0cHzRQBJSsqaGbJAvzdJN7xyiITLzuNwK4z7
ggsHNbRNi93YXy1n1GhQc9LhLJKi2fxULwkGgOSmYawUB/NZQ5mfotSiL5LJ0zx7
6Ts/f8nu9eLkarS8H4RDYkkgSdRiWQKdeYc/h6noHpk4GE1EhFp9e4YcLeURnu3g
```

2.5.5. Restarting All MIR Processes

1. Log into the Administration Console and select **Application** → **Components** in the navigation panel.
2. Click **Restart All MIR Processes** in the main page.

Components

Manage MIR Controller processes

The button below will restart all MIR system processes without rebooting the appliance.

To Shutdown or Reboot the appliance click [here](#).

3. After all processes have restarted, the Administration Console page will refresh and report status of the operation.



You may receive the same error dialogs as discussed in ???.

4. Validate all processes were restarted correctly, by navigating to **Application** → **Components** → **Processes**. All processes should show **started** under **Process Status** and **alive** under **Availability**.

MIR_Processes
View the state of currently running MIR processes on the Controller.

Page refresh in - 6

Component	Process Status	Availability	# Running
mir_lcd	stopped	dead	1
mir_web_admin	started	alive	1
mir_web	started	alive	1
mir_data	started	alive	1
mir_discovery_server	started	alive	1
mir_mbus	started	alive	1
mir_discovery_proxy	started	alive	1

2.5.6. Backing Up the TDCA

The TDCA is critical to continued operations in a MIR deployment. You need to ensure you can recover from a failure or administrative mistake that deletes the TDCA. For more information about the TDCA and MIR security infrastructure see *Chapter 6, Understanding the Trust Domain*.



The security of your MIR deployment rests on the security of the TDCA.

You should protect the backup files (`ca.p12` and `server.p12`) as you would any other extremely sensitive data within your infrastructure. The TDCA can create other certificates that are able to access any Agent within the Trust Domain. The server certificate from a Controller can be used to directly access Agents. Backups of these materials should be treated with appropriate physical and logical security measures to prevent compromise.

The TDCA is critical for disaster recovery. **Offsite backups and physical protections are best practices.** See *Chapter 6, Understanding the Trust Domain* for more information about the TDCA.

The passphrase set during key export is as critical as the backup file. If you have the `ca.p12` or `server.p12` file, but lose or forget the passphrase you set during backup, you *will not* be able to restore the TDCA or the server certificate.

1. To back up the TDCA using the Administration Console:
 - a. Log into the Administration Console and select **Application** → **SSL** → **Export Keys** in the navigation panel.

Key Exports

Export the private key and certificates for the TDCA (if present) and server. Exports are in PKCS12 format, encrypted to a passphrase supplied by the administrator. See the Administrator Guide for more information.

Passphrases for PKCS12 files must be at least 4 characters long.

Key/Cert Pair: **MIR TDCA** **MIR Server**

PKCS12 Passphrase:

Repeat PKCS12 Passphrase:

- b. The **Key Exports** page will be displayed. Select **MIR TDCA** and set a passphrase in the boxes provided.



This passphrase is used to encrypt the TDCA public and private keys. The strength of the passphrase directly corresponds to the degree of protection provided to the TDCA keypair.

Click **Export Key**. You will be prompted by your browser to download and save a file named `ca.p12`. Save the file to a safe location.

- c. Select **MIR Server** and, again, set a passphrase. The same precautions apply to this backup as to the TDCA backup.

Click **Export Key**. You will be prompted to download and save a file named `server.p12`. Save the file to a *safe location*.

2. To back up the TDCA through SSH:

- Log into the Controller, tar the keyfiles, and copy them to a *safe location*.

```
ssh -l [admin] [controller]
sudo tar cvfz /home/[admin]/mir_ca.tgz /opt/apollo/etc/mir/ssl/private/mir_ca*
exit
scp [admin@controller]:mir_ca.tgz ./
mv mir_ca.tgz [safe location]
```

2.6. End to End Function Check

2.6.1. Installing the Agent

To test the MIR system end-to-end you will need to install at least one Agent and one Console. The following sections describe how to quickly get an Agent and Console up and running. For complete details, see *Chapter 7, Agent Deployment*.

2.6.1.1. Installing an Agent Using the Windows Installer

The easiest way to install an Agent is with the Windows installer. This method lets you quickly perform a default install, but limits your customization options. By default the Agent is installed as a Windows Service listening on TCP port 22201.

You must have Admin-level privileges to install the Agent.

1. Create a folder named `MIR_Install` on the Host or on portable media.
2. Copy `AgentSetup.msi`, the Agent installer, from the MIR software CD to the `MIR_Install` folder.
3. Copy `conf.xml` – created in *Section 7.2.1.2, “Generating the Agent Configuration File”* – to the `MIR_Install` folder.
4. Start `AgentSetup.msi` and follow the prompts:
 - a. In the **Welcome to...** window, click **Next**.
 - b. In the **License Agreement** window, select **I Agree** and click **Next**.
 - c. In **Select Installation Folder**, accept the default installation destination unless you have reason to install to a different location. Clicking **Disk Cost...** will display the available and required drive space for the installation.

Click **Next** to continue.
 - d. In the **Confirm Installation** window, click **Next**. A progress bar will advance as the software is installed.
 - e. In the **Installation Complete** window, click **Close**.

Note that Agent installation does not require a system restart.

2.6.2. Installing the Console(s)

Next, install the Console on a system that is attached to the local network. You can use the same system as the one used above for the Agent installation. By default, the Console must be able to initiate TCP connections to the Controller on port 443. If your installation media includes MCIC, both the Administration Console and MCIC Console will be installed.

2.6.2.1. Installing the Console

The Console has the following requirements:

- Microsoft Windows 7, Windows Vista, and Windows XP SP2 or higher.
- 1 GB of RAM.
- Microsoft .NET 3.5 SP1 software installed.
- Microsoft Internet Explorer 6 or higher.
- 75 MB of free disk space.
- Administrator-level privileges for installation.

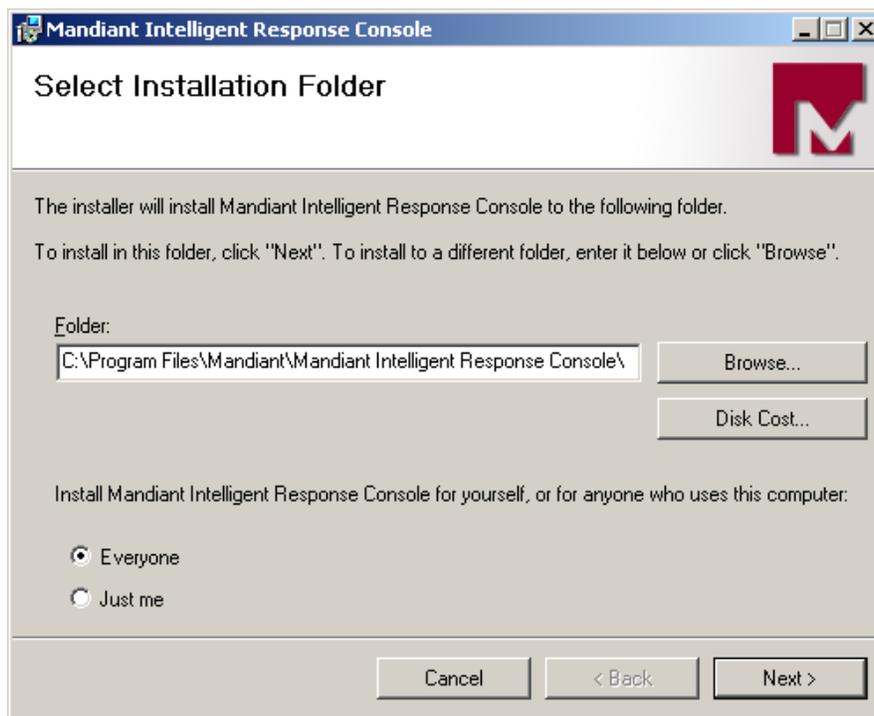


You may wish to export your custom IOCs before you uninstall the old Console.

1. Copy `ConsoleSetup.msi`, the Console installer, to your workstation.
2. Start `ConsoleSetup.msi` and follow the prompts:
 - a. In the **Welcome to...** window, click **Next**.
 - b. In the **License Agreement** window, select **I Agree** and click **Next**.
 - c. In the **Additional Tasks** window, enable the options you want to use, then click **Next**.
 - d. In **Select Installation Folder**, accept the default installation destination unless you have reason to install to a different location. Clicking **Disk Cost...** will display the available and required drive space for the installation.

Select **Everyone** or **Just me** depending on whether you want all users with access to this workstation to have the ability to use the Console, or want access to be restricted to just your own account.

Click **Next** to continue. A progress bar will advance as the software is installed.



- e. In the **Installation Complete** window, click **Close**.

2.6.3. Creating and Running a Test Job

To validate that the Controller, Console, and Agent are properly configured, you will create a Host Audit Job and obtain information from the Agent. The data will be retrieved by the Controller and stored. You will then be able to review that data using the Console.

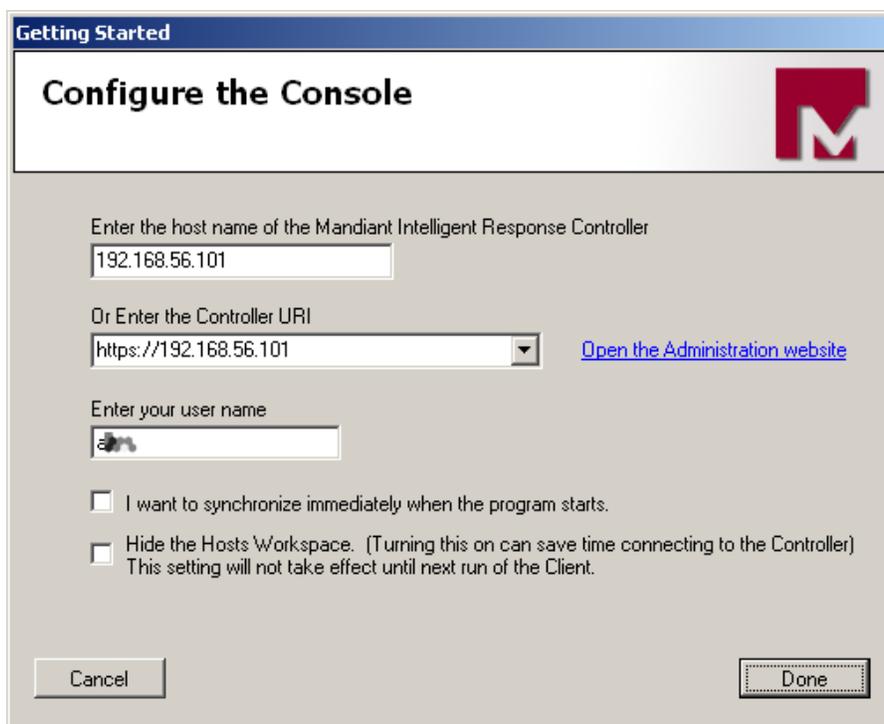
2.6.3.1. First-Run Configuration of the Console

Before starting the Console the first time, you will need the following information:

- The hostname or IP address for the Controller to which you will be connecting.
- Your user name.

To start using the Console:

1. Double-click the Desktop shortcut,  **MANDIANT Intelligent Response Console**. If the shortcut is not available, choose *Start → Program Files → MANDIANT → MANDIANT Intelligent Response Console .
2. If this is the first time the Console has run, a **Connect to Controller** configuration window will be shown:



- a. In **Enter the host name...** provide the host name or IP address of the Controller. As you type the **Controller URI** field will show a best-guess URI address which can be corrected as needed.
- b. In **Controller URI** provide the `https:` or IP address of the Controller.
- c. In **Enter your user name** provide your user name. When you connect, you will be asked for your password as well.

If you are an MIR Administrator testing your installation, this will be the account you created in *Section 2.4.4, "Adding a "Test" User Account"*.

- d. For testing, leave the **I want to synchronize immediately when the program starts** and **Hide the Hosts Library** options unchecked.

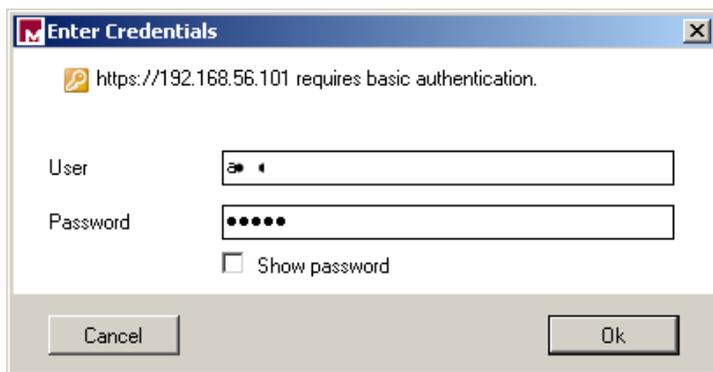
- e. Click **Done**.

If the Console has been previously configured, you will not be presented with the **Connect to Controller** window.

2.6.3.2. Logging On

If the Console has a valid TDCA and Controller address, you can log into the system:

1. In the tool bar, click  **Home**.
2. In the **Enter Credentials** window, provide your user name and password, then click **Ok**.



If the TDCA is not recognized, you will be presented with an alert window. You can continue to log in, but need to be aware that you are operating with reduced security.

3. The Console status bar will indicate the system is synchronizing. Depending on network traffic and the amount of data to be synchronized, this may take a few seconds.

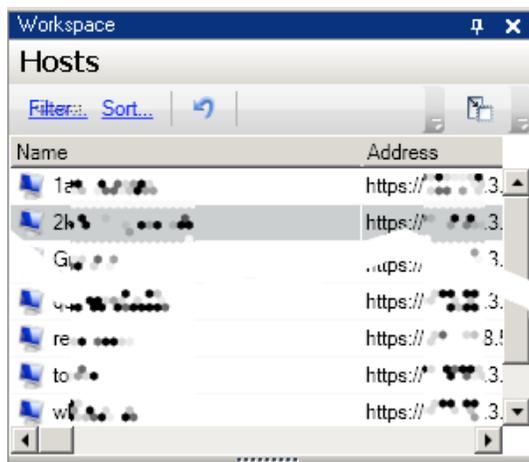
Host Audit Jobs collect data from deployed Agents and return it to the Controller for you to review and analyze. There are several methods for running a Job against a Host. In this example, you will select an existing Host to run a Job against. In this instance, it will be the Host you installed an Agent on in *Section 2.6.1.1, "Installing an Agent Using the Windows Installer"*.

First, find the Host. If it is available through *Section 2.6.3.3, "Selecting an Existing Host"*, below, follow the steps in *Section 2.6.3.4, "Manually Adding and Configuring a Host"* after it. For the latter, you will need to determine the Host's IP address.

2.6.3.3. Selecting an Existing Host

If the Controller is already aware of an Agent, the Agents Host resource is found in the Hosts Library:

1. In the **Workspace** panel on the left of the Console, click **Hosts**. The Hosts Library will be displayed at the top of the panel. (If the Hosts Library button has been hidden, choose **Libraries** → **Hosts**. The Hosts Library will be shown in a *Viewer* tab.)



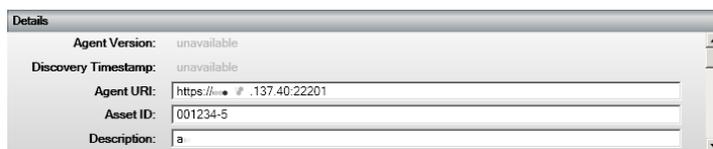
2. Double-click the Host name to display its current IP address, description, and any Audits that have been run against it.

2.6.3.4. Manually Adding and Configuring a Host

You can add a Host Resource manually if you know the Host's IP address or hostname:

1. Choose **File** → **New** → **Host**. A tab named **New Host** will open in the *Viewer/Editor* panel, where you will create the new Host Resource.
2. In the **Details** area the only *required* information is **Agent URI**. It has the format `https://[ip_address|hostname]:[port]`. By default Agents listen on TCP port 22201.

For example, if you are connecting to a Host with an IP address of 10.201.137.40, you would enter `https://10.201.137.40:22201` in **Agent URI**.



Providing an **Asset ID** or **Description** is optional. You can provide that information at this time if you know it. If not, it can be added later.

3. Click **Save** in the **New Host** tool bar.

You will be prompted to enter a name for the Host. You can enter any next you like; common convention is to enter the actual hostname of the host system, if one exists. The **New Host** tab name will be replaced by the new name.

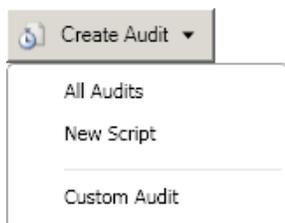
Click **Ok** to save this new Resource.

If the **Hosts** Library is visible, you'll see the new Host in its list as soon as the Controller sends it to the Console.

2.6.3.5. Configuring a Custom Audit Script

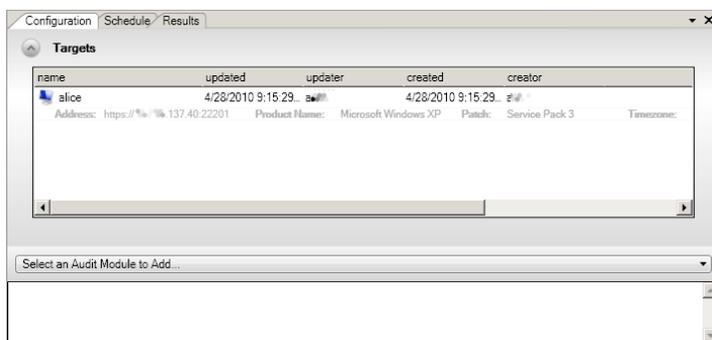
Once you have identified a Host, the next step is to create and configure an Audit. The Audit with its list of Hosts is called a Job.

1. In the **[Host Name]** tool bar, click  **Create Audit** and choose **Custom Audit**.



A new *Viewer/Editor* tab named **Custom Audit of [Host Name]** will be created.

2. Select the **Configuration** sub-tab. The **Targets** area will list the Host; below that is a **Select an Audit Module to Add...** selector.



3. From **Select an Audit Module to Add...**, choose the type of Items you wish to collect.

For installation testing, choose the **System Information** module.

4. Click  **Save** in the *Editor* tool bar. Provide a **Name** for the Job and click **Ok**. Note the *Viewer/Editor* tab name changes from **Custom Audit of [name]** to the new name.

Leave this tab open; you will be using it to run the Job and view the results.

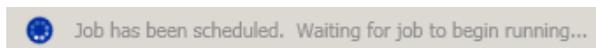
2.6.3.6. Running a Custom Audit Script and Viewing Results

If you wish to run a Job immediately after creating it (see *Section 2.6.3.5, "Configuring a Custom Audit Script"*), you can do so from the *Viewer/Editor* tab:

1. Click  **Run Immediately**.

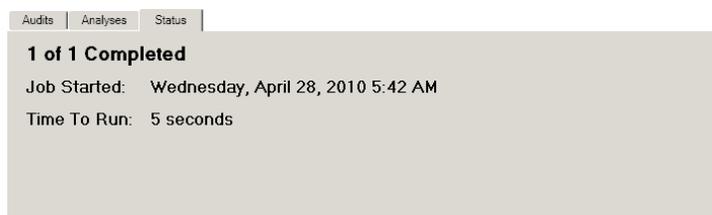
The Console will send the Job to the Controller, as indicated in the *Editor* status bar. While the Job queues and executes the *Editor* will prevent you from making changes or queuing additional jobs.

There may be a minute delay as the Job waits in the queue; it will be executed the next time the Controller checks for waiting Jobs. During this period the **Schedule** tab will be selected, showing the queued Jobs.

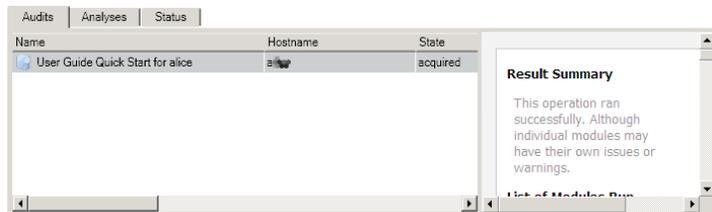


When the Job begins executing the *Viewer/Editor* will switch to the **Results** tab and list each Result Set that has or will be created. The **State** column will be periodically updated as long as the **Click to turn Automatic Refresh Off** button remains active (blue).

Selecting a Result Set will activate its Status sub-tab. The **Status** sub-tab shows the number of Host Audits completed, the time the Job started, and the time that has elapsed.



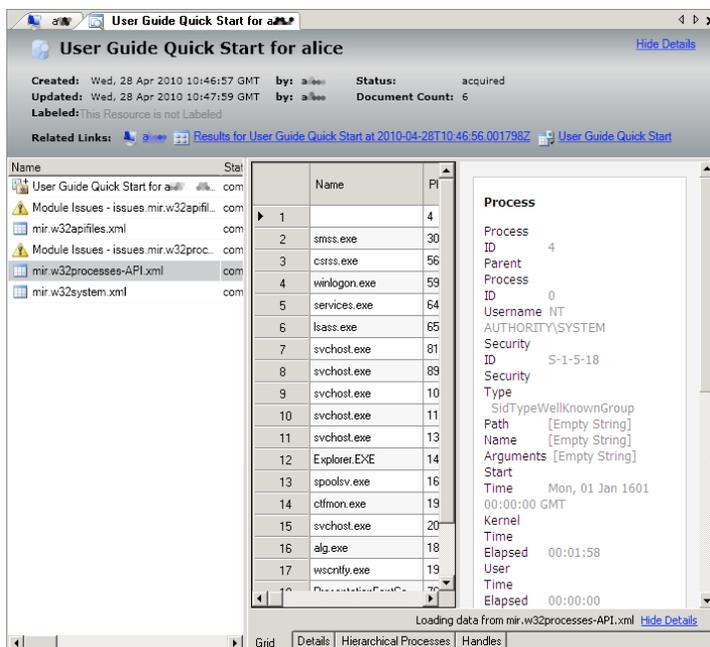
2. When **State** indicates **acquired**, click the **Audits** sub-tab. On the left is a list of each Audit Job. Click an entry to see a summary of the Audit Modules that were run.



3. In the **Audits** sub-tab, double-click Audit Job entry to display a detailed list of its documents. This is like clicking a link on a web page: the *Viewer/Editor* controls are replaced with content from a new "page." As with a web browser, you can use the  **Back** button to restore the previous view.

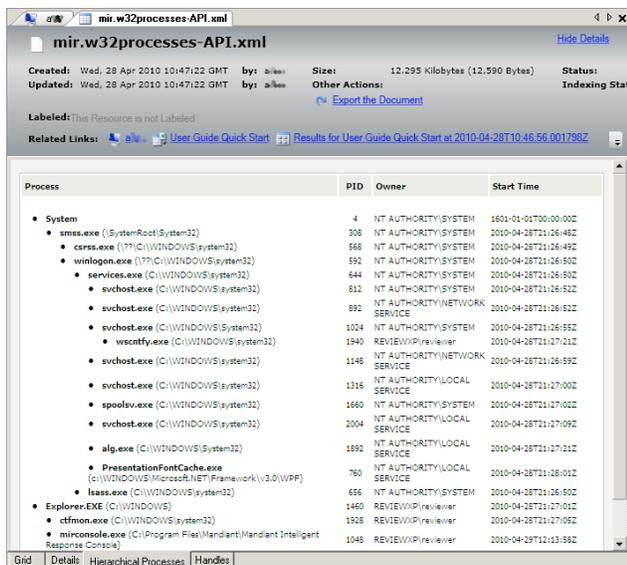
The new page shows a list of documents on the left and, when you click one of the documents, the contents of that document to the right. Double-clicking a document replaces the two-pane view with a full-size view of the document. You can use the  **Back** button to restore the two-pane view, and click it again to restore the *Viewer/Editor* page.

4. Clicking a document in the two-pane *Viewer* will give you a Console window looking similar to this:



- Double-click a document to display its contents. For many documents, a set of tabs at the bottom of the *Viewer* will provide different views of the document.

For example, opening a `mir.w32processes-API.xml` document and selecting the **Hierarchical Processes** tab will display information similar to this:



2.6.3.7. Function Check is Completed

Completing the steps above confirms basic operation for the MIR Controller, Console, and Agent components. Later chapters detail all of the configuration options available for the system, and provide additional information on advanced maintenance topics, such as security configuration, wide-scale agent deployments, and backup/restoration capabilities.



Chapter 3

Appliance Administration

Chapter 2, *Introducing MANDIANT Intelligent Response* provided an overview of the most critical configuration and administration tasks when bringing a new Controller online and testing it with the Console and Agent components. This and subsequent chapters provide an in-depth look into administration and maintenance topics.

The Controller appliance is associated with three levels of administration: appliance-level environmental items (network, logging configuration, etc), and MIR application-level items (application-level timeout behaviors, database backup, certificates, etc), and user/group management. This chapter discusses the appliance-level issues that should be addressed during standard MIR configuration and maintenance.



The administration interface supports Firefox 4 and greater. MSIE is not supported.

3.1. The Appliance Configuration Menu

Appliance configuration functions are found in the Appliance menu in the Administration Console. You can access the Administration Console by navigating to `https://[Controller URI or Hostname]/administration/`. The Administration Console is only available to users in the Administration group (see Section 5.1, “Managing Users”). The default admin account is always able to access this interface.

When you select **Appliance**, the interface should look like this:

The screenshot shows the Mandiant Administration Web Console interface. At the top, the Mandiant logo is displayed with the tagline "INTELLIGENT RESPONSE". Below the logo, the text "Administration Web Console - 172.16.12.128" is shown. The main content area is divided into two sections. On the left, a dashed box highlights the "Administration" menu, which includes "Appliance" (selected), "Backup", "Cleanup", "Config", "Diagnostics", "Packages", "Reboot", "Statistics", and "Upgrade". Below this is the "Application Users" section. On the right, the "Appliance Administration" section is visible, containing the text: "View the running status of the MIR appliance and change configurations. Consult the Administrator Guide for additional information." Below this text are two links: "Controller Version" and "Disk Statistics".

