CA DataMinder

Platform Deployment Guide

Release 14.5

CA Technologies Product References

This document references the following CA Technologies products:

- CA DataMinder
- CA Identity Manager

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Contents

Chapter 4: Clustered CMS Deployments 59

Chapter 5: Before You Start using CA DataMinder 71

Chapter 6: Advanced Encryption Mode 87

Chapter 7: iConsole Deployment 105

Chapter 10: Quarantine Manager 225

Chapter 11: Account Import 245

Chapter 12: Object Storage 307

Chapter 15: Policy Engine Hub Advanced Configuration 363

Chapter 16: Content Registration 389

Chapter 17: Content Services 403

Chapter 18: EML Conversion Utilities: Cnv2email and BB2email 435

Chapter 19: Binary Text Extractor 457

Chapter 20: Universal Extractor 471

Chapter 21: Mapping Events to Users 485

Chapter 22: Backing Up and Restoring the CMS 499

Chapter 23: Log Files 503

Chapter 24: Technical Information 523

Chapter 1: Introduction

This section outlines the recommended methods for installing and deploying CA DataMinder.

This section contains the following topics:

CA Support

To contact CA Support, go to:

http://ca.com/support

If you do contact CA Technical, they may ask you to supply:

- The infrastructure log file, wgninfra.out.
- Any relevant system log files. These take the format: stderr_200201200945.log.

These are all located in CA's \data\log subfolder of the Windows All Users profile; see Where are log files saved? (see page 510)

Deploying CA DataMinder

After installing CA DataMinder on your CMS, some further configuration is needed before continuing with the deployment. When this is complete, you can install CA DataMinder on as many gateways and client machines as your license agreement permits.

We also recommend that you make a full backup of your CMS at least once per week, and incremental backups on a daily basis.

All of these tasks are described in the following chapters. Known deployment issues are also covered in Technical information.

Database Tasks

Before installing the CMS or a gateway, your chosen database engine must already be installed and correctly configured on the target server. CA DataMinder currently supports Oracle and SQL Server database engines. For configuration guidelines, plus details about purging and backing up your database, see the *Database Guide*.

Architecture

CA DataMinder machines are organized into hierarchical branches, with the central management server (CMS) as the top level server. The CMS acts as a central repository for all policy details and captured data. Below the CMS, each branch of the hierarchy is optionally managed by a gateway, and each gateway can serve multiple client machines.

Database changes on the CMS are copied automatically via gateway servers to local databases on client machines. These changes can include policy updates and modifications to user and machine accounts. Similarly, at intervals defined in the machine policy, captured data and local policy changes are copied automatically up to the CMS, again via gateway servers.

You manage CA DataMinder using a console. You can deploy consoles on any machine in your CA DataMinder installation. You can also have console-only installations, that is, you can install the console without installing any CA DataMinder server software or client integration features—see Console-only installations.