Selecting the **Controller Version** link on the right will display the current version of Controller software installed on the appliance. Selecting **Disk Statistics** will display current disk usage for all volumes in the Controller.

The Controller configuration options are:

Backup

Manages data backup tasks.

Cleanup

Allows deletion of logs files to recover disk space.

Config

Provides access to network, time, NTP, and logging settings.

Diagnostics

Provides access to simple network troubleshooting utilities.

Packages

Lists the currently-installed packages.

Patches

Lists field patches, with their version and status.

Reboot

Reboots the Controller.

Statistics

Provides version information and usage statistics.

Upgrade

Used to upgrade the Controller software.

3.2. Controller Backups

The **Backup** menu configures backup and restoration of the Controller database. You may choose to backup to a local storage media, or to a remote storage server.

The **Backup** options are:

[default]

Configure backup destination.

Backup/Restore Status

View the backup/restoration activity log.

Restore

Restore a Controller backup.

The Controller provides facilities to backup its complete contents and configuration to a remote fileserver using either Network Filesystem (NFS), `rsync`, or a storage device directly

attached to the Controller using USB or Firewire. Backups may be subsequently restored, and will restore all aspects of the Controller to the state at the time of the backup.

This section discusses the Controller configuration for executing backup, and assumes you have either a correctly configured backup server that offers NFS or `rsync`, or a properly formatted USB or Firewire drive with enough storage to accept a backup image from the Controller. See the notes at the end of this section for further details.



The backup process does not back up the Trust Domain Certificate Authority or Server SSL certificates and keys. See *Chapter 6, Understanding the Trust Domain* for more information about backup and restore for the TDCA.

3.2.1. Configuring Local Storage Backup

USB or Firewire drives may be directly attached to ports on the back of the Controller appliance. To be used for backup, drives must be formatted with an XFS, `ext2`, or `ext3` filesystem. Limitations with `FAT32` and maximum file sizes make it incompatible with storing large system backups. `NTFS` is unsupported in this release as a backup format.

1. In the Administration Console, select **Appliance** → **Backup**.
2. On the right, select **Attached Storage**. Locally-attached non-primary media will be automatically detected.
3. From **Backup Target Device** select the storage device that will serve as your backup media.

Several common Linux formats are supported; Windows `vfat` and `fat32` formats are explicitly disallowed. Other alternative formats may be compatible with the Controller, but their functionality is not guaranteed.

If the drop-down says **No Attached Storage Devices** the most likely reason is that you are using a backup drive with an unsupported file system.

4. Start the backup process by clicking **Execute Backup**.



To ensure full backups with a minimum of potential data issues, users should not be logged into the Console application and the Controller should not be collected data from Agents during the backup operation.

3.2.2. Configuring Remote Server Backup

1. In the Administration Console, select **Appliance** → **Backup**.
2. On the right, select **Remote Storage**.
3. From **Backup Protocol** select either **nfs** or **rsync**.



`rsync` should be used if you wish to perform incremental backups.

In **Target Host** provide the hostname or IP address for the remote backup server.

In **Target Directory** provide the backup server directory path to which the backup will be saved.



When using `rsync`, providing an existing target directory will cause only items modified since the last backup to be transferred to the backup server.

When using `NFS`, providing an existing target directory will overwrite the contents of that directory.

4. If you are using a password-protected server, provide an appropriate **User Name** and **Password** for the backup process. The user must have write permissions for the target directory.
5. Start the backup process by clicking **Execute Backup**.



To ensure full backups with a minimum of potential data issues, users should not be logged into the Console application and the Controller should not be collected data from Agents during the backup operation.

3.2.3. Viewing the Backup Status

The Controller maintains a log of backup activity. You may view this log at any time:

- In the Administration Console, select **Appliance** → **Backup** → **Backup/Restore Status**.

The status display automatically updates every ten seconds. You may manually update the view by clicking **Refresh Status**.

3.2.4. Restoring the Controller Appliance from a Backup



Users should not access the Controller while it is being restored. Restoration destroys all data currently stored in the Controller and replaces it with that in the restore bundle.

1. Select **Appliance** → **Backup** → **Restore**. On the right, select **Local Storage** to restore from a USB or Firewire device; or **Remote Storage** to restore from a network backup.



The **Local Storage** option is available only if a disk has been attached to the Controller. Otherwise, the **Remote Storage** screen will be presented by default.

2. If you selected **Local Storage**, select the local storage device.

If you selected **Remote Storage**:

- a. From the **Backup Protocol** drop-down list, select **NFS** or **rsync** as suitable.
 - b. In **Target Host**, type the IP address or hostname of the backup server.
 - c. In **Target Directory**, provide the directory path on the remote host in which the backup files were stored.
 - d. Provide a **User Name** and **Password** for the backup server. The user must have read permissions for the target directory.
3. Test the network backup configuration by clicking **Test Restore**. The test may take several minutes.



While the test is running you will be forwarded to a status page that displays its progress.

4. If the test is successful, click **Execute Restore** to start the restoration process.



While the restoration is running you will be forwarded to a status page that displays its progress. You may also access this page using **Appliance** → **Backup** → **Backup/Restore Status**.

When the restore is complete, the Controller is ready for continued operations. You should check that the restored configuration (e.g. network settings, users, and so on) is correct and up-to-date. Any changes that had been made after the backup date of the restored data will, of course, have been overwritten by the restored data.



Any users who were accessing the system before the restore should restart their Consoles before continuing to use MIR. Also, note that during the restore operation the Administration Console may not be as responsive as it usually is. Refresh your browser view periodically if this happens.

3.2.5. Notes on Backup Server Configuration

While configuration of generic NFS and `rsync` services is beyond the scope of this guide, there are several requirements that should be provided to your backup server administrator:

- The backup server should be running either NFSv2 over UDP or have an active `rsync` daemon.
- On NFSv2 servers, the entry for the NFS share receiving the backups must enable the `no_root_squash` parameter.
- A directory should be created to contain Controller backups. Multiple target directories can then be created within that root backup directory if you want to retain multiple backups of the Controller.

A remote user must be configured with write permissions to this directory and its contents.

- When using `rsync`, the backup process will create the target directory inside of a top-level directory if it exists. So, if `/backups` exists, and `/backups/controller-incremental` is specified as the target directory, it will be created in `/backups` if it does not exist.

When using NFS, the target directory is not created: it must already be present on the backup server.

- If you want to perform incremental backups, use `rsync` and keep the same target directory specified from backup to backup.

`rsync` will only synchronize changed files, so after an initial backup, subsequent backups to the same target directory will be less costly and only transfer changes that have happened since the last backup.

The target directory that contains a Controller backup on the backup server will contain multiple objects and a nested directory structure. Do not modify any of these contents: if they are moved or changed then restoration will likely fail or will lose its integrity.

3.3. Log File Cleanup

Over time the log files generated by the Controller may consume excessive disk space. The **Cleanup** menu allows you to delete legacy log files and other objects that are no longer in use by the system.

3.3.1. Deleting Controller Log Files

1. In the Administration Console, select **Appliance** → **Cleanup**.
2. Click **Logs** on the right to show a list of Controller administrative and system log files.
3. Select the files you wish to delete.



For any particular class of log file, the current file has a `.log` extension. Rolled-over logs are enumerated, `.0` being the most-recent and successive numbers representing older backups.

4. Click **Delete Selected Files** to permanently delete the selected log files.



User data (e.g. collected Audits, host information, analysis data) is never deleted from the system by this administrative command.

3.4. Controller Configuration

The **Config** menu allows you to configure logging, network, and NTP/time settings; and controls the `SSHD` maintenance service on the Controller. It provides the following options:

AgentConfigFile

Generates valid configuration files for Agent installation.

Logging

Controls log rotation and syslog settings.

NTP

Configures Network Time Protocol.

Network

Configures network settings for the Controller Appliance.

SSHD

Starts or stops the MANDIANT Customer Support SSHD process.

Time

Configures time and date settings.

3.4.1. Generating Agent `conf.xml` Settings

This option is discussed in *Section 7.2.1, "The Agent Installer"*.

3.4.2. Controller Log Files

Log retention and network log services are configured through **Appliance** → **Config** → **Logging**. For more information on logs see *Section 4.1.2, "Viewing Current Log Files"*.

You can customize the rotation and retention settings for logs on the Controller. Note that log retention uses disk space and may become problematic if not properly configured.

3.4.2.1. Configuring MIR Local Log Rotation

1. In the Administration Console, select **Appliance** → **Config** → **Logging**.
2. Using the controls in the **MIR Log Rotation** section on the right, configure the following settings, showing (defaults):

Log Rotation Cycle (weekly)

Controls whether logs are rotated daily, weekly, or monthly.

Max Log Size (MB) (1)

Sets the maximum size for a log file before a rotation is forced, independent of the rotation schedule.

Number to Keep (7)

Sets the maximum number of rotated log files to retain; as a new log is rotated into archival format, the oldest is removed.

Compress Logs (Yes)

Controls whether logs remain plaintext, or are compressed using gzip.

3. Click **Update Settings** to record your changes.

See *Section 3.3, “Log File Cleanup”* for details on deleting old log files.

3.4.2.2. Configuring Syslog NG Logging

The Controller can export its logs via network connection using Syslog NG ¹.

Assuming you have a correctly configured log host that will accept logging traffic from the Controller:

¹See <http://www.balabit.com/network-security/syslog-ng>.

1. In the Administration Console, select **Appliance** → **Config** → **Logging**.
2. In the **Syslog NG All Logging** section, specify whether to use the **UDP** or **TCP** protocol for log export.
3. Add **Host** and **Port** information for the logging Host.
4. Click **Add New Host**.
5. The new log host will now appear in the **Available Logging Hosts** list. Highlight the host. Select the host from the list and click the **Assign Host** button.
6. The log host should now appear in the **Assigned Logging Hosts** list. If you wish to remove it, use the **Unassign Host** button.
7. Click **Update Settings** to record your changes.

When syslog-ng is configured, all logs will be exported from the system, including application logs and appliance-level event logs. See *Appendix B, Error Messages and Troubleshooting* for more information about available log files.

3.4.2.3. Configuring Common Event Format (CEF) Logging

The Controller can provide CEF-compliant ² log information to servers:

1. In the Administration Console, select **Appliance** → **Config** → **Logging**.
2. In the **Syslog NEG CEF Only** section, specify whether to use the **UDP** or **TCP** protocol for log export.
3. Add **Host** and **Port** information for the logging Host.
4. Click **Add New Host**.
5. The new log host will now appear in the **Available Logging Hosts** list. Highlight the host. Select the host from the list and click the **Assign Host** button.
6. The log host should now appear in the **Assigned Logging Hosts** list. If you wish to remove it, use the **Unassign Host** button.
7. Click **Update Settings** to record your changes.

See the Arcsight section of the *User Guide* for additional usage details.

3.4.3. Configuring Network Time Protocol

To ensure the Controller persistently syncs its clock to a remote time source, you need to configure Network Time Protocol (NTP) services:

²See <http://www.arcsight.com/collateral/CEFstandards.pdf>.

1. In the Administration Console, select **Appliance** → **Config** → **NTP**.

NTP Configuration
Configure appliance NTP settings.

Single IP NTP Servers

NTP Server #0:

NTP Server #1:

NTP Server #2:

NTP Server #3:

Multiple IP NTP Servers

NTP Servers #0:

2. By default, the Controller is set to synchronize with servers maintained by the *NTP Pool Project*. If your Controller can connect to resources on the internet, these settings may be enough.

Otherwise, remove those entries and configure the appliance to sync to your preferred time services.

3. When your configuration is complete click **Update NTP Configuration**.

You may also set the time manually; see *Section 3.4.6, “Setting the Time and Date Manually”*.

3.4.4. Network Settings

By default, the Controller will try to configure its primary interface (`eth0`) using DHCP, and will leave the `eth1` interface deactivated. You can use either DHCP or set the configuration for each interface manually.

If you configure both interfaces, it is strongly recommended to put each interface on a separate subnet so the Controller can benefit by dividing its bandwidth consumption across multiple networks. Placing one interface on a network primarily facing Console users and another interface on a network primarily facing Agents will help to split the data acquisition and end-user traffic loads, increasing overall performance.



This version of MIR does not allow you to specify the route that a particular type of traffic will take. Traffic will be dispatched according to standard routing rules. The default route for the Controller when it is first configured is on the `eth0` interface.

3.4.4.1. Configuring the Controller with Static Addresses

1. . In the Administration Console, select **Appliance** → **Config** → **Network**.

2. Provide a **Hostname** for the Controller.
3. For **Enable eth0**, select **Yes**.
4. For **Use DHCP for eth0?**, select **No**. This allows you to use manual configuration options for the eth0 interface.

Network Configuration

Configure appliance network interfaces. Changing the network configuration will cause the network interfaces to be restarted, along with all MIR Controller services. **WARNING: an invalid network configuration will make the appliance and this administration interface unreachable.** Carefully validate settings before proceeding. Refer to the Administrator Guide for additional information.

* Hostname:

Enable eth0?: Yes No

Use DHCP for eth0?: Yes No

MAC Addr for eth0: 08:00:27:fc:16:d2

Domain Name for eth0:

* IP for eth0:

* Netmask for eth0:

Gateway for eth0:

Broadcast for eth0:

MTU for eth0:

DNS Server Search Order for eth0:

DNS Server 0 for eth0:

DNS Server 1 for eth0:

Enable eth1?: Yes No

5. Set the network configuration appropriate to your MIR environment.



Double-check your entries: an incorrect network configuration may render the Administration Console unavailable.

The fields are as follows:

Domain Name (Required)

The domain name suffix for the Controller. If the Controller is `controller.MANDIANT.com` this field would be `MANDIANT.com`.

IP (Required)

The IP address for the interface. Example: `192.168.0.3`.

Netmask (Required)

The netmask for the interface, in dotted-quad format. Example: `255.255.255.0`.

Gateway (Required)

The IP address of the default gateway for the interface. Example: `192.168.0.1`.

Broadcast (Optional)

The broadcast address for the interface. If not specified, it defaults to setting all bits within the host portion of the network mask to all 1s. Example: `192.168.0.255`.

MTU (Optional)

The MTU size in bytes for the interface. If not specified, defaults to 1500.

DNS Server Search Order (Optional)

A space-delimited list of up to six domain names (256 characters) to be used when a hostname but no domain name is provided for lookup. Example: `company.com MANDIANT.com`.

DNS Server 0 (Optional)

The IP address of a DNS server for this interface. If one is not specified the Controller will not be able to resolve any host names. Example: `192.168.0.2`.

DNS Server 1 (Optional)

The IP address of a DNS server for this interface. Example: `192.168.2.2`.

6. Click **Update Network Configuration** when you have completed configuring the interface.

The Administration Console will display a message showing that network changes are being applied. If the IP address has been changed from its prior setting, an error will be displayed when the Administration Console page tries to refresh (typically within four minutes). In that instance you may manually point your browser to the address you set for the Controller as follows:

```
https://[new ip address]/administration/
```

If you wish to configure `eth1` as well, repeat this task; the same configuration options will be presented for its interface.

3.4.4.2. Configuring the Controller with DHCP-Persistent Leases

1. First, obtain the hardware address for the `eth0` interface:
 - a. In the Administration Console, select **Appliance** → **Statistics** → **Network Stats** in the left navigation pane.
 - b. Select **Interfaces** on the right. Details for each active interface will be displayed.

In the **Interfaces** information, find and note the **HWaddr** value for the `eth0` interface; if `eth1` has also been activated, make note of its value as well.

```
Network Status
Interfaces
eth0  Link encap:Ethernet HWaddr 08:00:27: : :d2
      inet addr:192.168.56.103 Bcast:192.168.56.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:1311 errors:1 dropped:0 overruns:0 frame:0
      TX packets:1339 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:329008 (321.2 KiB) TX bytes:781174 (762.8 KiB)
      Interrupt:10 Base address:0xd020
lo    Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      UP LOOPBACK RUNNING MTU:65536 Metric:1
```

2. Next, configure your network DHCP server to reserve a lease for the `eth0` (and, optionally, `eth1`) interface:
 - a. In the Administration Console, select **Appliance** → **Config** → **Network**.
 - b. Provide a **Hostname** for the Controller.
 - c. For **Enable eth0**, select **Yes**.
 - d. For **Use DHCP for eth0?**, select **Yes**. The Controller will now obtain its IP address from the DHCP server.
3. Click **Update Network Configuration** when you have completed configuring the interface.

If you wish to configure `eth1` as well, enable it similarly.

3.4.5. Disabling the SSHD Maintenance Service

The Controller appliance installs and enables a Secure Shell daemon (`SSHD`). It allows MANDIANT to help you with remote maintenance in the event of a support emergency. The `SSHD` service is configured to allow key-only authentication and MANDIANT Customer Support has the only keypair authorized to authenticate to the running service.



For customer support engineers to access your Controller you will need to authorize and enable a remote access service for MANDIANT personnel.

Under no circumstances does the service connect back to MANDIANT without your knowledge and approval.

- To disable the `SSHD` service, select **Appliance** → **Config** → **SSHD**, click **Stop SSHd** and then reboot the Controller (**Appliance** → **Reboot**).

3.4.6. Setting the Time and Date Manually

You can manually set the time on the Controller or do a “one shot sync” of the Controller to a remote time source using `ntpdate` by selecting **Appliance** → **Config** → **Time**.

Date/Time Configuration
Configure appliance Date/Time settings.

Manual Date/Time Entry

Current Date/Time: Jul / 4 / 2012 5 : 58

Manually entered time must be in UTC/GMT/Zulu.

Force Remote Date/Time Synchronization

Forces the appliance system and hardware clocks to immediately match the time provided by the currently-configured NTP servers. May be needed if the appliance time differs significantly from the network time.

For manual setting, set **Current Date/Time** and click **Update Appliance Date/Time**.

For a one-shot remote sync, click **Sync Appliance Date/Time**. This performs an immediate synchronization of the Controller with the NTP servers configured via **Appliance** → **Config** → **NTP** (see *Section 2.4.5, “Configuring Network Time Protocol”*).



Because of the Controller’s sensitivity to time settings we strongly recommend configuring and using NTP services. See *Section 2.4.5, “Configuring Network Time Protocol”*.

3.5. Controller Diagnostics

The **Diagnostics** menu provides access to troubleshooting tools that can be useful for diagnosing configuration or performance problems. This version of MIR provides two basic network utilities for diagnosing network configuration problems: `ping` and `traceroute`.

3.5.1. How to Use Ping

Select **Appliance** → **Diagnostics** → **Ping**. Enter an IP address or hostname and click **Run Ping**. The command behaves identically to the Unix `ping` command

3.5.2. How to Use Traceroute

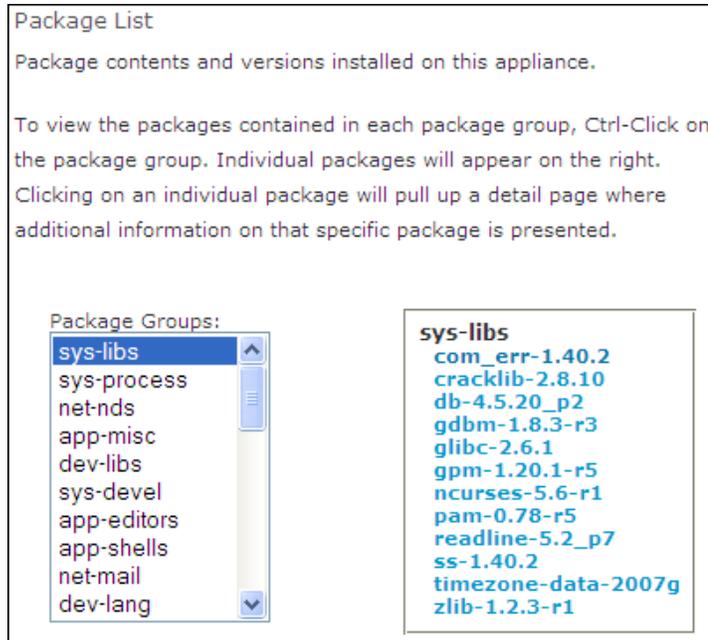
Select **Appliance** → **Diagnostics** → **Traceroute**. Enter an IP address or hostname and click **Run Traceroute**. The command behaves identically to the Unix `traceroute` command.

3.5.3. System Report

Please see the Appendices, *Error Messages and Troubleshooting*, for details on using the MIR System Evaluation Report command.

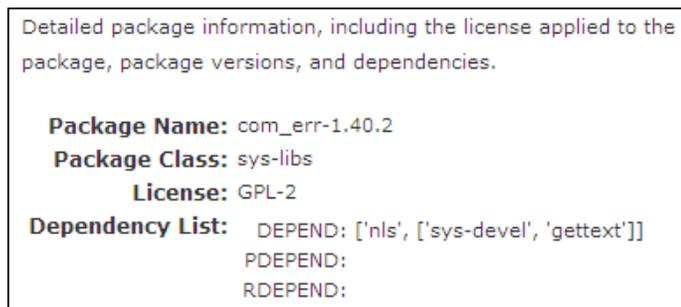
3.6. Troubleshooting Packages

The **Packages** menu provides information about software packages installed on the Controller appliance. The main page displays a scrolling list of package groups. Selecting a package group will provide a detailed list of software packages associated with that group.



Selecting a package from the detailed list on the right will display more specifics about the package itself.

In day-to-day Controller operations you typically will not need the information provided here. However, if you run into a support issue a MANDIANT support engineer may request certain details provided under this menu.



3.7. Field Patch Tracker

On occasion, MANDIANT releases field patches for the Controller software. These patch files provide minor bug fixes and functionality improvements between formal Controller Upgrade releases. As of the Controller v1.4.4 release, the Administration console can now list the field patches that have been applied to the Controller, and includes an integrity-checking feature that ensures your security remains uncompromised.

3.7.1. Viewing the Controller Field Patches List

1. In the Administration Console, select **Appliance** → **Patches**.
2. On the right, a list of patches will be displayed with the following information:

Name

The name of the patch file.

Version

The Controller version against which the patch was applied.

Status

The result of the patch integrity check (**Broken** or **OK**).

3.8. Rebooting the Controller

The **Appliance** → **Reboot** choice provides two options:

Reboot MIR Appliance

Reboots the Controller hardware. This process will take several minutes as the Controller safely shuts down, issues a hardware reboot command, and starts its processes again.

Shutdown MIR Appliance

Powers-off the Controller hardware. You will need to physically power-on the Appliance to restart the MIR Controller.



Rebooting or shutting down the system will abort current data collection and indexing operations. Data will not be corrupted, but search indexes may be incomplete, and data collection and analysis jobs will not contain complete results.

It is strongly recommended to have end users quit their Consoles before rebooting or shutting down the system.

3.9. Statistics

This menu provides access to information about the running Controller environment. The following options are available:

Network Stats

Provides full network interface information and network connection status for the Controller. Located in the menu side bar.

Processes Stats

Details currently running processes and provides `top` output for the Controller. Located in the menu side bar.

Controller Version

Controller software version.

Disk Statistics

Disk capacity information.

System Statistics

Uname, uptime, and load average data.

3.10. The Appliance Upgrade Menu

Upgrading the Controller is a simple process. MANDIANT will provide software updates to you as either a downloadable ISO image or on DVD media. Either can be used to upgrade a Controller with minimal effort.



It is strongly recommended that you advise users to not use the system during the upgrade process.

Additional details and requirements are provided with the upgrade software package. Always be sure to review the release notes and follow the advice of MANDIANT Customer Support.

3.10.1. Upgrading the Controller using DVD Media

1. Open the Controller face plate and press the DVD drive eject button. Place the DVD in the tray and close the drive.
2. Perform all prerequisite tasks as stated in the release notes that accompanied the software upgrade package.
3. In the Administration Console, select **Appliance** → **Upgrade**.

On the main page, select **Local DVD**.

4. Test the upgrade before committing to it: click **Test Upgrade**. Any errors encountered will be displayed on the page. The running system is not changed during the upgrade test.
5. If **Test Upgrade** reports success, continue the upgrade: click **Start Upgrade**.
6. When the upgrade process is finished, perform all post-install tasks as stated in the release notes that accompanied the software upgrade package.
7. Eject the DVD and remove it from the Controller. Close the drive, and close the Controller face plate.

3.10.2. Upgrading the Controller using an ISO Image

1. Place the downloaded ISO file on the same system you are using to access the Administration Console.
2. Perform all prerequisite tasks as stated in the release notes that accompanied the software upgrade package.
3. In the Administration Console, select **Appliance** → **Upgrade**.

On the main page, select **File Upload**.

4. Click the **Browse** button. Locate the downloaded ISO file on the workstation. Click **Open** to send the file to the Controller. When the file has been completely transferred, continue the installation process.
5. Test the upgrade before committing to it: click **Test Upgrade**. Any errors encountered will be displayed on the page. The running system is not changed during the upgrade test.

6. If **Test Upgrade** reports success, continue the upgrade: click **Start Upgrade**.
7. When the upgrade process is finished, perform all post-install tasks as stated in the release notes that accompanied the software upgrade package.

You may delete the ISO file when the upgrade is complete.



Contact MANDIANT Customer Support for more information about the product roadmap, software release schedule, and support downloads site.



Chapter 4

Application Administration

Chapter 3, Appliance Administration provided an in-depth look at appliance-level administration functions. This chapter discusses maintenance and operation of the MIR application that resides on the appliance, including application health monitoring, and configuration options.

Configuration options for MIR can be accessed in the Administration Console by selecting Application in the Navigation Bar.

MANDIANT Intelligent Response Administration Console	
View and update MIR Controller configurations.	
Discovery Status	Green
Database Status	Green
MIR Status	Green
TDCA Configured	Green

As with all functions available through the Administration Console, only a user with administrator rights can access these functions. The default admin account is always able to use this interface.

Various configuration options are available:

Components

Allows you to restart MIR processes, view their current status, and view and download application log files.

ConfigFiles

Displays statuses and allows you to alter various application configuration parameters.

Database

Allows you to reset all MIR databases, erasing all information stored on the Controller.

Discovery

Manages the Agent *Discovery* Service, which allows Agents to be automatically discovered on a network and identified to the Controller; also provides Agent upgrades over the network. This topic is discussed in more detail in *Chapter 7, Agent Deployment*.

MCIC

Allows you to reset MCIC.

RemoteAuth

Configures Active Directory remote authentication. MANDIANT strongly recommends NOT using Active Directory due to security concerns.

Resources

Allows you to create Scheduling Queues, Hosts, or Documents directly on the Controller.

SSL

Manages the Trust Domain and associated SSL certificates for the Controller. This topic is discussed in *Chapter 6, Understanding the Trust Domain*.

4.1. Components

The Components menu allows you to view and download application logs, view application process status, and restart all application processes. There are two menu options:

ComponentLogs

Allows you to view various MIR process logs.

Processes

Lists the status of all MIR processes.

4.1.1. Restarting MIR Processes

The main page of the Components menu contains a button labeled **Restart All MIR Processes**. Clicking this button will restart all MIR application processes.

Note that the appliance itself will not reboot or power cycle: only the application-level processes will restart. Users connected to the Controller at the time of the restart will experience a service disruption (typically in the form of timeout errors displayed by the Console), but can continue working after the processes come back online.

4.1.2. Viewing Current Log Files

Logs for each application service listed above can be viewed and displayed via the **Application** → **Components** → **ComponentLogs** menu. Select the desired application log from the drop-down and click **Display Log** or **Download Log** depending on your preference. For information regarding log rotation and size settings see *Section 3.4.2, “Controller Log Files”*.



If log configuration settings are set to large log sizes (10+ megabytes) your browser may have difficulties directly viewing the log contents.

4.1.3. Controller Process Status

You can view the current status of MIR application processes by choosing **Application** → **Components** → **Processes**.

Component	Process Status	Availability	# Running
mir_lcd	stopped	dead	1
mir_web_admin	started	alive	1
mir_web	started	alive	1
mir_data	started	alive	1
mir_discovery_server	started	alive	1
mir_mbus	started	alive	1
mir_discovery_proxy	started	alive	1

On a properly running Controller all processes should report **started** under **Process Status**, **alive** for **Availability**, and **1** (one) under **# Running**.

The following details the function of each MIR application process:

mir_agent

Logs *Discovery* data to the audit files.

mir_agent_controller

Controls all interaction with deployed Agents.

mir_agent_dispatcher

Controls dispatch of tasks to subsystems responsible for interacting with deployed Agents.

mir_agent_upgrade_proxy

Controls the Agent over-the-network (*Discovery*) upgrade process.

mir_analyzer_dispatcher

Dispatches analysis jobs to associated components.

mir_analyzer_service

Conducts analysis tasks.

mir_audit

Manages the system activity audit log.

mir_data

The data management service responsible for organizing and storing information in associated databases and on disk.

mir_discovery_server

The Agent *Discovery* server.

mir_discovery_service

The Agent *Discovery* Service. Registers new Agents and updates their records as information (such as network address) changes.

mir_indexer

Indexes acquired data so it can be made available to *mir_searcher*.

mir_lcd

Controls the LCD panel on the front of the Controller (deprecated).

mir_mbus

The back-end message bus that transmits messages between MIR Controller components.

mir_pound

Web caching service for the Controller.

mir_restore_web

Performs Controller backup and restore tasks..

mir_scheduler

Schedules and executes tasks within the system.

mir_script_runner

Parses Job Scripts and controls dispatch to either the Analysis engine or Agent management subsystem.

mir_searcher

Controls MIR's search engine.

mir_searcher_dispatcher

Controls dispatch of tasks to subsystems responsible for resolving search requests.

mir_web

The primary web service that Consoles interact with, plus the log files for *mir_web_admin*, *mir_web_files*, and *mir_web_static*.

mir_web_admin

Provides the admin UI interface.

mir_web_files

Provides file transfer services.

mir_web_static

Serves static content for the *mir_web* service.

mir_discovery_prox

Proxy server for *Discovery*.

4.2. Appliance Configuration Details

Various application settings can be controlled through the **Application** → **ConfigFiles** menu. The main page displays a set of expanding configuration topics. To change a setting update its value in the associated input box and click **Update MIR Configurations** at the bottom of the page.

All MIR processes should be restarted after configurations are updated.

MIR Configuration

Set MIR configuration options for the Controller and underlying appliance. Changes made here will not be reflected in the running system until the system has been restarted [here](#). Refer to the Administrator Guide for more information.

MIR Configurations

- **debug_level:**
Level of debugging output verbosity. 0 is appropriate for production systems.

Agent Controller
Web Service

- **betweenRequestsTimeOut:**
Number of seconds server will wait before closing an idle connection.

Web Service Location
Admin UI Parameters

Authorization
Discovery Service
Scheduler Service
Search Indexer

The Controller ships with MANDIANT's recommended default settings. Most should only be changed under the guidance of Customer Support. Information about each setting is available directly in the Administration Console. The following list describes the settings which you can change with minimal risk to Controller operations. The default value is also provided.

Agent Dispatcher → limit_simultaneous_host (False)

If **True**, restricts more than one audit from executing on an Agent at any given time. Will only dispatch a single job to an Agent and wait for it to complete before giving that same Agent another job.

Agent Dispatcher → global_concurrent_agent_sessions (500)

Sets the maximum number of concurrent connections the Controller will allow to Agents. This is a global setting across all jobs executing on the Controller.

Agent Controller → Quality of Service → bandwidth units (kpbs)

Units associated with parameter for setting max bandwidth.

Agent Controller → Quality of Service → bandwidth (none)

Value, when combined with units, that represents max bandwidth the Controller should consume when retrieving information from Agents. Controller will throttle at the TCP layer the aggregate consumption across the entire system.

Agent Controller Connection → maniConnectTimeout (15)

Number of seconds to wait before giving up on initial connection to Agent. Set this lower to make connection attempts to Agents fail faster.

Agent Controller Connection → connectMaxDelay (120)

For connection retries, sets the maximum delay between retry attempts. Note the Agent Controller uses an increasing backoff interval when multiple attempts are being made to an Agent.

Agent Controller Connection → keepAliveNumProbes (6)

Number of keep-alive probes sent to Agent on an established connection where the Agent is not responding before giving up.

Agent Controller Connection → maniConnectMaxRetries (1)

Number of times to retry connections to Agents that don't respond to an initial connection attempt.

Agent Controller Connection → maxResumeAttempts (10)

Sets the number of times to attempt to resume an interrupted HTTP transaction with the Agent. Set this lower to make Agent connection attempts fail faster and exit.

Agent Controller Connection → maniConnectMaxDelay (120)

Maximum number of seconds for connection retry backoff for an initial connection attempt. Set this lower to make connection attempts to Agents fail faster.

Agent Controller Connection → keepAliveIdleSeconds (120)

Number of seconds of idle activity on a TCP connection before sending a keep alive probe.

Agent Controller Connection → protocolTimeout (86400)

Number of seconds to wait for HTTP data to flow on a connection. Note setting this value too low may result in apparent job failures if the Agent is processing a job that takes a long time to return data (e.g. MD5 sums for an entire hard drive, extensive memory forensic operations).

Agent Controller Connection → connectMaxRetries (1)

Maximum number of retries when contacting Agent after initial communication. Initial connection behavior are controlled by "mani"-prefixed configuration variables.

Agent Controller Connection → connectTimeout (15)

Number of seconds to wait before giving up on a connection attempt to the Agent after initial communication. Initial connection behavior are controlled by "mani"-prefixed configuration variables.

Web Service → Web Service Location → port (443)

Controls which TCP port the Controller listens on to service Console requests.

Web Service → Admin UI Parameters → num_dns_servers (2)

Controls how many DNS servers may be configured in the Controller's network settings. Changing this value allows more to be entered via the Administration Console.

Web Service → Admin UI Parameters → num_ntp_servers (4)

Controls how many NTP servers may be configured in the Controller's network settings. Changing this value allows more to be entered via the Administration Console.

Web Service → Admin UI Parameters → num_ntps_servers (1)

Controls how many NTPS servers may be configured in the Controller's network settings. Changing this value allows more to be entered via the Administration Console.

Remote Web Service → IP Address Filtering (0.0.0.0/0)

You can also restrict access to the remote web service by specifying IP addresses to specifically admit or reject. Users are not allowed to access the remote web service from an IP address that is either denied or not allowed by the admin configuration. This configuration is done in the admin UI. Administrators can express allow and deny rules as a comma separated list of CIDR-notation IP ranges. By default, no IPs are denied and 0.0.0.0/0 is allowed (i.e. all IP addresses are allowed).

Discovery Service → Discovery Server Location → port (8077)

Controls which TCP port the Agent *Discovery Service* runs on. Note it must be set to a different port than the Web service.

Script Runner → max_concurrent_agent_sessions (5)

Controls how many Agents the Controller will talk to simultaneously per job. You can use this to throttle Controller activity to avoid overloading your infrastructure. Note this is a "per job" variable, not a global variable. See `global_concurrent_agent_sessions` for a way to limit the maximum number of global agent connections.

Search Indexer → num_indexers (3)

Sets the number of processes used to index acquired files and metadata

Search Indexer → category_exclude_regex (blank)

A regular expression identifying categories of files to be skipped while indexing. This can significantly increase indexing speed. See the *MIR User Guide Appendix B, Searches*, for further detail.

User Audit Logging → enabled (False)

Controls logging of host audit actions, with details including user name, hosts audited, scripts used, and date/time logging. See for details.

Authorization → Password Policy → remote_password_length_min (8)

Minimum length for passwords of users in the *RemoteOnly* group. If users are added to the RemoteOnly group after they are created, the user password may need to be changed to comply with this restriction.

4.3. The Application Database Menu

This menu provide access to:

Database Maintenance

Provides commands for vacuuming and re-indexing the Controller database.

Database Reset

Deletes the Controller database.

Database Statistics

Provides an overview of the database used, dead, and free space.

4.3.1. Database Maintenance

Over time the Controller database may become slow and bulky. Periodic database maintenance will improve its speed and reduce its size by releasing old records, releasing reclaimed space, or re-indexing its contents.

4.3.1.1. Lightly Vacuuming the Database

When database rows are deleted, they are not automatically released for re-use by the database. Performing a minor vacuum operation will make these rows available. Though it will not reduce the size of the database, it will temporarily reduce the growth rate of the database, as well as provide some speed improvements.

1. In the Administration Console, select **Application** → **Database** → **Database Maintenance**.
2. On the right, click **Vacuum**.

While the database is being vacuumed, Agents and Console users may remain active.

4.3.1.2. Fully Vacuuming the Database

When database rows are deleted, they are not automatically released for re-use by the database. A full vacuum operation will remove these rows from the database, reducing the size of the database and improving Controller performance.

1. In the Administration Console, select **Application** → **Database** → **Database Maintenance**.
2. On the right, click **Full Vacuum**.



While the database is being vacuumed, Agents and Console users will not be able to connect to the Controller.

4.3.1.3. Re-indexing the Database

The database engine is sometimes unable to re-use its index pages when rows are deleted. This results in reduced performance and increased database size. Re-indexing the database will force the database to rebuild its indices, improving performance and reducing the number of index pages.

1. In the Administration Console, select **Application** → **Database** → **Database Maintenance**.
2. On the right, click **Reindex Database**.



While the database is being re-indexed, Agents and Console users will not be able to connect to the Controller.

4.3.1.4. Re-creating the Database Indexes

The database engine is sometimes unable to re-use its index pages when rows are deleted. This results in reduced performance and increased database size. Re-creating the database

indexes will force the database to start with a fresh index, restoring performance and reducing the number of index pages.

1. In the Administration Console, select **Application** → **Database** → **Database Maintenance**.
2. On the right, click **Re-Create Database**.



While the database is being re-created, Agents and Console users will not be able to connect to the Controller.

4.3.2. Resetting the Database



Resetting the Controller database is a destructive operation. All information, save for that stored with the Agent *Discovery* Service, will be deleted. Only users, groups, and permissions will be preserved; everything else will be reset to factory default.

To reset the Controller database, select **Application** → **Database** → **Database Reset**. On the right, click **Reset Database** with some amount of trepidation: if there was anything important in the database it will soon be unrecoverable.

4.3.3. Database Statistics

Before performing database maintenance, you should look at database statistics to determine whether a potentially lengthy vacuum or re-indexing process is necessary.

The database statistics report will list the following information:

Space on Disk

The amount of disk space used by the database. The size of the database is not necessarily the most important measure of its performance: dead and free database row space is more important.

Used

The number of database rows in use, as a percentage of the database size. When most of the database is in use, the database file size will increase to permit more rows to be stored.

Dead

The number of database rows that have been deleted or updated, as a percentage of database size. Deleted rows are not available for re-use until the database is vacuumed. Lightly vacuuming dead rows will not reduce the database size, but will make the dead space available for re-use.

Free

The number of database rows that are available for new data, as a percentage of database size. When there is very little free space left, it is often best to run a full vacuum; otherwise, the database size will be increased to permit more rows to be stored.

4.4. Discovery

Application → **Discovery** provides control of the Agent *Discovery* service and facets of remote Agent management. The following commands are available:

Agent Upgrade

Provides the ability to upgrade Agents over the network. Please see *Section 7.2.3, “Configuring Agent Upgrades via Discovery Services”* for details.

Bulk Load

Allows you to save and transfer the current state of the *Discovery* service to another *Discovery* instance on another Controller. See *Section 7.1.4, “Discovery Service Administration”* for details.

Discovery Query

Provides the ability to search for an Agent on the network. See *Section 7.1.4, “Discovery Service Administration”* for details.

Discovery Reset

Deletes the information contained in the *Discovery* database. See *Section 7.1.4, “Discovery Service Administration”* for details.

Discovery Search

Allows you to find Agents based on an IP address or range of addresses. See *Section 7.1.4, “Discovery Service Administration”* for details.

Duplicate Agent Certs

Detects duplicate Agents on the network. Again, see *Section 7.1.4, “Discovery Service Administration”* for details.

4.5. MCIC

The MCIC database can be reset (with the option of retaining a backup) by selecting **Application** → **MCIC**.

If MCIC has not been installed, selecting **Reset MCIC** will fail with a “script not found” error.

The screenshot shows the 'MCIC Reset' configuration page. On the left is a navigation menu with 'Application' selected. The main content area contains the following text and form elements:

MCIC Reset
Use the form below to reset data for the local MCIC installation. Note that if this MIR installation does not have MCIC installed the script will have no effect.

Path to the webclient database. Leave blank for default.

Path to store a backup of the webclient db. Leave blank to skip creating a backup.

Delete MCIC acquisition files data?

4.6. RemoteAuth: Active Directory Authorization



Configuring these settings will connect MIR to an external authentication provider. Although this may simplify account management, MANDIANT **strongly** recommends against enabling this capability.

In many cases, attackers target compromised directory infrastructure (particularly Active Directory controllers); if your MIR installation is connected to a compromised authentication provider, the attacker may gain access to all MIR capabilities, enabling

them to potentially disrupt incident response activities or use MIR's data gathering features for malicious purposes.

This feature was created to enable certain MIR customers with highly restrictive, mandatory authentication integration policies to make use of the product. MANDIANT advises all users to carefully consider the risks of sharing credentials between MIR (or any incident response infrastructure) and other applications and avoid it whenever possible.

The **Application** → **RemoteAuth** menu manages remote authentication services against the Controller. Options are:

Active Directory

Configures Microsoft Active Directory remote authentication.

Active Directory Test

Confirms that the settings configured in Active Directory are functional.

4.6.1. Active Directory

Configure Microsoft Active Directory by selecting **Application** → **RemoteAuth** → **Active Directory**.

New Realm

The authentication realm (MYCOMPANY.COM)

Is This The Default Realm

Controls whether the authentication realm is the default realm.

Admin Server

The Active Directory server (dc1.mycompany.com)

New KDC

Entries for Kerberos Key Distribution Centers (dc1.mycompany)

Default Domain

The default domain used to map AD realms to kerberos realms (mycompany.com)

New Domain

Additional domains for mapping AD realms to kerberos realms (.mycompany.com)

4.6.2. Active Directory Test

Select **Application** → **RemoteAuth** → **Active Directory Test** to confirm that the settings provided in **Active Directory** provide full functionality. When selected, the main page will display a table of tests and test results. When possible, failures are fully explained as an aid to troubleshooting. Click **Check Active Directory Configuration** to re-test.

4.7. Resources

The **Application** → **Resources** menu allows you to view, create, and modify certain data objects within the Controller. The most important aspect of this is the management of

Scheduling Queues. However, it can also be useful for troubleshooting if a support issue arises. You can directly manage:

Documents

Search for existing documents on the Controller, or upload a new document to the Controller.

Queues

Create Job Queues.



These utilities should be used with caution. Documents should only be directly modified if you explicitly understand what you are doing. Contact Customer Support for assistance.

4.7.1. Documents

Select **Application** → **Resources** → **Documents** to view the document management interface. It has two sections: **Create New Document** and **Document List**. Each section can be hidden or shown using the *Hide/Show* links above the sections.

The **Create New Document** section allows you to create new documents.

Fill in the appropriate fields to create the Document and select content from you local system to inject into the Document by **Browse**. Set the **Indexing State**, and pick a **Source Host** and **Source Result Set** with which to associate the Document.

Again, this should only be used if you have complete understanding of the modification you are making. Contact Customer Support for assistance.

Create New Document

Name:

Content Type:

Document Content:

Indexing State:

Source Host:

Source Result Set:

Document List shows a list of documents in the system matching filter criteria supplied by the administrator. Resources can be filtered by name or Content Type. The filters themselves can be set to a complete match (e.g. "="), or a partial match (e.g. "like" or "startswith").

Document List

Filter/Search For Documents:

Filter By Resource Name:

Filter By Content Type:

= like startswith =

SORT ON A COLUMN BY SELECTING THE COLUMN NAME

	Name	Source Host	Indexing State	
Edit Delete	mir.w32volumes.551b3223.xml	Source Host	indexed	application/vn
Edit Delete	mir.w32volumes.3808447c.xml	Source Host	indexed	application/vn

2 Records Matching Your Filter Criteria

The **Document List** shows a grid of matching documents and a series of fields describing the Document:

Edit/Delete

Click the appropriate link to edit or delete the selected document.

Name

The name of the Document.

Source Host

A link to the Host object the Document was collected from.

Indexing State

The search engine indexing state for the document. One of **indexing**, **indexed**, or **not indexed**.

Content Type

The type of Document, e.g. `application/vnd.MANDIANT.mir.w32volumes+xml` (a w32volumes audit Document)

Document Content

A link to the contents of the Document. This will directly display XML in your web browser if your browser supports raw XML display without a style sheet (note: Firefox works particularly well for this).

4.7.2. Queues

Scheduling Queues describe a schedule that Jobs associated with that Queue use to execute. A Queue might describe a time-based recurring schedule, like “Every Hour on the Hour”, or “Every Tuesday at Noon.” These Queues, once defined, show up under the Scheduling tab in the Job *Viewer/Editor* in the Console. A Queue must first be defined in the Administration Console before it can be used in the Console.

There are two types of Queues: *Event Queues*, and *Non-Event Queues*:

- Event Queues launch jobs when a specific event happens within the MIR system. As of this release of MIR, only the **When a new Host is Discovered Event** is defined. It is triggered every time a new Host record is created in the system by the *Agent Discovery Service*. Jobs scheduled to this Queue are automatically run against new Hosts as they are added by the service.
- Non-Event Queues launch jobs based on a schedule as defined on the Queue. “Every Tuesday at Noon”, “Run Immediately”, and “Every Hour on the Hour” are examples of Non-Event Queues.

Click **Application** → **Resources** → **Queues** to see the management interface. It has two sections: **Create New Queue** and **Queue List**. Each section can be hidden or shown using the *Hide/Show* links above the sections.

Create New Queue displays the queue creation tool in the main page area.

If you select **Yes** for **Event Queue** you will be asked to select an event to associate with the Queue. The only event supported in this version of MIR is **When a new Host is discovered**.

If you select **No** for **Event Queue**, you will be presented with schedule fields that you can set for the Queue. Event schedules are set using cron-style syntax. Allowable field values are as follows:

Minute

0–59

Hour

0–23

Day of Month

1–31

Month

1–12

Weekday

0–7 (0 and 7 are Sunday)

Nth Weekday of Month

1–5 (e.g. “1st Tuesday of Month” is coded as Weekday=2, Nth=1.)

Additional Field Values:

- An asterisk (“*”) specifies “first to last” for any given field.
- Number ranges are permitted. (e.g 7–10 would mean 7, 8, 9, 10).
- Comma separated list both with and without ranges are permitted (e.g., “1, 2, 3, 5–10”).
- A “step value” may be specified with a range using a trailing slash and skip value (e.g. “/2”). For example, specifying “0–12/2” in the Hours field would specify the queue should fire every other hour between midnight and noon. Steps may also be specified after an asterisk.

Names can not be used for Month and Weekday fields.

Once you have completed defining your Queue schedule or event, click **Create Queue**.

Create New Queue

Name:

Event Queue?: Yes No

Minute:

Hour:

Day Of Month:

Month:

Weekday:

Nth Weekday of Month:

Queue List displays Queues in the system that match specified criteria. You can filter by the event associated with a Queue or the name of the Queue.

The **Queue List** shows a grid of matching Queues and a series of fields describing the Queue.

Queue List

Filter/Search For Queues:

Filter By Event:

Filter By Resource Name:

= like

SORT ON A COLUMN BY SELECTING THE COLUMN NAME

	Name	Soonest Possible Run Time	Event
Edit Delete	Run Immediately	2008-05-18 16:01:00	
Edit Delete	Every Hour on the Hour	2008-05-18 17:00:00	
Edit Delete	Every Hour on the Half Hour	2008-05-18 16:30:00	

A series of Queues are defined on the Controller by default, including the **Run Immediately Queue**. The **Run Immediately Queue** should not be modified: it is used extensively by the Console to schedule Jobs for immediate execution. All others can be modified or deleted to meet your needs.

Edit/Delete

Select these links to Edit or Delete a Queue.

Name

Name of the Queue. This is displayed to users through the Console. Clicking on the link will display the raw XML for the Queue.

Soonest Possible Run Time

The next time the Queue will “fire,” launching Jobs assigned to it.

Event

If the Queue is an event queue, specifies the event that will fire the Queue.

4.8. SSL

MANDIANT Intelligent Response uses a Public Key Infrastructure (PKI) to authenticate and encrypt transactions across network connections between its various components. The SSL menu provides commands for managing all aspects of the Controller keys and certificates. The following actions are available:CA:: Configures SSL key length and lifespans.

CRL

Generates the CRL and revokes certificates.

Certs

Manages the TDCA certificate and Controller Appliance certificate.

DisplayCRL

Displays raw CRL data.

DisplayCert

Displays raw certificate data.

Export Keys

Performs a backup of the TDCA certificate and Controller Appliance certificate. The backup is protected by PKCS12 encryption.

Import Keys

Restores a backed-up TDCA certificate or Controller Appliance certificate.

Keys

Allows deletion of the TDCA certificate and key, and Controller Appliance certificate and key.

Signing

Signs or deletes CSRs, including the Controller Appliance CSR.

Please refer to the *Chapter 6, Understanding the Trust Domain* chapter for a thorough discussion of the use of PKI, Trust Domains, certificates, and keys.



Chapter 5

MIR Users and Groups

Users may be administered via the Administration Console. Administrators can create and delete users, assign users to pre-defined groups, add and revoke administrator privileges for the appliance, and manage user account credentials. These features are available through the Users menu in the Navigation Bar.

5.1. Managing Users

To deploy the MIR system you must create and manage user and administrator accounts on the Controller. The Controller is pre-configured with one administrative-level user (admin) and one default group that provides rights for accessing the Administration Console (Administration).

This version of MIR provides two different levels of access: administrator and user. Administrator access is controlled by membership in the Administration group: all users in the Administration group have full administrator rights; all other users have user-only rights. All user-level access is identical and provides read/write permissions to all data stored within the Controller.

All user management functions are available under the **User** menu in the Navigation Bar.



All users and groups are case sensitive. User “Dmerkel” and “dmerkel” are different users.

5.1.1. Creating User Accounts

1. Log into the Administration Console using this URI:

```
https://[Controller URI or Hostname]/administration/
```

2. Select **Users** → **Create User**.
3. On the right, you will be presented with a user administration form.

The screenshot shows a web interface for 'User Administration'. At the top, it says 'Create and manage user accounts.' Below this is a section titled 'Create New User' containing several input fields: 'User Name:', 'Active Directory Realm:' (with a radio button selected for 'Local User'), 'First Name:', 'Last Name:', 'New Password:', and 'Confirm New Password:'. Below the 'Create New User' section is a 'Groups' section. It features two lists: 'Available Groups' (containing Administration, ReadOnly, SearchOnly, and RemoteOnly) and 'Assigned Groups' (containing Users). Between these lists are navigation buttons: '>>>', '>', '<', and '<<<'. At the bottom of the 'Groups' section is a 'Create User' button.

4. Provide information for the new account. Note that **User Name** is case-sensitive, and must be a unique identifier.
5. Assign group memberships. By default, all new users are assigned to the **Users** group; if the account is to be an administrator account, add the **Administration** group as well.

An account may also be **ReadOnly**, allowing a user to only read items stored on the Controller. They can not search, run Jobs, or modify any data. **SearchOnly** is similar, except the user can run searches.

A **RemoteOnly** account allows a user to access URLs through the remote web service, and denies all other types of access:

- Cannot access any other MIR resources including the Admin UI and whole REST API.
- Must have a password of minimal length (by default 8 characters, but configurable in the Admin UI).
- Cannot access the MCIC UI or any direct MCIC URL.
- Can only access select MCIC functions through the remote service:
- Current version information (`/sys/version/`).
- Create acquisition POST URL (`/acquisitions/create`)

See the *User Guide* for further information about this account type.

6. Click **Create User** to complete. The **Users** → **Current Users** selector will now include the new account.

5.1.2. Modifying or Deleting a User Account

1. Log into the Administration Console using this URI:

`https://[Controller URI or Hostname]/administration/`

2. Select **Users** → **Current Users** , then click the user account you wish to modify.
3. Click **Delete** if you want to remove the user account.

OR

Click **Edit** if you want to modify the user account. The account information will be displayed on the main page. Click **Update User Record** to commit the changes.

5.1.3. Setting and Resetting User Credentials

User passwords may be reset by modifying the user's record as described above. You can also reset a user password by selecting **Users** → **Password Change** from the Navigation Bar. However, you must know the user's existing password to use this method.

Users may reset their own passwords via the Console. See the *User Guide* for additional information.



The administrator account password can not be changed by using the Password Change option. Instead, from **Current Users** select **admin**, then click **Edit**.

5.2. Managing Groups

The Controller ships with five user groups by default: Administration, Users, SearchOnly, and RemoteOnly. In this version of MIR, only the Administration group provides any access management capability.

All Administration members have administrator access to the Controller. Other groups have the following permissions, in order of increasing capability:

RemoteOnly

Read-only access to specific MCIC URLs and functions. Other URLs and functionality, including the Admin UI, REST API, and MCIC UI, are blocked. Passwords have a minimum length requirement. No searches. No administration privileges. Can not change password for self. Can create an acquisition POST URL.

ReadOnly

Read-only access to the Controller through the Console. No searches. No administration privileges. May change password for self.

SearchOnly

Read-only access plus the ability to search against the Controller through the Console. No administration privileges. May change password for self.

Users

Full read/write access to the Controller through the Console. No administration privileges. May change password for self.

Administration

Full administration privileges. When accessing Controller through Console has ReadOnly privileges. May change password for self and others.

You can not modify these groups in this version of MIR.

5.3. User Audit Logging

User Audit Logging, enabled through **Application** → **ConfigFiles**, allows you to log Console auditing actions. The log files may be written to the Controller or to a remote syslog server. User Audit Logging is disabled by default to avoid impacting system performance.

5.3.1. Enabling and Disabling User Audit Logging

1. Log into the Administration Console using this URI:

```
https://[Controller URI or Hostname]/administration/
```

2. Select **Application** → **ConfigFiles** . In the main display, click **User Audit Logging**.

To enable User Audit Logging, set **enabled** to **true**.

To disable User Audit Logging, set **enabled** to **false**.

3. Click **Update MIR Configurations**.

4. Select **Application** → **Components** and click **Restart All MIR Processes**.

By default, User Audit Logs are recorded to the Controller. See *Section 5.3.4, "Configuring User Audit Remote Syslog"* if you wish to use a remote syslog server.

5.3.2. Viewing User Audit Logs

1. Log into the Administration Console using this URI:

```
https://[Controller URI or Hostname]/administration/
```

2. Select **Application** → **Components** → **ComponentLogs**. In the main display, select the **mir_user_audit** log and click **Display Log**.

Note that any event that modifies User Audit Logging settings are recorded regardless of whether logging is enabled. This insures you can determine whether a user has changed the settings.

5.3.3. Configuring the User Audit Log Format

As an aid to filtering syslog messages, you may customize the User Audit Log format string.

1. Log into the Administration Console using this URI:

```
https://[Controller URI or Hostname]/administration/
```

2. Select **Application** → **ConfigFiles**. In the main display, select **User Audit Logging**, then **Log Format Configuration**.
3. In the agent text entry, provide a new log template using freeform text and the following special variables:

%(response_queue)s

Where on the Controller the Agent response resides.

%(id)s

The identification number assigned to the task.

%(job_name)s

The name of the job that was run.

%(script)s

The full name of the script that was run against the host.

%(user)s

The user name of the person who ran the audit.

%(host_name)s

The name of the host system that was audited.

%(host_addr)s

The IP address and port on which the Agent was listening.

%(agent_version)s

The Agent's version number.

The default log format is:

```
user='%(user)s' agentname='%(host_name)s' agentaddr='%(host_addr)s' \
agentver='%(agent_version)s' job='%(job_name)s'
```

Which results in log entries like this:

```
Jun 30 06:19:50 MIRApp User-Audit: 9856dd8a957d11df983b000c29c28f6c
  user='mthomas' \
  agentname='mandiant-66a49d' agentaddr='https://192.168.52.136:22202'
  agentver='1.3.2300' \
  job='marks indicator job'
```

Note that the date, application, and user-audit UUID are automatically prefixed to the syslog message. When a syslog message is too long for the UDP protocol, it will be broken into smaller chunks; all chunks will have the same UUID, allowing you to re-assemble the message.

When using a syslog daemon which handles string-matching, you can filter the User Audit Log messages for "User-Audit:".

5.3.4. Configuring User Audit Remote Syslog

1. Log into the Administration Console using this URI:

`https://[Controller URI or Hostname]/administration/`

2. Select **Application** → **ConfigFiles**. In the main display, select **User Audit Logging**, then **Syslog Configuration**.
3. Configure the syslog facility and UDP settings as appropriate for your servers and network:
 - facility**
The syslog facility keyword.
 - remote_chunksize**
The maximum UDP message size for sending remote log messages.
 - local**
Enable (**true**) or disable (**false**) local logging.
 - local_chunksize**
The maximum message size for local log messages.
4. If you are using a remote syslog server, select Remote Syslog Configuration and configure appropriately:
 - host**
The IP address of the remote syslog server.
 - enabled**
Enable (**true**) or disable (**false**) remote logging.
5. Click **Update MIR Configurations**.
6. Select **Application** → **Components** and click **Restart All MIR Processes**.



Chapter 6

Understanding the Trust Domain

MANDIANT Intelligent Response uses a Public Key Infrastructure (PKI) to authenticate and encrypt transactions across network connections between its various components.

This PKI is self-contained: the Controller houses its own Certificate Authority (CA), and issues certificates to its various services, as well as a Certificate Revocation List (CRL) to manage revocation of access for components. The self contained CA and CRL in conjunction with any issued certificates are referred to as the Trust Domain.

This chapter details how the various transactions occur between MIR components to allow you to better understand problems that may arise and troubleshoot those situations. Fully describing the functioning of PKI is beyond the scope of this chapter. We will attempt to clarify terms and basic concepts; however, to more fully understand PKI we recommend consulting an external resource, such as *Understanding PKI: Concepts, Standards, and Deployment Considerations* by Carlisle Adams and Steve Lloyd (Addison-Wesley, 2003; ISBN-13 978-0672323911).

In addition to an overview of the Trust Domain, this chapter also provides more information about backup, restoration, and management of the Trust Domain Certificate Authority (TDCA) and associated server certificates.

6.1. The Trust Domain

MIR uses PKI technologies to create the necessary components to manage aspects of authentication, authorization, and encryption within the system. A Trust Domain (TD) is maintained by creation of a Certificate Authority (CA), which is subsequently used to issue credentials to participating entities (e.g. Controllers) and to validate that an entity is a member of a given Trust Domain. Generally speaking, a Trust Domain defines a set of entities that are authorized to interoperate with one another on some level.

6.1.1. Trust Domain and Subscribing Entities

The Trust Domain is defined by the collection of Trust Domain Entities (TDE) and currently-participating Subscribing Entities (SE).

Trust Domain Entities include Controllers and Agent *Discovery* Service providers; they are typically entities that perform some service that requires them to be strongly identified and authenticated to other entities.

Subscribing Entities are those that need to authenticate an entity within the Trust Domain, and may include the end user (through the Console) and Agents. Each TDE is also an SE for certain operations. For example, when a Controller is operating in a cluster (unsupported in this version of MIR) it needs to authenticate other Controllers within that cluster to see if those Controllers are authorized to communicate with it. When a Controller is authenticating another Controller, it is an SE of the Trust Domain. The table below provides an overview of the TDE and SE roles for components within MIR and for the end user.

Controller

Trust Domain Entity Role	Controllers must be authenticated when they connect to other Controllers, Agents, and Agent <i>Discovery Services</i> . Controllers authenticate one another when operating in a cluster. Agents must authenticate Controllers before accepting commands for acquisitions. Agent <i>Discovery Services</i> must authenticate a Controller before providing information about Agents or executing any other requested command.
Subscribing Entity Role	Controllers must be able to authenticate other Controllers when operating in a cluster.

Agent

Trust Domain Entity Role	None
Subscribing Entity Role	Agents must be able to authenticate Controllers and Agent <i>Discovery Services</i> . Controllers are authenticated before an Agent will accept commands to acquire data, etc. Agent <i>Discovery Service</i> is authenticated before the Agent will disclose any of its network configuration settings.

Console

Trust Domain Entity Role	None
Subscribing Entity Role	Consoles authenticate Controllers before permitting a user to enter their credentials (username and password, etc).

Agent *Discovery Service*

Trust Domain Entity Role	ADS authenticates itself to Agents before Agents will transmit their network configuration information.
Subscribing Entity Role	ADS must authenticate Controllers before accepting any commands, such as disclosing the network configuration information for an Agent.

End User

Trust Domain Entity Role	Users must authenticate themselves to a Controller before they can be provided access to that Controller or Controller cluster.
Subscribing Entity Role	Users must authenticate the Controller they connect to (primarily through the Console) before they input their credentials to gain access.

6.1.2. Trust Domain Operations

As previously mentioned, the Trust Domain is implemented through use of PKI technologies. The “trust anchor” for the domain is a Certificate Authority (CA) – an entity that brokers trust between the various components of the system. The CA issues credentials to TDEs so that SEs can validate them.

More concretely, the Trust Domain Certificate Authority (TDCA) issues certificates to Controllers and Agent *Discovery Services* services so that Agents can authenticate them, thus preventing unauthorized users from controlling Agents that you have deployed within your network.

All of this is accomplished through four administrative activities for a MIR deployment:

Creation of a TDCA

A TDCA keypair and certificate must be created for the Trust Domain.

Issuing Certificates to Controllers

Every Controller in the Trust Domain must receive a signed certificate from the TDCA in order to work with Agents in the Trust Domain.

Maintaining a CRL for the TDCA

Every Controller and every Agent in the Trust Domain must receive the CRL in order to know if any certificates (e.g. Controllers) have been removed from the Trust Domain.

Using the TDCA and CRL with Agents and Consoles

Agents and Consoles must receive a copy of the certificate for the TDCA. Agents must also receive a copy of the TDCACRL.

6.1.2.1. Creation of a TDCA

A Certificate Authority must be created to anchor the Trust Domain. The role of the CA is to issue identities and credentials to Trust Domain Entities (e.g. Controllers) by signing certificates for those TDEs. The *ITU X.509* standard is used for certificates, which specifies the format, signing algorithms and standards, and methods for validation of a certificate.

The CA is a component of one Controller within a Trust Domain and is created and maintained through the Administration Console (see *Section 2.5, “Configuring the Initial Trust Domain”*). If there are multiple Controllers within the same Trust Domain, one is designated the *Master* and provides the CA for the entire Trust Domain.

The TDCA is made up of a public/secret encryption keypair and an X.509 certificate that contains the public key and additional meta-data about the TDCA. The TDCA certificate can be used by any Subscribing Entity (e.g. an Agent or an end user) to verify whether a certificate presented by a Trust Domain Entity (e.g., a Controller) is valid; therefore, all Subscribing Entities must receive a copy of the TDCA certificate before they can begin to operate with components that are in the Trust Domain.

6.1.2.2. Issuing Certificates to Controllers

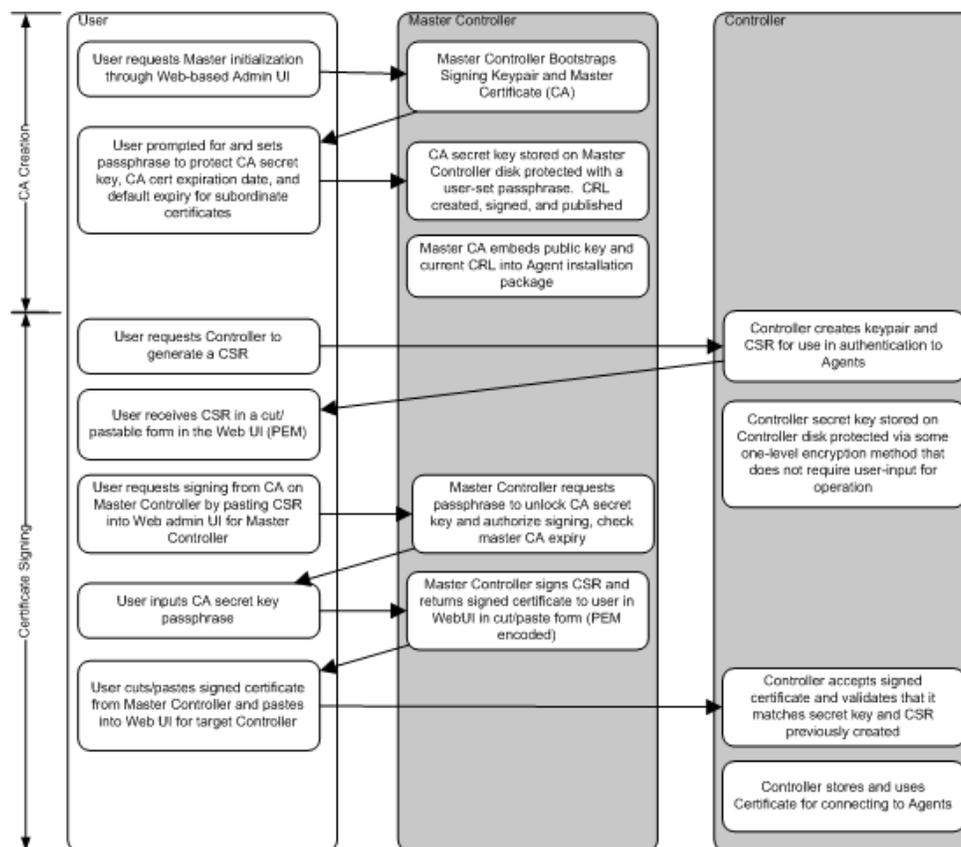
The primary operation of the TDCA is to sign certificates for TDEs, thereby making them members of the Trust Domain. This is accomplished according to the appropriate standards associated with use of X.509 certificates within a functioning PKI. In short, an entity that

wants to become a TDE (for example, a new Controller) must generate a public/secret key pair along with a Certificate Signing Request (CSR) and submit the CSR to the TDCA. The CSR must be digitally signed by the TDCA, using the secret key corresponding to the TDCA's public key (which is embedded in the TDCA certificate). Signing the CSR results in creation of a certificate; this is then passed from the TDCA back to the requesting entity (e.g. the Controller). Once the newly created certificate is installed in the requesting entity, it becomes a TDE and may now operate within the Trust Domain.

If you are only using a single Controller within your Trust Domain all of this is accomplished automatically for you when you create the TDCA and a corresponding SSL certificate for your Controller (see *Section 2.5.2, "Generating Keys and Certificates"*). If, however, you have multiple Controllers that you want to configure to talk to the same set of Agents within the same Trust Domain, then one Controller must be designated as the Master (e.g., the one configured with a TDCA), and other Controllers must have their certificates signed by it (see *Section 6.2.2, "Managing Server Certificates"*).

The overall process for creation of the TDCA, creation of a CSR, and signing a CSR to generate a Controller certificate is outlined graphically below.

Figure 6.1. Establishing the TDCA and Issuing a Certificate

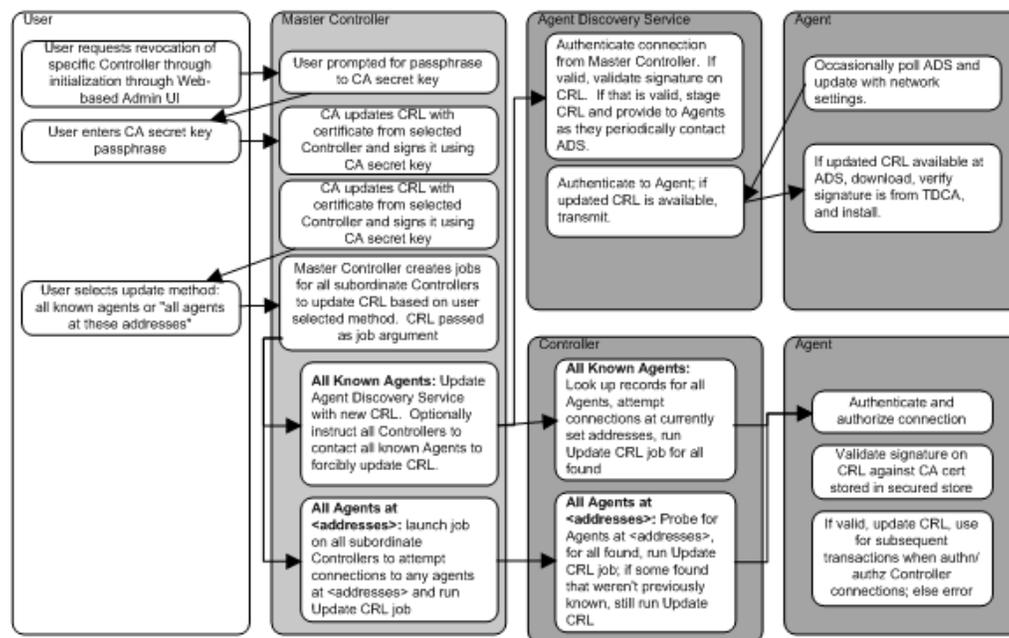


6.1.2.3. Maintaining the TDCA CRL

The TDCA can issue certificates to entities to make them part of the Trust Domain. However, it must also be able to revoke those certificates, allowing Subscribing Entities to identify when a given Trust Domain Entity should no longer be trusted.

For example, if a user had two Controllers within their Trust Domain and subsequently decided to remove one (e.g., sending it away for maintenance or decommissioning it), a method must be in place to ensure that Controller's certificate is no longer trusted by the Trust Domain. This is accomplished by publication of a Certificate Revocation List (CRL). The CRL is created and signed by the TDCA when an administrator wants to remove an entity from the Trust Domain. The CRL contains a list of all certificates previously issued by the TDCA that are no longer valid. Subscribing Entities can obtain a copy of this list and use it in conjunction with the TDCA certificate to validate whether a TDE has a valid, non-revoked certificate when it communicates with an SE. The diagram below outlines how this works.

Figure 6.2. Managing the Certificate Revocation List



 This version of MIR requires Agent *Discovery Service* in order to update CRLs for Agents: the Controller does not support running a “CRL Update” job manually against Agents at specific addresses.

6.1.2.4. Using the TDCA and CRL

Subscribing Entities must be able to authenticate members of a Trust Domain in order to interact with it. For example, Agents need to be able to authenticate a Controller before accepting commands from it to acquire data; and users need to be able to authenticate a Controller before typing their username and password into the Console to gain access.

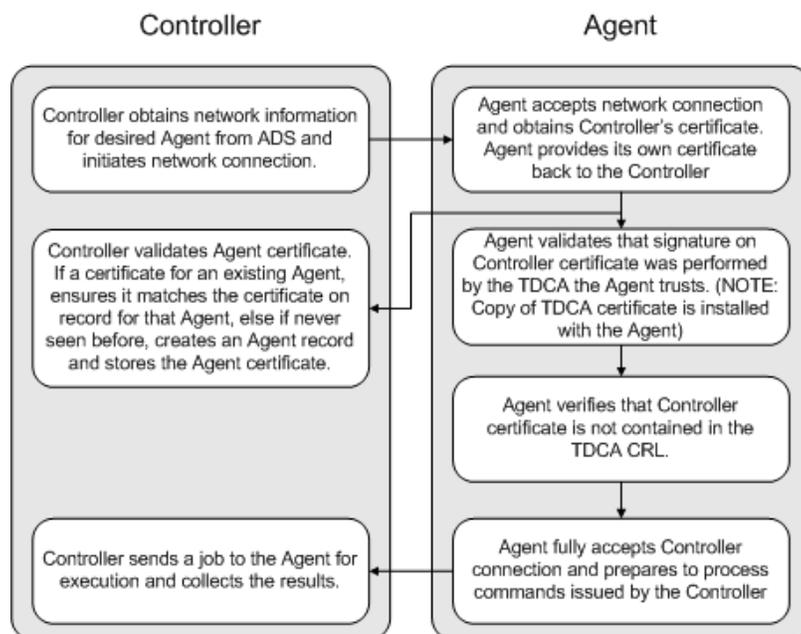
An SE authenticates a Trust Domain Entity by validating that the TDE's certificate was issued by the TDCA. To do this the SE must have a copy of the TDCA certificate. The TDCA certificate is **NOT** part of the installation package for the Agent and Console software.

When a TDE connects to an SE (or vice versa), the TDE provides a copy of its certificate to the SE and performs a challenge operation to validate that it also has the secret key that corresponds to the public key embedded in that certificate. The SE then validates the TDE certificate by verifying the digital signature embedded inside of it. It uses its copy of the TDCA certificate to perform this operation. It then checks the TDE's certificate against the

CRL. If the TDE certificate is not on the CRL, the connection is then authenticated and the SE can continue its communication with the TDE. All of this is accomplished using the Secure Sockets Layer (SSL) protocol.

The figure below illustrates an authentication exchange between a Controller and an Agent.

Figure 6.3. Controller/Agent Authentication



6.1.3. The Identity of Subscribing Entities

The sections above describe how the TDCA is set up, how Trust Domain Entities (primarily Controllers) are issued certificates, and how Subscribing Entities (primarily Agents) use them to authenticate the Controller.

However, SEs also need to have a concept of identity so that they can be correctly identified during a transaction. In other words, when a Controller connects and authenticates to an Agent, it needs a way to identify that Agent on an ongoing basis. This is accomplished by having the Agent create and maintain its own certificate for use in transactions with the Controller.

In this case the certificate is not issued by the TDCA: instead, the Agent creates a “self-signed” certificate when it is installed and activated on the target system. This certificate is then registered with the Controller either via Agent *Discovery Services*, or when the Controller first connects to the Agent.

6.1.3.1. Agent Enrollment

The process of registering an Agent’s certificate with a Controller is referred to as “enrollment.” The major components of that process are detailed below.

1. Upon installation, the Agent must have the TDCA certificate installed with the Agent software to ensure the Agent can authenticate TDEs. The Agent is also configured with

the network information necessary to contact the Agent Discovery Services for the Trust Domain.

2. At first startup the Agent creates a self-signed certificate; that is, a secret/public key pair with a corresponding certificate that encompasses the public key, whereby the certificate was digitally signed by its own private key. The public/secret key pair is 2048 bits in length, generated through use of software adhering to *Public Key Cryptography Standard #1*.
3. Two possible methods are now used to "enroll" the Agent. Enrollment is the process of registering an Agent's existence with the Controller and recording its certificate.

An Agent may be enrolled either via the Agent *Discovery Services* or through a direct connection from a Controller to the Agent:

a. Enrollment via Direct Connection from Controller to Agent

- i. When a Controller connects to an Agent, it transmits its own certificate and receives a copy of the Agent's self-signed certificate. These certificates are used to establish an SSL connection, which is then used for all subsequent communication between Controller and Agent for that session.
- ii. The Controller looks up the Agent certificate to see if an Agent record exists for it.

If the record does not exist, the Controller creates a record for the Agent and stores a copy of the Agent's certificate inside of it. Any data retrieved from the Agent (e.g. a process listing, files from the hard drive, etc) is associated with the newly-created Agent record. This process is referred to as enrollment.

If the Controller has seen the Agent certificate before, it identifies the Agent record for that certificate: in other words, the Agent has already enrolled. Any data retrieved from the Agent is associated with this pre-existing Agent record.

b. Enrollment via Agent Discovery Service

- i. If the Agent is configured to contact an Agent *Discovery Service*, the Agent initiates a connection and authenticates the Agent *Discovery Service* according to its certificate by using its local copy of the TDCA certificate and TDCA CRL.
- ii. The Agent *Discovery Service* records the Agent's certificate and network configuration settings (these are transmitted by the Agent to the Agent Discovery Service). When Agent *Discovery Service* is polled by a Controller for a list of new Agents, the record for the Agent, along with its certificate, is provided to the Controller.
- iii. The Controller looks up the Agent certificate to see if an Agent record exists for it.

If the record does not exist, the creates a record for the Agent and stores a copy of the Agent's certificate inside of it. Any data retrieved from the Agent (e.g. a process listing, files from the hard drive, etc) is associated with the newly created Agent record. This process is referred to as enrollment.

If the Controller has seen the Agent certificate before, it identifies the Agent record for that certificate: in other words, the Agent has already enrolled. Any data retrieved from the Agent is associated with this pre-existing Agent record.

6.1.3.2. Time Sensitivity in the Trust Domain

As with any similar security infrastructure, certain authentication operations in the Trust Domain are sensitive to the time settings. For example, if an Agent is deployed on a system with a clock that is too far ahead or behind of the Controller, then issues may arise when it creates its self-signed certificate and attempts to register itself via Agent *Discovery Services* (see *Chapter 7, Agent Deployment* for more information).

6.2. Administering the Trust Domain

Trust Domain administration is accomplished through the Administration Console. The **Application** → **SSL** menu provides access to all of the functions needed to create, operate, and maintain the TDCA and its associated components.



Section 2.5, “Configuring the Initial Trust Domain” describes the basic steps required to set up a simple Trust Domain. This section details all available Trust Domain administration functions and describes deployment scenarios that may be encountered in larger, more complex environments.

6.2.1. Managing the TDCA

The Administration Console provides the ability to create and delete a Trust Domain Certificate Authority (TDCA). There must be at least one TDCA for every MIR deployment. All Controllers within the same Trust Domain must have a certificate signed by the TDCA. A Controller hosting the TDCA is referred to as a *Master Controller*.

All TDCA operations (creating the TDCA, signing certificates, creating/managing the CRL) must be accomplished directly on the Master Controller. Other Controllers within the same Trust Domain (that is, additional Controllers that have been authorized to connect to and

interact with the same set of Agents) will not have a TDCA installed on them. Instead they will have their SSL certificates signed by the TDCA managed from the Master Controller.

6.2.1.1. Configuring SSL Parameters

The first step in managing a TDCA is to set up the basic SSL parameters that will govern how the TDCA is created and the settings assigned to certificates signed by the TDCA. To set up these parameters do the following:

1. Choose **Application** → **SSL** → **CA** in the Navigation Bar.

The main page will display options for SSL key length and validity duration.

2. Configure the three options. The default values will allow you to deploy successfully.

SSL Key Length

The length of the public/private key pairs (in bits) associated with all entities in the system. MANDIANT recommends using 2048 bits for maximum security. Other valid values include 1024 and 4096. Note that longer key lengths increase the processing time for network connections.

New Certificates Valid For

The duration (in days) that Controller certificates are valid for. This affects how often you will need to re-install certificates, and is identical in concept to certificate duration for web servers.

New CRLs Valid For

The duration (in days) that Certificate Revocation Lists are valid for. Once a CRL expires, all parties relying upon it must fetch or receive a new copy. MANDIANT recommends the default setting of 180 days.

3. Click **Update Certificate Authority Configs**.



Certificate and CRL validity periods affect how frequently you will need to install new certificates into your Controllers, and how frequently Agents must update their CRLs.

If you have deployed Agents and properly configured the Agent *Discovery Service*, CRL updating for Agents is trivial. If, however, ADS is not active in your environment then a CRL update may necessitate Agent re-installation.

Consult *Chapter 7, Agent Deployment* for additional details.

6.2.1.2. Creating the TDCA



Do not re-create the TDCA if one exists. See the warning for *Section 6.2.1.4, "Deleting the TDCA"*, below.

Before creating a new TDCA ensure that the date and time on your Controller are correctly set. The TDCA certificate will have a validity period that is set based on existing system settings. Incorrect date and time information may create several problems during normal operation.

To create a new TDCA do the following on your Master Controller:

1. Choose **Application** → **SSL** → **Keys** in the Navigation Bar.

The main page will display key and certificate status for the TDCA and the server keys and certificate for your Controller.

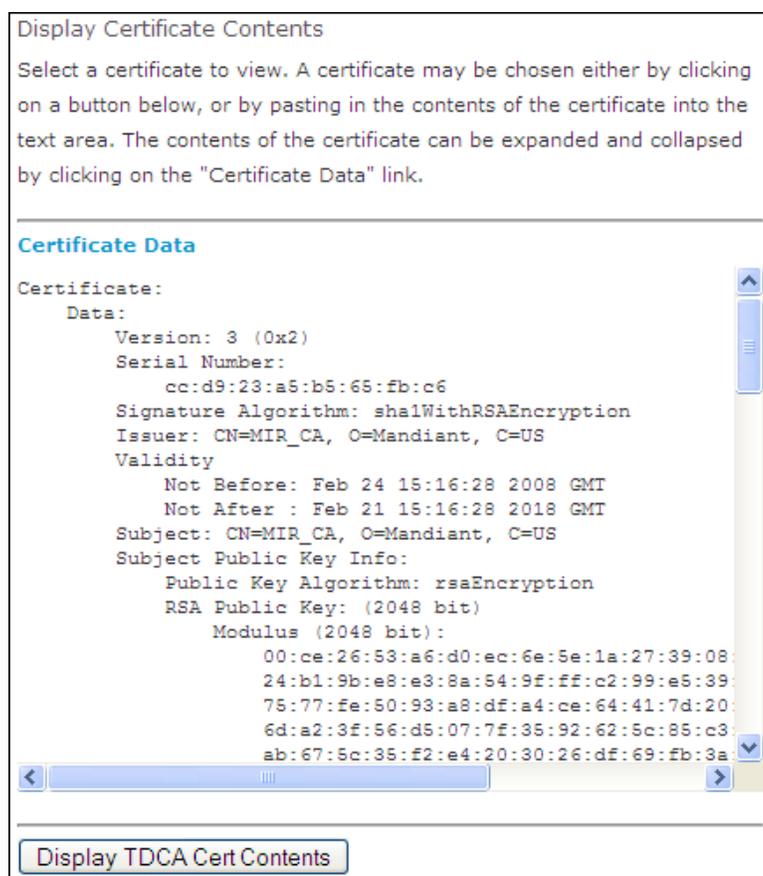
2. To create the TDCA key pairs and certificate, click the **Create TDCA Keys/Cert** button. Once complete, the TDCA status will change to **Exists**.

This process results in two entities being created: a public/secret keypair, and an X.509 certificate for the TDCA. The certificate contains the public key from your TDCA keypair and additional metadata, such as the name of the CA (this has been generated by the system: you do not need to configure it), its creation date, and validity period.

6.2.1.3. Viewing the TDCA

To view the contents of the TDCA certificate:

1. Choose **Application** → **SSL** → **DisplayCert**. Click the **Display TDCA Cert Contents** button to see a text-version of the TDCA certificate.



6.2.1.4. Deleting the TDCA



Deleting a TDCA will significantly impair operations within your Trust Domain. It may require you to re-deploy Agents and Consoles with new TDCA information. You should

only delete a TDCA if you are absolutely certain you are no longer going to be using it. We strongly recommend backing up the TDCA before deleting it from a running Controller in case you require it at a later date. See *Section 6.2.4, “Backup and Restore of Keys and Certificates”*.

To delete the TDCA:

1. Choose **Application** → **SSL** → **Keys** from the navigation bar.
2. Click the **Delete TDCA Keys/Cert** button.
3. Click **OK** to confirm you wish to delete the TDCA keys.

6.2.2. Managing Server Certificates

The primary purpose of the TDCA is to allow Agents to authenticate Controllers when they connect. This prevents Agent hijacking, allowing MIR to operate securely.

In order for a Controller to make connections to Agents within a Trust Domain it must have certificate installed on it that has been signed by the TDCA for that Trust Domain. When installing a certificate on a Master Controller – that is, a Controller that is also hosting the TDCA – most of this process is transparent. You only need to choose **Application** → **SSL** → **Keys** and then click the **Create Server Keys/Cert** button.

However, if you have multiple Controllers in the same Trust Domain, the process requires additional steps. This section provides additional guidance for this scenario:

6.2.2.1. Creating Server Keys for a Non-Master Controller

1. Choose **Application** → **SSL** → **Keys** in the navigation bar.
2. Click the **Create Server Keys/Cert** button.



This is the same process to fully create and install a key pair and a certificate on a Master Controller that is also hosting the TDCA. However, in this scenario the TDCA is not present. The key pair has been created, but a certificate signed by the TDCA is not yet installed.

6.2.2.2. Creating a Server CSR for a Non-Master Controller

A CSR is a Certificate Signing Request. It is submitted to a Certificate Authority in order to generate a certificate. The CSR contains information about the requested certificate, including the name of the system requesting the certificate and the generation date of the CSR.

The CSR is automatically generated when new Server keys are generated. To view the CSR:

- Choose **Application** → **SSL** → **Signing**. The box at the bottom of the main page will contain the CSR.

6.2.2.3. Signing and Installing the CSR on a Non-Master Controller

The CSR displayed in the step above (*Section 6.2.2.2, “Creating a Server CSR for a Non-Master Controller”*) must be submitted to the TDCA for signing. The easiest way to do this is to bring

up the Administration Console for both the Master Controller which contains the TDCA, and the Controller you are creating a new certificate for (referred to as *Subordinate Controller* for this exercise).

1. In the Subordinate Controller console:
 - a. Choose **Application** → **SSL** → **Signing** to display the CSR.
 - b. Copy the contents of **Server CSR** to the clipboard (highlight the contents and selecting **Edit** → **Copy**.) Highlight *everything* in the **Server CSR** box, including the lines that say `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST`.
2. In the Master Controller console:
 - a. Choose **Application** → **SSL** → **Signing**.
 - b. Paste the clipboard into the **Paste CSR Here** box. Check that the pasted contents include the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines.
 - c. Click the **Sign CSR** button.

If successful, the signed certificate will be displayed in the **Signed Certificate** box.
 - d. Copy the contents of the **Signed Certificate** box, including the lines labeled `BEGIN CERTIFICATE` and `END CERTIFICATE`.
3. In the Subordinate Controller console:
 - a. Choose **Application** → **SSL** → **Certs**.
 - b. Paste the clipboard into the **Server Certificate** box. Check that the pasted contents include the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines.
 - c. Click the **Update Certificate** button.

The certificate is now fully installed.

6.2.3. Managing the CRL

The Certificate Revocation List is the primary method that entities in the Trust Domain use to tell if a certificate is still valid. If, for example, you had deployed multiple Controllers and needed to take one out of service you would want to revoke its ability to access your deployed Agents. To do that you would revoke its certificate and redistribute the CRL.

6.2.3.1. Revoking a Certificate

To revoke a certificate you must first have a copy of it. You can obtain a copy of a Controller's certificate by choosing **Application** → **SSL** → **Certs** for the Controller with the certificate you want to revoke.

The **Server Certificate** box will display the encoded Controller certificate. Select and copy the entire contents, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines, and do the following:

1. Choose **Application** → **SSL** → **CRL** for the Master Controller.
2. Paste the certificate into the **Revoke** box.
3. Click the **Revoke Certificate** button.
4. The **TDCA Certificate Revocation List** box will now display the contents of the CRL. Update it by clicking the **(Re)Generate CRL** button.



With this release of MIR, manually inserting a CRL for distribution by a discovery service is not supported. This will be remedied in a subsequent release.

To display a more readable form of the CRL contents, choose **Application** → **SSL** → **DisplayCRL**. Click the **System CRL Contents** button to see the active CRL.

Display CRL Contents

Select a CRL to view. A CRL may be chosen either by clicking on a button below, or by pasting in the contents of the CRL into the text area. The contents of the CRL can be expanded and collapsed by clicking on the "CRL Data" link.

CRL Data

```

Certificate Revocation List (CRL):
  Version 2 (0x1)
  Signature Algorithm: sha1WithRSAEncryption
  Issuer: /CN=MIR_CA/O=Mandiant/C=US
  Last Update: Feb 24 15:23:24 2008 GMT
  Next Update: Feb 23 15:23:24 2009 GMT
  CRL extensions:
    X509v3 CRL Number:
      2
No Revoked Certificates.
  Signature Algorithm: sha1WithRSAEncryption
  6a:38:32:89:fb:19:26:f3:d2:4a:a6:3f:99:14:e9:74:17:
  b9:73:ce:19:e6:e7:e7:c0:7d:ea:86:0b:dc:76:11:63:d6:
  6a:14:d1:a4:ce:15:f1:c9:42:65:21:cc:b6:d4:e2:69:34:
  35:de:2f:fe:a9:49:b5:1e:16:3f:1d:90:52:89:69:60:2a:
  1d:5a:d8:f2:5b:d2:71:57:79:38:cc:eb:aa:34:17:4c:8c:
  98:5b:94:37:41:4e:92:fb:6d:f1:fb:cb:af:f3:10:7c:c6:
  77:b7:2c:1e:35:09:eb:6e:08:f1:64:d5:b3:35:5b:a2:4b:
  28:13:be:3d:56:08:99:64:6e:74:71:86:d1:be:9f:8c:e3:
  01:04:57:16:10:50:15:10:10:10:10:10:10:10:10:10:
    
```

System CRL Contents

6.2.4. Backup and Restore of Keys and Certificates

The TDCA is extremely important to any MIR deployment. If it is lost then new Controllers cannot be added to a Trust Domain and certificates within the Trust Domain cannot be revoked. The only remedy in the case of a lost TDCA is to create a new one and re-deploy all Agents within the Trust Domain. In order to avoid this problem, the Controller provides a method to backup the TDCA. Though less critical, the server certificate and keypair may also be backed up.



The security of your MIR deployment rests on the security of the TDCA.

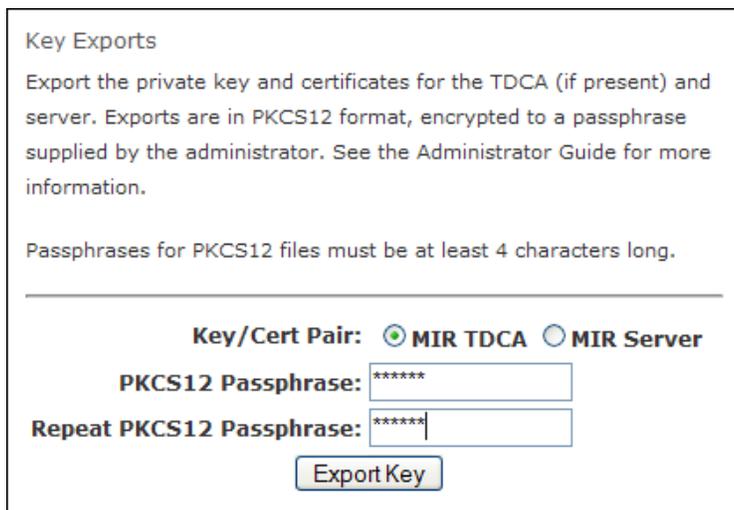
You should protect the backup files (`ca.p12` and `server.p12`) as you would any other extremely sensitive data within your infrastructure. The TDCA can create other certificates that are able to access any Agent within the Trust Domain. The server certificate from a Controller can be used to directly access Agents. Backups of these materials should be treated with appropriate physical and logical security measures to prevent compromise.

The TDCA is critical for disaster recovery. **Offsite backups and physical protections are best practices.** See *Chapter 6, Understanding the Trust Domain* for more information about the TDCA.

The passphrase set during key export is as critical as the backup file. If you have the `ca.p12` or `server.p12` file, but lose or forget the passphrase you set during backup, you *will not* be able to restore the TDCA or the server certificate.

Backing up the TDCA using the Administration Console

1. Log into the Administration Console and select **Application** → **SSL** → **Export Keys** in the navigation panel.



2. The **Key Exports** page will be displayed. Select **MIR TDCA** and set a passphrase in the boxes provided.



This passphrase is used to encrypt the TDCA public and private keys. The strength of the passphrase directly corresponds to the degree of protection provided to the TDCA keypair.

Click **Export Key**. You will be prompted by your browser to download and save a file named `ca.p12`. Save the file to a safe location.

3. Select **MIR Server** and, again, set a passphrase. The same precautions apply to this backup as to the TDCA backup.

Click **Export Key**. You will be prompted to download and save a file named `server.p12`. Save the file to a *safe location*.

Backing up the TDCA through SSH

- Log into the Controller, tar the keyfiles, and copy them to a *safe location*.

```
ssh -l [admin] [controller]
sudo tar cvfz /home/[admin]/mir_ca.tgz /opt/apollo/etc/mir/ssl/private/mir_ca*
exit
scp [admin@controller]:mir_ca.tgz ./
mv mir_ca.tgz [safe location]
```

Restoring the TDCA using the Administration Console

1. Copy the `ca.p12` or `server.p12` files to the system on which you are using the Administration Console.
2. Choose **Application** → **SSL** → **Import Keys**.
3. Select **MIR TDCA** or **MIR Server**, as appropriate.
4. Type the passphrase for the backup file.
5. Click **Browse** and select the backup file. Click **Import Key**.

The selected backup file will, if you provided the correct passphrase, be uploaded and installed into the running system.

6. Choose **Application** → **Components** and click **Restart All MIR Processes**.

Restoring the TDCA key through SSH

1. Copy the key to your Controller account, stop the Controller, restore the key, and restart the server:

```
scp mir_ca.tgz [admin]@[controller]
ssh -l [admin] [controller]
sudo su
/etc/init.d/mir stop
cd /
tar xvfz /home/[admin]/mir_ca.tgz
/etc/init.d/mir restart
exit
```

1. Log into the Administration Console.
2. Select **Application** → **SSL** → **Keys** in the navigation panel.
3. Click **Delete Appliance Key** followed by **Create Appliance Key** to recreate your appliance keys.
4. Select **Application** → **SSL** → **CRL**.
5. Click **(Re)Generate CRL** to recreate the certificate revocation license.

6. Select **Application** → **Components**.

7. Click **Restart All MIR Processes** to effect the changes.



If you restore a TDCA to a Controller you need to also restore an SSL server certificate issued from that TDCA or create a new server certificate once the TDCA is installed. See *Section 6.2.2, “Managing Server Certificates”*.



Chapter 7

Agent Deployment

To collect information from systems using MANDIANT Intelligent Response an Agent must first be deployed on those systems. In most installations, the Controller and Agent communicate over the network. Alternately, the Agent can be used in *Portable Use Mode*, collecting data and capturing it to local storage (typically a USB key) for later import into a Controller. In either case, Agents must be properly configured, deployed, and managed to properly collect information from target systems.

This chapter discusses the management of Agent deployments, including the Agent *Discovery* Service on the Controller, the Agent installation package materials, and the use of enterprise software management systems to install and uninstall Agent software.



This guide assumes that you have already installed and performed basic configuration of your Controller. Additionally you must have properly configured your Trust Domain. See *Chapter 6, Understanding the Trust Domain* for more information.

7.1. The Agent *Discovery* Service

It is important to understand how the Agent *Discovery* Service works. A properly configured *Discovery* service will make it significantly easier for users to find, manage, and collect information from deployed Agents.

Discovery provides a method for Agents to automatically “report in” to a Controller, identifying their network information, certificate, and Agent version. In a dynamically addressed environment, the *Discovery* service allows an Agent to be tracked automatically if the IP address of the target system it is on changes over time.

Use of *Discovery* is not required to use MIR. It does, however, greatly reduce the configuration overhead needed to begin the collection of information from deployed Agents. Some features are only available if the *Discovery* service is configured and deployed Agents can successfully communicate with the service.



Use of the *Discovery* service in a MIR deployment is optional. However, for Agents to receive updates to their Certificate Revocation Lists (see *Section 6.2.3, “Managing the CRL”*) *Discovery* must be used. When *Discovery* is not used users must add and configure host objects within MIR, including host address information (see the *User Guide* for more information).

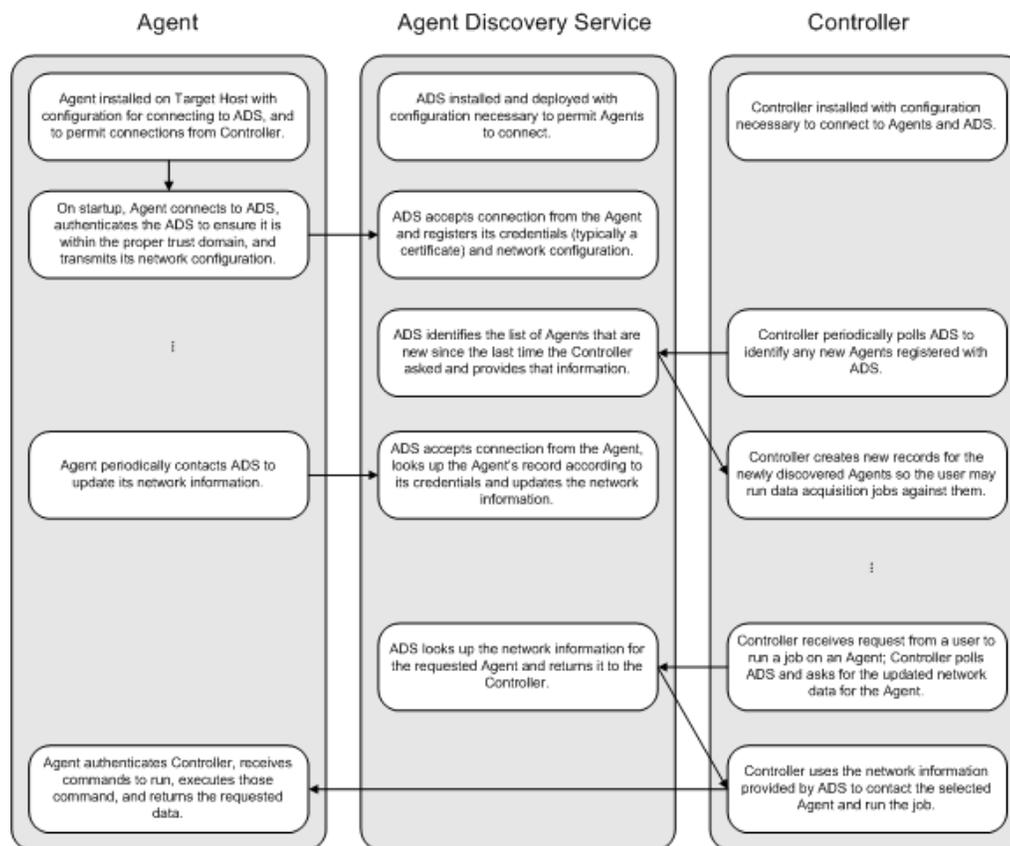
7.1.1. Functional Overview

A *Discovery* service instance runs on each Controller. Agents are configured with the location of a *Discovery* service at the time of install. When the Agent is installed it communicates with the Controller’s *Discovery* service and registers itself by transmitting information about its identity, network configuration, and software version. The Agent will periodically contact the service to update its configuration information, ensuring the Controller remains up-to-date.

As Agents register themselves, they are made available to Console users through the **Hosts** Library. The figure below details the various interactions between Agents, *Discovery*, and

the Controller. Note that for this version of MIR the *Discovery* service resides directly on the Controller, and that Agents may only interact with a single *Discovery* instance.

Figure 7.1. ADS Functional Overview



7.1.2. Network Requirements

The *Discovery* service needs the Agent to initiate TCP connections to the Controller on a specified port (TCP 8077 by default). Also, the Controller must be able to initiate TCP connections to Agents on a specified port (TCP 22201 by default) to launch Audit Jobs.

Controller ⇒ Agent TCP 22201

The Controller initiates connections to deployed Agents when performing Audits. The default listening port on the Agent is TCP 22201. Data collections are done across these connections; the speed of data acquisition is directly related to the bandwidth available between Controller and Agent.

Agent ⇒ Controller TCP 8077

OPTIONAL: Agents initiate connections to the Controller to register themselves using the Agent *Discovery* Service. If the Agent *Discovery* Service is not being used, this connection is not needed.



It is strongly recommended that you assign a hostname in DNS to your Controller and use that to configure Agents for *Discovery*. This will allow you to change the IP

address of the Controller without having to re-deploy Agents. Agents cannot self-update their *Discovery* settings; changes to their configuration requires an update to the Agent configuration file and subsequent re-installation.

7.1.3. Discovery Service Configuration

To use *Discovery*, the service must be configured on the Controller before Agent deployment. The Agent installation package must then be properly configured with the correct *Discovery* service settings information.

7.1.3.1. Configuring Agent *Discovery* Service on the Controller

1. Configure the Controller network settings per *Section 3.4.4, "Network Settings"*. Ensure you assign a fully qualified domain name to the Controller's interface to make ongoing maintenance and configuration of Agent *Discovery* Service easier.
2. Configure the TDCA per *Section 6.2, "Administering the Trust Domain"*.
3. The default listening port for the Agent *Discovery* Service on the Controller is TCP 8077.

If you wish to change this to a different TCP port you can configure that setting through the Administration Console, by choosing **Application** → **ConfigFiles** in the left navigation panel, and then selecting **Discovery Service** → **Discovery Server Location** on the right.

4. Ensure Agent *Discovery* Service is running by choosing Application in the left navigation panel. **Discovery Status** will be **green** if Agent *Discovery* Service is operating correctly.



Administration Web Console - **192.168.56.103**

<p>Administration</p> <hr/> <p>Appliance</p> <p>Application</p> <hr/> <ul style="list-style-type: none"> * Components * ConfigFiles * Database * Discovery * Resources * SSL <hr/> <p>Users</p>	<p>MANDIANT Intelligent Response Administration Console</p> <p>View and update MIR Controller configurations.</p> <hr/> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 60%;">Discovery Status</td> <td style="text-align: right;">Green</td> </tr> <tr> <td>Database Status</td> <td style="text-align: right;">Green</td> </tr> <tr> <td>MIR Status</td> <td style="text-align: right;">Green</td> </tr> <tr> <td>TDCA Configured</td> <td style="text-align: right;">Green</td> </tr> </table>	Discovery Status	Green	Database Status	Green	MIR Status	Green	TDCA Configured	Green
Discovery Status	Green								
Database Status	Green								
MIR Status	Green								
TDCA Configured	Green								

7.1.4. Discovery Service Administration

The Controller provides several functions for administering a running *Discovery* service through the Administration Console. *Discovery* administrative functions can be accessed using **Application** → **Discovery**.

7.1.4.1. Bulk Load

The current state of the *Discovery* service can be saved and transferred to another *Discovery* instance on another Controller. It can also be used as a backup for the *Discovery* service. Since Agents typically continually poll *Discovery* and re-register themselves, loss of *Discovery* data is usually not an operational emergency. Nonetheless, using the bulk load feature the information can be saved and restored. These functions are accessed through **Application** **Discovery** → **Bulk Load**.

Discovery state is saved by clicking the **Discovery Data** link. To import *Discovery* state information, click **Browse** in the **Reload Discovery Service State** section. Once you have selected the *Discovery* state file, click **Import**.



Existing *Discovery* data will be replaced with the information in the bulk load file.

Import/Export Discovery Service State

Export Discovery Service State

Right-click on the link below, and select "Save Link As..."

[Discovery Data](#)

Reload Discovery Service State

Click the "Browse" button below to locate the Discovery data file you wish to reload into the local discovery service. Click the "Import" button to load the file.

7.1.4.2. Discovery Query

You can directly query the *Discovery* database for a specific Agent based on Agent certificate using the Administration Console. Choose **Application** → **Discovery** → **Discovery Query**. The main page displays an input box that accepts Agent certificates in PEM-encoded format. To perform a query, paste in the PEM-encoded certificate of the Agent you are searching for and click **Query Discovery**. The query will return the current IP address and port number of the Agent that corresponds to the submitted certificate. See *Chapter 6, Understanding the Trust Domain* for more information about Agent certificates.

7.1.4.3. Discovery Reset

In some circumstances it may be necessary to delete the information in the *Discovery* database. To reset *Discovery*, choose **Application** → **Discovery** → **Discovery Reset**, then click **Reset Discovery Database**. (A normal database reset as described in *Section 4.3, “The Application Database Menu”* does not remove the information contained in the *Discovery* database.)



Resetting the *Discovery* database is a common troubleshooting action when issues arise with Agents not being discovered or duplicate/multiple Agent entries. MANDIANT Customer Support can provide further help and guidance on this topic.

7.1.4.4. Discovery Search

In addition to the certificate query capability accessed via the **Discovery Query** menu, the *Discovery* service provides a search service available through the Administration Console that allows you to find Agents based on an IP address or range of addresses.

To perform a search, input an IP address or address range using one of the supported formats outlined in the main page and click the **Search Discovery** button. All matching Agent records will be returned, including Agent certificate, IP address, hostname, port number, Agent version, and last update time for that Agent record. Multiple individual IPs may be searched for, in addition to multiple ranges by separating arguments with commas. Consider the following examples:

Specify Network Range Using Dotted-decimal Netmask Notation

```
192.168.0.0/255.255.255.0
```

Specify Network Range Using CIDR Notation

```
192.168.0/24
```

Specify a Network Range Using Encompassing IP Addresses

```
192.168.0.0-192.168.0.255
```

Search for Individual IPs as well as a Network Range

```
172.16.12.134, 66.24.26.83, 192.168.0/24
```



This version of MIR supports IPv4; IPv6 is not fully supported.

7.1.4.5. Duplicate Agent Detection

Duplicate Agent Certs detects identical Agents on the network.

7.2. Agent Installation

Once you have correctly configured your Controller, Trust Domain, and *Discovery* service, you can begin to deploy Agents. The Agent can be installed interactively using the standard Windows installer, or silently installed using `msiexec`, provided by Microsoft. This section details system requirements and the structure and function of the Agent installer package.

Agent Requirements:

- 32 bit versions of Windows 7, Microsoft Vista, Windows 2003 SP2, Windows 2000 SP4, and Windows XP SP2 or higher.

- 64 bit versions of Windows 7, Windows 2008 R2, and Microsoft Windows 2003 SP2.
- 512 MB of RAM.
- 12 MB of free disk space.
- By default, the Controller must be able to initiate TCP connections to the Agent on port 22201.
- To test Agent *Discovery Service*, the Agent must be able to initiate TCP connections to the Controller on port 8077.
- Administrator-level privileges for installation.



On Hosts that do not support memory-related audits, such as Windows Vista 64 bit, an Issue Document will be returned indicating an error occurred and that the Audit could not determine the OS version.



A long-standing bug in the Windows operating system causes any program using the `windows\temp` directory to fail when the directory has reached its 64K space allotment for files or folders.

In such case, the following message is written to the Agent log file (see *Appendix B, Error Messages and Troubleshooting*):

```
WARNING [Discovery CheckForUpdates]- Agent was not able to
create a temporary download folder. Agent will not update.
(13)
```

The solution is to manually remove some or all of the files in the `windows\temp` directory before attempting to install the Agent again.

7.2.1. The Agent Installer

The Agent installer package is a standard MSI file (`AgentSetup.msi`) that can be installed interactively at the console of a target system or silently using the `msiexec` utility from Microsoft. Other software management utilities that are able to manage installation of MSI files and associated configuration files (such as Microsoft Systems Management Server) should similarly be able to deploy the Agent.

7.2.1.1. Standard, Service, and Portable Installations

The Agent supports three different installation modes: *standard*, *service*, and *portable*. In standard and service modes all internal security information is encrypted using Microsoft's *Data Protection API*¹. DPAPI encrypts information using a machine-specific key: this means that files encrypted with one system's key can not be decrypted on a different system. As a result, standard- or service-installed Agents can not be copied or moved to another system.

In portable mode the Agent's security information is not encrypted, and can thus be copied to multiple systems.

Standard

Properties	Agent installed to target location, Agent not running after install, security information encrypted using DPAPI.
Uses	Install Agent on target, but do not activate. Use when you want a persistent install and you wish to run the Agent in daemon mode, or

¹More information on Microsoft's DPAPI is available at <http://msdn.microsoft.com/en-us/library/ms995355.aspx>.

when you want to control the Agent from the command line. See the for command-line details.

Service

- Properties Agent installed to target location, Agent running as a Windows service after install, security information encrypted using DPAPI.
- Uses Installs and activates the Agent on a target system. Agent persists through system reboots. Use when you want a permanently running MIR Agent on a target. Useful when making the Agent part of a default system image for enterprise deployment.

Portable

- Properties Agent installed to target location, Agent not running after install, security information not encrypted, allowing Agent files to be copied to another system for use.
- Uses Install an Agent to a location that allows it to be copied to other systems. May be used to place Agent files on a network share or removable media for copying them to target systems. Allows you to put an Agent on a target without running the full installer.

The tasks below describe specific installation steps for common deployment scenarios. As with any software deployment, we recommend testing deployment of Agent software before executing enterprise-wide installations.



If a previous version of an Agent exists on the Host, the installer will perform an “upgrade” installation. During such an upgrade the old Agent is removed prior to installing the new Agent. If the installation process subsequently fails the installer will not find an existing Agent and will assume it is performing a full, fresh install. As a result, the previous user configuration files and certificates are overwritten.

7.2.1.2. Generating the Agent Configuration File

You will need administration access to the Controller.

1. Using any computer connected to your Controller network, log into the Controller Console:

```
https://[Controller URI or Hostname]/administration/
```

2. Choose **Appliance** → **Config** → **AgentConfigFile**.. This opens the **Agent conf.xml Generation** page.
3. Configure values for the Agent installation package.

You must provide a valid **Controller DNS Name** value. All other settings may be left in their default state.



We recommend that you review your previous Agent package and where possible, use the values from the `system.settings.xml` and `discovery.xml` files.



We recommend leaving **Use Agent Hiding?** set to **No**. If you would like to enable *Agent Hiding*, please contact MANDIANT Product Support *Section 1.3, “How to Get Support”*.

4. Click **Update conf.xml** to generate a new configuration file. A **Processing** timer will be displayed, and the page will be reloaded.
5. Right-click **Download conf.xml** and select **Save Link As...** Use the filename `conf.xml`.



You may send the generated **conf.xml** file to MANDIANT Product Support for review. We're happy to help you!

6. When creating an Agent installation package (see below), replace the original `conf.xml` with the customized `conf.xml`.

You are now ready to create an upgrade installation package.

7.2.1.3. Agent conf.xml Generation Settings

Discovery Server Interaction

Contains settings for the Agent's *Services* entry in the Windows Management Console.

Controller DNS Name

The Controller's *Discovery* Service address and port.

Time to Retry Discovery

The number of times the Agent will attempt to contact the Controller before taking a time-out.

Discovery Callback Interval (minutes)

The length, in minutes, of an error retry time-out.

MIRAgent Service Settings

Contains settings for the Agent's *Services* entry in the Windows Management Console.

MIRAgent Service Name

The short-form name of the Agent. This value should not be modified.

MIRAgent Service Display Name

The long-form name of the Agent. This value should not be modified.

MIRAgent Service Description

The description displayed in the Agent Service details. This value should not be modified.

MIRAgent Behaviors

Contains settings for the Agent's behavior on the network and user interface.

Issues Document Verbosity

When creating an Audit Issues Document, detailing the operation and success or failure of the Agent Module, include all information at and above this level. **debug**, **info**, **warning** and **error** are acceptable values providing progressively less information (and less network traffic.)

Use Agent Hiding?

On 32-bit Windows XP and Vista hosts, hides the process name from *Task Manager*, hides the installer from *Add/Remove Programs*, and hides the installation folders. When **false**, the user can easily see and kill the Agent process and can un-install the Agent.

On 64-bit Vista and on Windows 7 hosts, this setting hides the installer from *Add/Remove Programs* and hides the installation folders. The process name is not hidden in the *Task Manager*.

Hide from Add/Remove Programs?

Hides the Agent installer from the Windows *Add/Remove Programs* configuration utility. When **false**, the user can easily un-install the Agent. This setting is applicable only when *Use Agent Hiding?* is *false*.

Firewall Exception Mode

Determines how the Agent will handle Windows Firewall exceptions. Exceptions will be removed when the Agent is uninstalled or when `MIRAgent.exe` is executed with the `-cleanup` or `-dissolve` options.

Auto

Add an exception when starting the MIRAgent process and remove it when stopping.

ServiceAdd

Add an exception once only when installing as a service and remove it when uninstalling..

AlwaysAdd

Add an exception once.

NeverAdd

Do not add an exception.

Agent Listening Port

Sets the port on which the Agent will listen for Controller communications.

MIRAgent Logging Settings

Contains settings for Agent log reporting.

Use Agent Logging?

Sends Audit and Agent logging information to the Controller. When **false**, logging information will not be available in Consoles, and network traffic will be reduced.

Audit Log Level

Includes all logging information at and above this level. **debug**, **info**, **warning** and **error** are acceptable values providing progressively less information (and less network traffic)

Max Log Rollover Index

The maximum index number for rolled-over log files. When rolled-over log index numbers are incremented, the oldest log file will be deleted if its index value exceeds the maximum index number.

Min Log Rollover Index

The starting index number for rolled-over log files.

Log Size Rollover Trigger (KB)

The maximum size of a log file before the current log file renamed to `[filename].[ext].[min log rollover index]`. When this happens, existing rolled-over log index numbers are incremented.

Agent Log File Name Prefix

Sets the name of the host log file. See *Appendix B, Error Messages and Troubleshooting* for location and content details.

Agent Log File Name Extension

Sets the extension of the host log file name.

7.2.1.4. Installing an Agent Using the Windows Installer

The easiest way to install an Agent is with the Windows installer. This method lets you quickly perform a default install, but limits your customization options. By default the Agent is installed as a Windows Service listening on TCP port 22201.

You must have Admin-level privileges to install the Agent.

1. Create a folder named `MIR Install` on the Host or on portable media.
2. Copy `AgentSetup.msi`, the Agent installer, from the MIR software CD to the `MIR Install` folder.
3. Copy `conf.xml` – created in *Section 7.2.1.2, “Generating the Agent Configuration File”* – to the `MIR Install` folder.
4. Start `AgentSetup.msi` and follow the prompts:
 - a. In the **Welcome to...** window, click **Next**.
 - b. In the **License Agreement** window, select **I Agree** and click **Next**.
 - c. In **Select Installation Folder**, accept the default installation destination unless you have reason to install to a different location. Clicking **Disk Cost...** will display the available and required drive space for the installation.

Click **Next** to continue.
 - d. In the **Confirm Installation** window, click **Next**. A progress bar will advance as the software is installed.
 - e. In the **Installation Complete** window, click **Close**.

Note that Agent installation does not require a system restart.

The remaining scenarios make use of a Windows command-line utility, allowing for detailed customization of the installation:

7.2.1.5. Installing an Agent using `msiexec.exe`:

`msiexec` is a command-line installation utility included with Microsoft’s Windows Installer package, and is present by default on Windows systems. The Agent installer can be installed using `msiexec`, which allows you to bypass the interactive installation GUI, and complete installation in a single step. There are also some install-time configuration options that are otherwise not accessible via the interactive installation process.

`msiexec` options include the ability to set the amount of interaction with the installer, set the logging level for `msiexec` output, and specify restart actions after the install

(or uninstall) completes. For more information on `msiexec` options consult Microsoft's documentation, or type `msiexec /?` in a Windows command prompt to see abbreviated help text.

Basic `msiexec` usage via a Windows command prompt is as follows:

```
msiexec.exe {[/i] | [/x] | [/?]} [package]
```

Some of the more commonly used `msiexec` options for Agent installs are listed below:

```
msiexec.exe {[/i] | [/x] | [/?]} [/q[n|b|r|f]] [/quiet] [/norestart] [package] [properties]
```

/i

Install *package* to the system.

/x

Remove *package* from the system.

/qn | /qf

Selects a user interface level for the installation.

/qn

No user interface.

/qf

Use the full user interface.

/quiet

Use quiet mode. Do not prompt for user interaction.

/norestart

Do not restart the system after installation is complete.

/?

Prints a help message.

package

The file path to the `.msi` installation file.

properties

Settings to be used by the installer, as listed below.

TARGETDIR

Specifies where to install the Agent. Default is: `C:\Program Files\MANDIANT\MANDIANT Intelligent Response Agent`.

ALLUSERS

Specifies whether to perform a service mode install. Default value is 1 ("yes").

INSTALLSERVICE

Specifies whether to perform a portable mode install. Default value is 0 ("no"). Options are 0 ("no"), 1 ("install as service and start it"), and 2 ("install as service but do not start it").

Typical Installation Steps

1. Place the Agent installer (`AgentSetup.msi`) file on the host.
2. On the host system, open a Windows command prompt.
3. Change to the directory where you placed the Agent installer file (e.g. `cd directory`).
4. Assuming a quiet/non-interactive installation using a target directory of `C:\MIRAGENT`, installed for all users on the system, type the following at the command prompt:

```
msiexec /quiet /i C:\AgentSetup.msi ALLUSERS=1
```

Scenario-specific examples follow:

Installing a Standard Persistent Agent on a Single Workstation

1. Place the Agent installer and associated configuration files on the host system (assuming `C:\AgentSetup.msi`).
2. Open a Windows command prompt. Type the following:

```
msiexec /quiet /i C:\AgentSetup.msi ALLUSERS=1
```

3. The Agent will fully install on to the host system as a service listening on TCP port 22201, with an installed firewall exception for the Windows firewall (if active). The Agent will be placed in `C:\Program Files\MANDIANT\MANDIANT Intelligent Response Agent`.

Installing an Agent to a Portable Drive

1. Place the Agent installer and associated configuration files on the host system (assuming the installer is at `C:\AgentSetup.msi` and will be installed to `E:` drive).
2. Open a Windows command prompt. Type the following:

```
msiexec /quiet /i C:\AgentSetup.msi TARGETDIR=C:\MIRAgent
```

3. Copy the installation to the removable media device:

```
copy C:\MIRAgent E:\MIRAgent
```

4. The Agent will fully install on to the host system as a service listening on TCP port 22201, with an installed firewall exception for the Windows firewall (if active). The Agent will be placed in `C:\Program Files\MANDIANT\MANDIANT Intelligent Response Agent`.

Installing an Agent to a Network Share, then Copying to Targets

1. Place the Agent installer and associated configuration files on a system where you can mount your target network share (assuming the installer will be placed at `C:\AgentSetup.msi`).

2. Open a Windows command prompt. Type the following:

```
msiexec /quiet /i C:\AgentSetup.msi MULTIINSTALL=1  
INSTALLSERVICE=0 INSTALLFIREWALLEXCEPTION=0
```

3. Copy the Agent installation directory (C:\Program Files\MANDIANT\MANDIANT Intelligent Response Agent) to a network file share where you have read and write permissions.

To install the Agent on a host system that has read access to the share mount:

4. Copy the Agent directory from the network share to the host system.
5. Start the Agent as either a daemon or a service on the host system (see *Appendix A, Agent Command-line Reference* for complete information on using the Agent from the Windows command prompt).
6. To perform a basic installation and activation of the Agent as a service, open a Windows command prompt on the host system, change to the directory containing the Agent, and type the following:

```
miragent.exe -i -start
```

Installing an Agent to a Master Image. If you wish to make the Agent a part of a default system image that you use throughout your enterprise you can add it to a system you are using as a template.

The sequence of steps you follow is important since an Agent, once it starts on a host system, creates a certificate for authenticating with the Controller for your MIR deployment. If you take an image of your template system with these credentials in place your deployment will likely not work and will require you to reset the credentials on all deployed Agents.

The following instructions will address this issue. If you require additional assistance preparing the Agent to be part of an enterprise-wide system template please contact MANDIANT Customer Support.

1. Place the Agent installer on your template system (assuming C:\AgentSetup.msi).
2. Open the Windows command prompt and type the following:

```
msiexec /quiet /i C:\AgentSetup.msi MULTIINSTALL=1  
INSTALLSERVICE=0 INSTALLFIREWALLEXCEPTION=0
```

Now you need to install the Agent as a service without starting the service itself:

3. Open a Windows command prompt and navigate to the Agent directory (C:\Program Files\MANDIANT\MANDIANT Intelligent Response Agent).
4. Install the Agent as a service by typing the following:

```
miragent.exe -i
```



This technique is used because the Agent needs to be installed without running it. When the Agent first runs it creates credentials for use with the Controller.

To ensure new systems created from your template system each receive unique credentials they need to be created when the new system first starts. This can only occur if the credentials are not already present.

5. Shut down the template system and create a master image using the drive imaging/templating technology of your choice. When a new system created from this master image is first started, the Agent will run, creating new, unique credentials.



The sequence of events and timing for the above technique is critical to install success. Please test your installation by verifying that systems created with your master image receive unique credentials. You can verify this by viewing systems communicating with *Discovery*.

7.2.2. Using a Portable Agent

To use the portable Agent, take the USB drive to a host system and insert it. You may then open a Windows command prompt and directly use the Agent in either local mode, or you may start it as a daemon or a service. If you use the Agent as a daemon or a service, it will create new certificates to use for security purposes with the Controller. A good practice is to ensure you reset these each time you take the portable Agent to a new system.



Client scripts must be converted for use by the portable Agent.

Use the following process to always “start clean” when using a portable Agent on a new system:

1. Insert the USB drive into the host system.
2. Open a Windows command prompt and change to the drive/directory containing the Agent (assuming `E:\miragent` for this example).
3. You may now start the Agent as either a daemon or a service. For instance:

```
miragent -portable -o -script SystemAudit.Batch.xml
```

See *Appendix A, Agent Command-line Reference* for complete information on using the Agent from the Windows command prompt.

7.2.3. Configuring Agent Upgrades via *Discovery Services*



Agent Upgrades via Discovery Services is an alpha feature. While it has been used successfully to deploy Agent upgrades in several large enterprises (40 000+ hosts), we strongly urge users to review postings in the MANDIANT Customer Forums to better understand its limitations. Search for the topic “MIRAgent Auto Upgrade for Fun and Profit” for details.

On occasion, MANDIANT may provide updated Agent software. When possible, these updates will be provided as a secure bundle that allows Agents to upgrade themselves over the network. Existing configurations—keys, discovery, etcetera—will remain unchanged by the upgrade.

7.2.3.1. Upgrading Agents through *Discovery Service*

1. Choose **Application** → **Discovery** → **Agent Upgrade**.
2. In the **Import Agent Upgrade Bundle** area, click **Browse...** and use the standard file selector window to select the upgrade bundle.
3. Click **Import**. The Controller will unzip the bundle, validate its contents and sign it.

The bundle will then be listed in the **Imported Bundles** area, along with configuration controls.
4. Some bundles allow configuration. Click **Configure** to view and adjust the bundle's configuration options. Click **Apply** when done.
5. Select **Active** and click **Apply** to make the bundle available to Agents.

To subsequently disable the bundle, clear the **Active** checkbox and click **Apply**.
6. Use the **Netmask** entry to limit the bundle's availability to specific subnets.

7.2.3.2. Customizing the *Discovery Service Upgrade*

The Agent upgrade bundle contains several files. Do not modify these files:

- Agent installer(s), with a `.msi` extension.
- A manifest containing validation information, named `_manifest.xml`.
- An optional signature named `_manifest.xml.sha1`.

The upgrade bundle may also contain `conf.xml`, a default configuration; you may modify this file.

7.2.3.3. Limiting Upgrade Bandwidth Usage

1. Choose **Application** → **Discovery** → **Agent Upgrade**.
2. In the **Configuration** area, set **Maximum number of connections** and **Maximum file transfer rate (bps)** appropriately for your network capacities.

Note that total bandwidth is the product of (connections × transfer rate). Ten connections at 100 kbps will consume 1000 kbps.

7.2.4. Uninstalling Agents

If an Agent is not running in hidden mode (ie. it was installed with `<key name="hidden">false</key>` **and** `<key name="hide_from_addremove">false</key>`) it may be uninstalled using the Windows Add/Remove programs control panel.

If an Agent is running in hidden mode (`<key name="hidden">true</key>` **or** `<key name="hide_from_addremove">true</key>`) it may be uninstalled by running the Agent installer (`AgentSetup.msi`) again: you will be prompted to repair or remove the Agent.



It can sometimes take up to 5 minutes for all of the Agent directories to be removed. It is advised that after a manual uninstall or dissolve audit, you wait at least 5 minutes before installing a new Agent; this gives the Controller and Windows time to complete the operation.

7.2.5. Third Party Security Products

The MIR Agent has been tested with several third party security products, including those included by default with Microsoft Windows. As with any enterprise agent-based technology, host-based security software must be configured to permit the Agent to perform the tasks it needs to perform. Most notably:

Network

The Agent must be permitted to receive TCP connections on the port it has been configured to listen on (TCP 22201 by default). The Agent must also be able to connect outbound to the *Discovery* service it has been configured for, as specified in the Agent's `conf.xml` installation configuration file.

Agent File Installation

Agent files are placed in `C:\Program Files\MANDIAN\MANDIAN Intelligent Response Agent` by default. The Agent must be able to execute from this directory, and must also have read and write access.

Drivers

When a memory audit is first run against an Agent, that Agent installs a driver that is required for direct memory access. This driver is placed inside of `TARGETDIR\mktools.sys`, as specified at Agent install time, and is named `MANDIAN_tools`.

Registry Modification

The Agent may install one or more registry keys depending on the options selected at install time, and must therefore be able to make appropriate Windows API calls to create those registry keys. Note the Agent does not do anything unique; it utilizes well documented, standard Windows methods to install and register itself.

User Permissions

In order for the Agent to fully access all information you may wish to retrieve it must be installed to run as a user with Administrator permissions. Failure to do so may cause errors for several audits, including: *w32disk-acquisition*, *w32memory-acquisition*, *w32rawfiles*, *w32rawfiles-acquisition*, *w32processes-memory*, and *w32registry-raw*.

7.2.5.1. Digital Signatures

The Agent installer, the Agent executable, and all associated modules have been signed by a Certificate Authority that is present by default on Microsoft Windows systems. MANDIAN uses Thawte for its code signing certificates. The MANDIAN certificate is issued off of the Thawte Code Signing CA, which chains to the Thawte Premium Server CA. If you encounter any issues validating digital signatures on the Agent installer or any MIR component, please contact MIR Customer Support.



Appendix A

Agent Command-line Reference

A.1. Commands and Flags for Using the Agent

Commands

```
MIRAgent.exe {[-?] | [-i] | [-u] | [-d] | [-o [basedir]] | -regencert} [secondary flags]
```

Flags for Conducting Local Audits

```
MIRAgent.exe -o [basedir] [-f] [-script filename] [-encoding {gzip | aff | none}] [-notimestamp] [-auditissuefilterlevel {debug | info | warning | error}] [-allowmultiple] [-portable] [-cleanup]
```

Flags for Starting the Agent in Daemon Mode

```
MIRAgent.exe -d [-bind ip address[:port number]] [-allow network/mask, ...] [-settings filename] [-fw {on | service | auto | off}] [-allowmultiple] [-certpath path] [-keypath path] [-capath path] [-crlpath path] [-credspath path] [-memcert] [-disablediscovery] [-auditissuefilterlevel {debug | info | warning | error}] [-encoding {gzip | aff | none}] [-portable] [-cleanup]
```

Flags for Installing the Agent as a Windows Service

```
MIRAgent.exe -i [-start] [-servicename name] [-servicedisplay displayname] [-settings filename]
```

Flags for Uninstalling the Agent if it was Installed as a Service

```
MIRAgent.exe -u [-servicename name]
```

Flags for Regenerating Agent Credentials

```
MIRAgent.exe -regencert [path] [-regencreds]
```

Flags for Displaying On-screen Help for Agent Commands

```
MIRAgent.exe {-i | -u | -d | -o | -regencert} -?
```

A.1.1. Local Audit Flags

```
MIRAgent.exe -o [basedir] [-f] [-script filename] [-encoding {gzip | aff | none}] [-notimestamp] [-auditissuefilterlevel {debug | info | warning | error}] [-allowmultiple] [-portable] [-cleanup]
```

-allowmultiple

Allows multiple Agents to run on the same system. Each Agent must be bound to a different port using the `-bind` notation.

-auditissuefilterlevel debug | info | warning | error

Filter audit issues at the specified level. Default: `info`.

-cleanup

Allows the Agent to clean up any artifacts on the Host system that were created when an Audit was executed. For example, executing a memory Audit installs a driver on the Host machine. Specifying this parameter removes the driver when the Agent has finished executing.

-encoding gzip | aff | none

Sets the compression method for storing or transmitting Audits. Default: `gzip`.

-f

Forces the creation of the directory specified by the `-o basedir` argument.

-notimestamp

Do not add a timestamp to the output directory tree.

-portable

Installs configuration files in the local application folder, leaving no trace on the Host computer after audit completion.

-o [basedir]

Instructs the Agent to Audit the system on which it is installed, and to store the results locally. If `basedir` is not specified, the results are stored in the current working directory.

Audit files are stored in `basedir/Audits/machine/date`, where `machine` is the name of the system on which the Agent is running, and `date` is a date-time stamp in the form `YYMMDDHHMMSS`.

-script filename

Executes the specified Agent command script file.

A.1.1.1. Examples

```
miragent.exe -o -script scriptfile.xml
```

Executes an Agent command file and writes the script output to `.\hostname\date`.

```
miragent.exe -o c:\temp -script scriptfile.xml
```

Executes an Agent command file and *if the C:\temp directory exists*, writes the script output to `C:\temp\hostname\date`.

```
miragent.exe -o c:\audits -f -script scriptfile.xml
```

Executes an Agent command file and writes the script output to `C:\temp\hostname\date`. If the `C:\audits` directory does not exist, it will be created.

```
miragent.exe -o -script scriptfile.xml -encoding none
```

Executes an Agent command file and writes the script output to `.\hostname\date`. . . The script output is not compressed.

```
miragent.exe -o -script scriptfile.xml -cleanup
```

Executes an Agent command file and writes the script output to `.\hostname\date`. . . After the Agent has completed its tasks, it removes its artifacts from the system.

A.1.2. Daemon Mode Flags

```
MIRAgent.exe -d [-bind ip address[:port number]] [-allow network/mask,...] [-settings _filename] [-fw {on | service | auto | off}] [-allowmultiple] [-certpath path] [-keypath path] [-capath path] [-crlpath path] [-credspath path] [-memcert] [-disablediscovery] [-auditissuefilterlevel {debug | info | warning | error}] [-encoding {gzip | aff | none}] [-portable] [-cleanup]
```

-allow *network/mask,...*

Specifies an IP Allow List using a comma-separated list of network addresses and netmasks in CIDR (Classless Internet Domain Routing) notation. By default the Agent allows connections from all IP addresses. When `-allow` is set, only connections from the specified networks or systems are permitted.

-allowmultiple

Allows multiple Agents to run on the same system. Each Agent must be bound to a different port using the `-bind` notation.

-f

Forces the creation of the directory specified by the `-o basedir` argument.

-bind *ip address[:port number]*

Sets the IP address and TCP port number that the Agent will listen on. By default, the Agent listens on TCP port 22201.

-capath *path*

Tells the Agent to use the public certificate file located at the specified directory. This overrides any settings file specified and is saved in that settings file. By default the value is `miragentcert.pem.protected` and is found in the current working directory.

-cleanup

Allows the Agent to clean up any artifacts on the Host system that were created when an Audit was executed. For example, executing a memory Audit installs a driver on the Host machine. Specifying this parameter removes the driver when the Agent has finished executing.

-credspath *path*

Tells the Agent to use the credentials file located at the specified directory. This overrides any settings file specified and is saved in that settings file. By default the value is `creds.protected` and is found in the current working directory.

-certpath *path*

Tells the Agent to use the certificate file located at the specified directory. This overrides any settings file specified and is saved in that settings file. By default the value is `cert.protected` and is found in the current working directory.

-crlpath *path*

Tells the Agent to use the CRL (Certificate Revocation List) file located at the specified directory. This overrides any specified settings file and is saved in that settings file. By default the value is `mircrl.pem.protected` and is found in the current working directory.

-d

Instructs the Agent to begin listening in daemon mode. By default the Agent listens on TCP port 22201. `-d` is often used with `-bind` to set the listening port; and with `-allow` to restrict IP addresses.

-discovery

Disables the Discovery service. When the service is disabled, the Agent will not attempt to contact the Controller.

-encoding gzip | aff | none

Sets the compression method for storing or transmitting Audits. Default: `gzip`.

-fw on | service | auto | off

Automatically configures the Windows firewall to allow connections to the Agent.

- `on` adds a persistent exception for the Agent.
- `service` allows connections only when the Agent is running as a service.
- `auto` allows connections to the Agent while it is running in daemon mode and will remove the exception when the Agent exits.
- `off` prevents the Agent from setting any exceptions itself. You will need to manually set firewall exceptions in this case.

-keypath *path*

Tells the Agent to use the private key file located at the specified directory. This overrides any settings file specified and is saved in that settings file. By default the value is `miragentkey.pem.protected` and is found in the current working directory.

-memcert

Tells the Agent to use a certificate generated in memory only. Each time the Agent is started a new certificate is generated and remains in memory as long as the Agent is active.

-portable

Installs configuration files in the local application folder, leaving no trace on the Host computer after audit completion.

-settings *filename*

Instructs the Agent to use a specific settings file. If *filename* does not exist, the Agent stores its current settings in that file. By default, Agents do not store settings.

This flag needs to be used at least once with a default set of flags in order for a settings file to be created.

A.1.2.1. Examples

```
miragent.exe -d -bind 127.0.0.1:22222
```

Runs the Agent as a daemon listening local connections on port 22222.

```
miragent.exe -d -bind 127.0.0.1:22222
```

Runs the Agent as a daemon listening to all connections on port 22222.

miragent.exe -d -encoding gzip

Runs the Agent as a daemon, using gzip compression on its data files.

miragent.exe -d -memcert

Runs the Agent as a daemon, using SSL certificates generated in memory.

miragent.exe -d -cleanup

Runs the Agent as a daemon, removing any artifacts when the daemon exits.

A.1.3. Service Mode Flags

MIRAgent.exe -i [-start] [-servicename *name*] [-servicedisplay *displayname*] [-settings *filename*]

MIRAgent.exe -u [-servicename *name*]

-i

Installs the Agent as a Windows service.

-servicename *name*

Used with **-i**, sets the short name of the Agent service when it is installed.

-servicedisplay *displayname*

Used with **-i**, sets the display name of the Agent service when it is installed.

-start

Instructs the Agent to start when it has been previously installed as a service. The **-start** flag is only available with **-i** to both install and start the Agent with one command.

-u

Stops and uninstalls the Agent if it is running as a service.

A.1.3.1. Examples

miragent.exe -i -servicename MyShortName

Runs the Agent as a service with the short name *MyShortName*.

miragent.exe -i -servicename MyShortName -servicedisplay MyDisplayName

Runs the Agent as a service with the short name *MyShortName* and the display name *MyDisplayName*.

miragent.exe -i -start

Runs the Agent as a service and then starts the service.

miragent.exe -u -servicename MyShortName

Stops and uninstalls the Agent that has the short name *MyShortName*.

A.1.4. Special Usage

A.1.4.1. Start Agent as a Service Using a Settings File

1. Start the Agent in daemon mode, specifying a new settings file:

```
miragent -d -settings mysettingsfile.xml
```

2. Stop the daemon by tapping **Enter**.
3. Install the Agent as a service and start it, using the new settings file:

```
miragent -i -start -settings mysettingsfile.xml
```

A.1.4.2. Install the Agent onto a Removable Media Device and Execute the Agent on a New Host

1. Install the Agent on a 32-bit or 64-bit host to create a 32-bit or 64-bit portable Agent, respectively:

```
msiexec /quiet /i C:\AgentSetup.msi PORTABLE=True TARGETDIR=C:\MIRAgent COMMONAPPDATAFOLDER=C:\MIRAgent
```

2. Copy the installation to the removable media device:

```
copy C:\MIRAgent E:\MIRAgent
```

3. Install the removable media device at the target machine and run the Agent:

```
cd e:\MIRAgent miragent -portable -capath mircacert.pem -crlpath mircrl.pem -fw auto
```

or

```
cd e:\MIRAgent miragent -portable -memcert -settings service.settings.xml
```

or

```
cd e:\MIRAgent miragent -portable -memcert -capath mircacrt.pem -crlpath mircrl.pem -fw auto
```

A.1.4.3. Install the Agent and Burn a CD

1. Install the Agent on the device, making sure to select the **Portable Install** option.
2. Burn the installation folder to a CD
3. On the target machine, run the Agent using the following command:

```
miragent -d -memcert -capath mircacert.pem -crlpath mircrl.pem -fw auto
```

A.1.5. A Note on Default Behaviors

Agent commands have default behaviors as follows:

Interface and Port Binding

The Agent binds to all available network interfaces unless a more specific binding is explicitly specified on the command line or in a settings file. The default port is TCP 22201.

Firewall Exceptions

Firewall exceptions are set while the Agent is active as a service. The exception is added when an Agent is installed.

Compression

All audits are collected and packaged using gzip unless otherwise specified.

Output Base

Directory audits collected with the `-o` flag are stored in `.\Audits` according to Host name and collection date.



Appendix B

Error Messages and Troubleshooting

Enterprise environments can be very complex and enterprise security software is in turn often complex in order to meet some of the challenges posed by those environments. Failures and errors can occur for a variety of reasons. Clearly communicating these issues is an important aspect of enterprise products.

This appendix outlines the error reporting infrastructure contained in MIR, and provides guidance on identifying and resolving problems you may encounter.

B.1. Errors, Issues, and Logs

MIR uses three concepts to report information about anomalies in the operation of the system: error notification, issues documents, and log files.

Error notifications are typically communicated to you via a notification in a user interface, such as the Console or the web-based Administration Console. Issues Documents are generated as part of an audit, and report anomalies, failures, and other problems at each level in the process of collecting information from a MIR Agent. Log files capture information about the operation of a system component and may be useful in diagnosing errors and failures.

B.1.1. Errors

The Console and Administration Console directly report errors encountered during application operation. These errors are typically caused by direct failures, such as Console's inability to contact the Controller due to a network problem, or the failure of an administrative change to take effect due to insufficient user privileges. In these instances the error is reported directly to you via a message in the user interface or a dialog box.

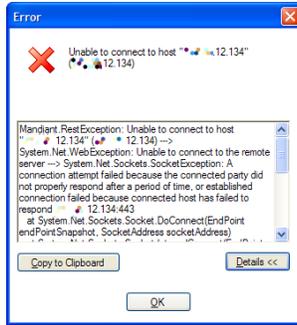
If an error is non-fatal, the Console will ask you to acknowledge the error before continuing. If the error is fatal, you will be asked to acknowledge the error; the Console is then shut down.

B.1.1.1. Console Errors

The Console provides three methods for communicating anomalies: error dialogs, activity status notifications, and warnings.

Error Dialogs

An error dialog is displayed when a problem is encountered that requires notification and acknowledgement. A summary of the problem is shown, along with buttons to display further information, copy the error message to the clipboard, and to acknowledge the error (and thus dismiss the dialog). Some errors are fatal and will prevent further operation of the Console until remedied. Examples of issues that will display error dialogs include connection failures and authentication failures.



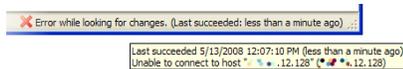
Download Status Notifications

The Console includes a **Downloads** Pane that reports on the status of long running operations, such as data import and export activities. If these activities change to an unexpected state (for example, if a long running import or client script is cancelled by the user) a notification is reported by turning the status bar for that activity red. Status text is also displayed in the status bar at the bottom of the Console window. Hover your mouse pointer over the status bar to see complete message text.



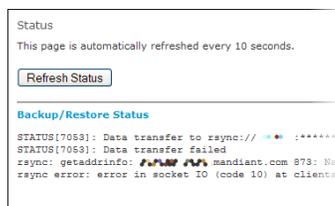
Warnings

Warnings are important information that does not require immediate attention, but may indicate an anomaly or other potential error within the system. Warnings are typically displayed in the Console status bar. For example, if there is an error when the Console is communicating with the Controller, that is communicated to the user as a warning. While it is an error condition, it is not fatal: you can continue to view the Console, but the information may not be completely up-to-date.



B.1.1.2. Administration Console Errors

The Administration Console displays errors, status, progress, and processing messages in the main page frame. When errors are encountered they are generally reported directly from the subsystem that encountered the error.



B.1.2. Issues Documents

Issues Documents are created as part of running a Host Audit, and are included as part of the Audit Results for a given Job. Issues report errors, anomalies, warnings, or informational items that were encountered as a result of attempting to collect the information specified in a Job script. These issues could occur at two levels in the system:

Agent Module

Issues generated by the auditor. They might include operating system level information (info, warning, or errors) which help to describe why some information may not have been retrieved. Errors at this level are often not fatal and usually do not cause the overall script to fail.

Audit Execution

Issues at this level are sometimes fatal, indicating that the overall script failed to run. This could be due to a number of reasons: the script failed to parse, the Controller couldn't reach the Agent, there was authentication failure between Controller and Agent, the Agent failed to parse a Job Script, or etceteras.

Items within an issues document have the following format:

Level Message Code Context

Level

The severity of the issue item (e.g. INFO, FATAL)

Message

The verbatim message describing the error. This message may be directly generated by MIR, or it may be generated by the target host operating system and placed directly in the issue item.

Code

A numeric value that may be useful for MIR Customer Support in diagnosing your issue.

Context

A string that may or may not be populated in the issue item. When present it helps identify where in the MIR software the issue was encountered.

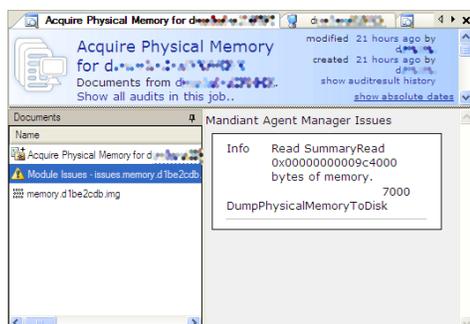
MIR defines two types of documents to capture these events: module issues and batch result issues, as detailed below.

B.1.2.1. Module Issues

Module issues contain a variety of items, often reported by the Host operating system as the MIR Agent collects information. Some fatal errors that interrupt the execution of an Audit Module are reported here; more often you will see informational or non-fatal items

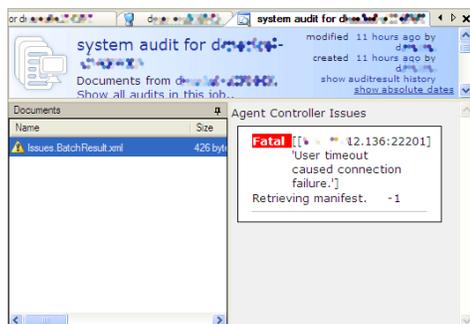
reported as a result of attempting to access a specific piece of information. For example, when retrieving a file listing, errors encountered on individual file listing line-items would be non-fatal, but would be recorded as issues. The file listing action might complete, but the Issues Document for that module would report why you may not have retrieved all of the expected information.

Module issues are displayed as a document in an Audit Result, and are named `Module Issues - [module result document].xml`. The screenshot below illustrates an Issues Document for a *Memory Acquisition* Module. Note the issue reported is a non-fatal informational item.



B.1.2.2. Batch Result Issues

Batch result issues record information that impact the entire Audit Result. Some of these may be informational, but they are more commonly fatal. Batch result issues are displayed as a document in an Audit Result and are named `Issues.BatchResult.xml`. There will only be one batch result issues document per audit result. The figure below illustrates a batch result issues document for an Audit where the Controller could not contact the Agent.



B.1.3. Logs

Every component of MIR creates one or more log files to record activity, errors, or other information that may be useful in auditing or troubleshooting system activities. This section identifies available log files by component and their use. With the exception of Controller activity logs, most logs are primarily used for troubleshooting errors and working with MANDIANT Customer Support.

B.1.3.1. Agent Logs

The Agent log contains informational messages about connections from Controllers, attempts to contact Agent Discovery Service, and other standard operating messages. If errors or other

problems are encountered, they are also written to the log. The log file can be retrieved from the Agent remotely via a *File Acquisition Audit*.

The location of this file varies dependent on the Host OS:

Windows 7 and Windows Vista

```
%SYSTEMDRIVE%\ProgramData\MANDIANT\MANDIANT Intelligent
Response Agent\MIRAgent.log
```

Windows XP and 2K3

```
%SYSTEMDRIVE%\Documents and Settings\All Users\Application Data
\MANDIANT\MANDIANT Intelligent Response Agent\MIRAgent.log
```

Windows 2K

```
%SYSTEMDRIVE%\Documents and Settings\All Users.WINNT
\Application Data\MANDIANT\MANDIANT Intelligent Response Agent
\MIRAgent.log
```

The log file has the following format:

```
Date Time [ThreadID] LogLevel [Context] Message
```

```
05-14-2008 17:27:09 [0x000006c8] INFO [Win32Service
AgentManagerService HttpdAgentManager Discovery]- The Discovery
Server settings file 'discovery.xml' has been loaded.
05-14-2008 17:27:09 [0x000006c8] INFO [Win32Service
AgentManagerService HttpdAgentManager Discovery]- The Discovery
service has started.
05-14-2008 17:27:09 [0x00000ad4] INFO [Discovery]- Starting
Discovery request.
05-14-2008 17:29:30 [0x00000ac0] INFO [HttpdAgentManager
shttpd]- Connection from 172.16.12.128:56158 on socket 352
05-14-2008 17:29:31 [0x00000ac0] INFO [HttpdAgentManager]-
Verifying client certificate...
V0
AC71725DBF21E0EE
sha1WithRSAEncryption
/CN=MIR_CA/O=Mandiant/C=US
May 14 09:26:50 2008 GMT
May 12 09:26:50 2018 GMT
/CN=MIRApp/O=Mandiant/C=US
05-14-2008 17:29:31 [0x00000ac0] INFO [HttpdAgentManager]- ...
client certificate verified.
```

B.1.3.2. Controller Logs

The Controller maintains logs for every application service that provides MIR functionality. Logs are directly accessible via the Administration Console or exportable from the system via **syslog-ng**.

Console logs are standard syslog-formatted log files stored in plain text. Log messages are generally of the format:

```
Date Time LogLevel [Context] Message
```

The following table outlines the MIR Controller services and their function:

mir_agent	Logs Discovery data to the audit files.
mir_agent_controller	Controls all interaction with deployed Agents.
mir_agent_dispatcher	Controls dispatch of tasks to subsystems responsible for interacting with deployed Agents.
mir_agent_upgrade_proxy	Controls the Agent over-the-network (Discovery) upgrade process.
mir_analyzer_dispatcher	Dispatches analysis jobs to associated components.
mir_analyzer_service	Conducts analysis tasks.mir_auditManages the system activity audit log.
mir_data	The data management service responsible for organizing and storing information in associated databases and on disk.
mir_discovery_server	The Agent Discovery server.mir_discovery_serviceThe Agent Discovery Service. Registers new Agents and updates their records as information (such as network address) changes.
mir_indexer	Indexes acquired data so it can be made available to mir_searcher.
mir_lcd	Controls the LCD panel on the front of the Controller (deprecated).
mir_mbus	The back-end message bus that transmits messages between MIR Controller components.
mir_pound	Web caching service for the Controller.
mir_restore_web	Performs Controller backup and restore tasks.
mir_scheduler	Schedules and executes tasks within the system.
mir_script_runner	Parses Job Scripts and controls dispatch to either the Analysis engine or Agent management subsystem.
mir_searcher	Controls MIR's search engine.
mir_searcher_dispatcher	Controls dispatch of tasks to subsystems responsible for resolving search requests.
mir_web	The primary web service that Consoles interact with, plus the log files for mir_web_admin, mir_web_files, and mir_web_static.
mir_web_admin	Provides the admin UI interface.
mir_web_files	Provides file transfer services.
mir_web_static	Serves static content for the mir_web service.
mir_discovery_proxy	Proxy server for discovery.

B.1.3.3. Console Logs

The location of Console logs varies dependent on the Host OS:

Windows 7 and Windows Vista

Crash Logs	%SYSTEMDRIVE%\Users\[username]\AppData \Local\MANDIANT Corporation\MANDIANT
------------	--

	Intelligent Response\[Console version]\[*.log]
Error Messages	%SYSTEMDRIVE%\Users\[username]\AppData\Roaming\MANDIANT Corporation\MANDIANT Intelligent Response\Mirconsole.apperrors.log
Log Messages	%SYSTEMDRIVE%\Users\[username]\AppData\Roaming\MANDIANT Corporation\MANDIANT Intelligent Response\Mirconsole.log
User Settings	%SYSTEMDRIVE%\Users\[username]\AppData\Local\MANDIANT_Corporation\mirconsole.exe_StrongName_[hash]\[Console version]

Windows XP, Windows 2K3, Windows 2K

Crash Logs	%SYSTEMDRIVE%\Documents and Settings\[username]\Local Settings\Application Data\MANDIANT Corporation\MANDIANT Intelligent Response\[*.log]
	%SYSTEMDRIVE%\Documents and Settings\All Users\Application Data\MANDIANT Corporation\MANDIANT Intelligent Response\[*.log]
Error Messages	%SYSTEMDRIVE%\Documents and Settings\[username]\Application Data\MANDIANT Corporation\MANDIANT Intelligent Response\Mirconsole.apperrors.log
Log Messages	%SYSTEMDRIVE%\Documents and Settings\[username]\Application Data\MANDIANT Corporation\MANDIANT Intelligent Response\Mirconsole.log
User Settings	%SYSTEMDRIVE%\Documents and Settings\[username]\Local Settings\Application Data\MANDIANT_Corporation\user.config

Like most other log files in the system, the Console logs follow the format:

Date Time LogLevel [Context] Message

The Console writes three logs into the directory it was installed into:

Console.log

Records general errors, warnings, and information about Console operation. Used for general Console troubleshooting.

Console.History.log

Records full details of any exception errors that interrupt Console operation. Used for troubleshooting of a specific exception or error.

ChangeDetection.log

Records change notifications received by the Console. Used for troubleshooting collaboration and view update issues.

B.2. System Reports

MCIC provides a robust MIR system evaluation report tool, available to Administrators only, through the command-line interface or through MCIC. This tool queries the controller, compiling a package of configuration, logs, and state information that helps MANDIANT Product Support to identify issues.



MCIC provides Administrator-only access to MIR System Eval reports through **Controller** → **System Reports**. Please see the MIR User Guide (Chapter *MCIC User Interface*) for details.

1. Download the latest revision of the `MIRSystemEval` script from

```
https://forums.mandiant.com/topic/now-presenting-the-mirsystemeval-script#post-1522
```

2. Copy the script to each Controller, set execute permissions, and run the script:

```
cp source/MIRSystemEval.sh ~/MIRSystemEval.sh
sudo chmod 755 MIRSystemEval.sh
sudo bash ./MIRSystemEval.sh -l -pd
```

The script will create an archive named `MIRSystemEvalReport.hostname.date.tgz`.

3. Send an email to `MIRsupport@MANDIANT.com` to request instructions for uploading the archive. This file is normally far too large to send through email. MANDIANT runs a drop-box at `transit.MANDIANT.com`, and typically receives the archives through that service.
4. MANDIANT Product Support will contact you to discuss the issues you have encountered and to identify fixes or work-arounds that will resolve them.

The `MIRSystemEval.sh` script creates a `.tgz` archive containing:

BashHistory.txt

Bash command history for root and users.

BasicInfo.txt

A snapshot of the controller and its settings.

cksum_system.txt

Checksums of system components.

dmesg.txt

Boot information captured during the last reboot.

logs.tgz

Various MIR system logs.

ls_IR_of_opt.txt

Lists the contents of /opt.

map.tgz

Various MCIC system logs.

mcic_server_install.log

The MCIC installation/upgrades log.

merge_timings.txt

Lists indexer merge timings.

messages.tgz

Various MCIC system messages.

perf_files.txt

Performance timing statistics.

perf_times.txt

Performance timing statistics.

ps_aef.txt

Output from the `ps` command.

webclient.tgz

Various MCIC system logs.

In addition to the contents of the archive, Product Support will also ask for the following information:

- Version numbers for the Controller, Console, Agents, and MCIC.
- Whether any Controllers are configured for Single Discovery.
- Whether any Controllers are configured for Sniping.
- A description of custom or modified IOCs.

In the rare case when log collection (`logs.tgz`, `messages.tgz`, `webclient.tgz`) is not desired, you may use the `--nologs` option:

```
MIRSystemEval.sh
```

NAME

```
MIRSystemEval -- create an archive package of MANDIANT MIR Controller configuration, logs, and state information.
```

SYNOPSIS

```
MIRSystemEval [-lp] [-u username] [-g groupname] [-o path]
```

DESCRIPTION

```
The MIRSystemEval utility creates a tarball archive containing a variety of MIR Controller troubleshooting information. By default, it collects system logs that are often extremely large.
```

```
The following options are available:
```

```
-l
  Collect logs.

-P
  Collect performance statistics.

-u username
  Set the owner of the archive to the specified user name.

-g groupname
  Set the group owner of the archive to the specified group name.

-o path
  Save the archive to the specified directory path.
```

EXIT STATUS

The MIRSystemEval utility exits with one of the following values:

```
0 The tarball archive was created.
1 An error occurred.
```



Appendix C

Credits

In addition to the open source projects listed below, MANDIANT Intelligent Response contains technologies from the Sleuthkit Project and the public packer database provided by Bobsoft and the PEiD Project. MANDIANT wishes to thank Bobsoft, the PEiD Project team, and Brian Carrier for their excellent work. Learn more about them at:

Sleuthkit

<http://www.sleuthkit.org>

PEiD Public Packer Database

<http://www.peid.info>

C.1. Component License Notices

MANDIANT Intelligent Response makes use of several Open Source and Shared Source technologies. The following license declarations are included in compliance with the licenses for those components. Note that these declarations govern MANDIANT's use of these technologies. Your purchase and use of MIR is governed solely by the *MIR End User License Agreement* (EULA) and software purchase agreement, which comply with the notices contained herein.

C.1.1. Be.HexEditor

Project URL. *<http://www.gotdotnet.com>*

License. This license governs use of the accompanying software ("Software"), and your use of the Software constitutes acceptance of this license. You may use the Software for any commercial or noncommercial purpose, including distributing derivative works. In return, we simply require that you agree:

1. Not to remove any copyright or other notices from the Software.
2. That if you distribute the Software in source code form you do so only under this license (i.e. you must include a complete copy of this license with your distribution), and if you distribute the Software solely in object form you only do so under a license that complies with this license.
3. That the Software comes "as-is", with no warranties. None whatsoever. This means no express, implied or statutory warranty, including without limitation, warranties of merchantability or fitness for a particular purpose or any warranty of title or non-infringement. Also, you must pass this disclaimer on whenever you distribute the Software or derivative works.
4. That no contributor to the Software will be liable for any of those types of damages known as indirect, special, consequential, or incidental related to the Software or this license, to the maximum extent the law permits, no matter what legal theory it is based on. Also, you must pass this limitation of liability on whenever you distribute the Software or derivative works.

5. That if you sue anyone over patents that you think may apply to the Software for a person's use of the Software, your license to the Software ends automatically.
6. That the patent rights, if any, granted in this license only apply to the Software, not to any derivative works you make.
7. That the Software is subject to U.S. export jurisdiction at the time it is licensed to you, and it may be subject to additional export or import laws in other places. You agree to comply with all such laws and regulations that may apply to the Software after delivery of the software to you.
8. That if you are an agency of the U.S. Government, (i) Software provided pursuant to a solicitation issued on or after December 1, 1995, is provided with the commercial license rights set forth in this license, and (ii) Software provided pursuant to a solicitation issued prior to December 1, 1995, is provided with "Restricted Rights" as set forth in FAR, 48 C.F.R. 52.227-14 (June 1987) or DFAR, 48 C.F.R. 252.227-7013 (Oct 1988), as applicable.
9. That your rights under this License end automatically if you breach it in any way.
10. That all rights not expressly granted to you in this license are reserved.

C.1.2. libcurl

Project URL. <http://curl.haxx.se>

License. COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996 - 2007, Daniel Stenberg, <<daniel@haxx.se>>.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

C.1.3. libxml2, libxslt, libexslt

Project URL. <http://xmlsoft.org>

License. License (libxml2) Except where otherwise noted in the source code (e.g. the files hash.c, list.c and the trio files, which are covered by a similar licence but with different Copyright notices) all the files are:

Copyright © 1998-2003 Daniel Veillard. 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", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE DANIEL VEILLARD BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Daniel Veillard shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from him.

License (libxslt)Copyright © 2001-2002 Daniel Veillard. 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", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE DANIEL VEILLARD BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Daniel Veillard shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from him.

License (libexslt)Copyright © 2001-2002 Thomas Broyer, Charlie Bozeman and Daniel Veillard. 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", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of the authors shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from him.

C.1.4. lxml

Project URL. <http://codes.peak.net/lxml>

License. lxml is copyright Infracore and distributed under the BSD license (see doc/licenses/BSD.txt), with the following exceptions:

Some code, such as selftest.py, selftest2.py and src/lxml/_elementpath.py are derived from ElementTree and cElementTree. See doc/licenses/elementtree.txt for the license text.

test.py, the test-runner script, is GPL and copyright Shuttleworth Foundation. See doc/licenses/GPL.txt. It is believed the unchanged inclusion of test.py to run the unit test suite falls under the "aggregation" clause of the GPL and thus does not affect the license of the rest of the package.

The doctest.py module is taken from the Python library and falls under the PSF Python License.

Copyright (c) 2004 Infracore. 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. Neither the name of Infracore 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 INFRACORE 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 ElementTree/XML Toys Library

Copyright (c) 1999-2003 by Secret Labs AB

Copyright (c) 1999-2003 by Fredrik Lundh

By obtaining, using, and/or copying this software and/or its associated documentation, you agree that you have read, understood, and will comply with the following terms and conditions:

Permission to use, copy, modify, and distribute this software and its associated 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 Secret Labs AB or the author not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

SECRET LABS AB AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL SECRET LABS AB OR THE AUTHOR BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM 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.

C.1.5. log4net, Lucene

Project URLs. <http://logging.apache.org/log4net/>

<http://lucene.apache.org>

License. Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of

this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the

Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

You must give any other recipients of the Work or Derivative Works a copy of this License; and

You must cause any modified files to carry prominent notices stating that You changed the files; and

You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied,

including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

C.1.6. OpenSSL

Project URL. <http://www.openssl.org>

License. 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 © 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 adhered 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.]

C.1.7. PyLucene

Project URL. <http://pylucene.osafoundation.org>

License. Copyright (c) 2004 - 2005 Open Source Applications Foundation. 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", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

C.1.8. SQLAlchemy

Project URL. <http://www.sqlalchemy.org>

License. This is the MIT license: <http://www.opensource.org/licenses/mit-license.php>

Copyright (c) 2005, 2006, 2007 Michael Bayer and contributors. SQLAlchemy is a trademark of Michael Bayer.

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", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

C.1.9. Twisted

Project URL. <http://twistedmatrix.com>

License. Copyright (c) 2001-2006

Allen Short, Andrew Bennetts, Apple Computer, Inc., Benjamin Bruheim, Bob Ippolito, Canonical Limited, Christopher Armstrong, David Reid, Donovan Preston, Eric Mangold, Itamar Shtull-Trauring, James Knight, Jason A. Mobarak, Jonathan Lange, Jonathan D. Simms, Jp Calderone, Jürgen Hermann, Kevin Turner, Mary Gardiner, Matthew Lefkowitz, Massachusetts Institute of Technology, Moshe Zadka, Paul Swartz, Pavel Pergamenschik, Ralph Meijer, Sean Riley, Travis B. Hartwell

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", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

C.1.10. zlib

Project URL. <http://www.zlib.net>

License. Copyright © 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>

C.1.11. nginx

Project URL. <http://nginx.org/>

License. Copyright © 2002-2010 Igor Sysoev

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.

THIS SOFTWARE IS PROVIDED BY AUTHOR 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 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.

C.1.12. PCRE

Project URL. <http://www.pcre.org/>

License. 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 University of Cambridge nor the name of Google Inc. 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 OWNER 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.

C.1.13. `ipaddr`

Project URL. <http://code.google.com/p/ipaddr-py/>

License. Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

You must give any other recipients of the Work or Derivative Works a copy of this License; and

You must cause any modified files to carry prominent notices stating that You changed the files; and

You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

C.1.14. commons-codec

Project URL. <http://commons.apache.org/codec/>

License. Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

You must give any other recipients of the Work or Derivative Works a copy of this License; and

You must cause any modified files to carry prominent notices stating that You changed the files; and

You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages

of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

C.1.15. google-gson

Project URL. <http://code.google.com/p/google-gson/>

License. Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

You must give any other recipients of the Work or Derivative Works a copy of this License;
and

You must cause any modified files to carry prominent notices stating that You changed the files; and

You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer

failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

Index

A

- Agents
 - auto upgrade, 98
 - deployment, 85
 - installation, 21

C

- Certificates, 17
 - backing-up, 20
 - generating CRLs, 18
 - generating signed certificates, 18
- Configuration
 - admin account, 13
 - CRL, 18
 - NTP, 15
 - server certificate, 18
 - server keys, 18
 - SSL, 17
 - TDCA certificate, 17
 - TDCA keys, 17
- Controller
 - backup status, 33
 - cleanup logs, 35
 - configuration
 - backups, 32, 32, 34
 - CEF logging, 37
 - dhcp, 40
 - log files, 35, 35
 - log rotation, 36
 - network settings, 38
 - ntp, 37
 - restore from backup, 33
 - sshd maintenance service, 41
 - static address, 38
 - Syslog NG, 36
 - time and date, 41
 - diagnostics
 - ping, 42
 - traceroute, 42
 - DVD upgrade, 45
 - ISO upgrade, 45
 - packages, 42
 - patches, 43
 - reboot, 44
 - restarting MIR processes, 19
 - restoring from backup, 33
 - shutdown, 44

- statistics, 44
- Customer Support
 - see Support, 2

I

- Installation
 - Controller, 7

K

- Keys
 - generating server keys, 18
 - generating TDCA, 17

L

- Log files
 - viewing, 48

M

- MANDIANT Intelligent Response
 - architecture, 4
 - components, 4
 - requirements
 - environment, 6
 - firewall, 6
 - general, 5
 - network, 6
- MCIC
 - database
 - reset, 56
- MIR
 - agent installation, 89
 - configuration file, 91
 - portable agent, 98
 - types of, 90
 - using msixexec, 94
 - using windows installer, 94
 - agents
 - installation, 90
 - portable, 98
 - regarding security products, 100
 - signatures, 100
 - uninstall, 99
 - Agents
 - command-line reference, 101
 - audit logging, 66
 - configuration, 50
 - database
 - maintenance, 54
 - mcic reset, 56
 - reset, 55
 - statistics, 55

- diagnostics
 - system report, 115
- discovery, 55
- discovery query, 88
- discovery service
 - administration, 88
 - bulk load, 88
 - configuration, 87
 - duplicate agent detection, 89
 - overview, 85
 - requirements, 86
 - reset, 89
 - search, 89
- documents, 58
- error reporting, 108
- issues documents, 110
- logs, 111
- managing groups, 65
- queues, 59
- remote authorization, 56
- resources, 57
- restarting processes, 48
- see MANDIANT Intelligent Response, 4
- SSL, 62
- system report, 115
- trust domain
 - backup and restore, 81
 - managing certificates, 79
 - managing CRLs, 80
 - managing TDCA, 76
 - overview, 69
 - user management, 63
 - viewing log files, 48
- MIR Administration Console
 - first run, 24
 - installing, 22
 - logging in, 25
- MIRAgent
 - command line, 101
- P**
- Processes
 - restarting, 19, 48
 - status, 48
- Product Support
 - see Support, 2
- S**
- SSL
 - configuring, 17
- Status
 - processes, 48
- Support
 - books and articles, 3
 - downloads, 1
 - email, 2
 - forums, 2
 - requirements, 2
 - telephone, 2

MANDIANT CORPORATION
WWW.MANDIANT.COM

© 2012, MANDIANT Corporation. All rights reserved.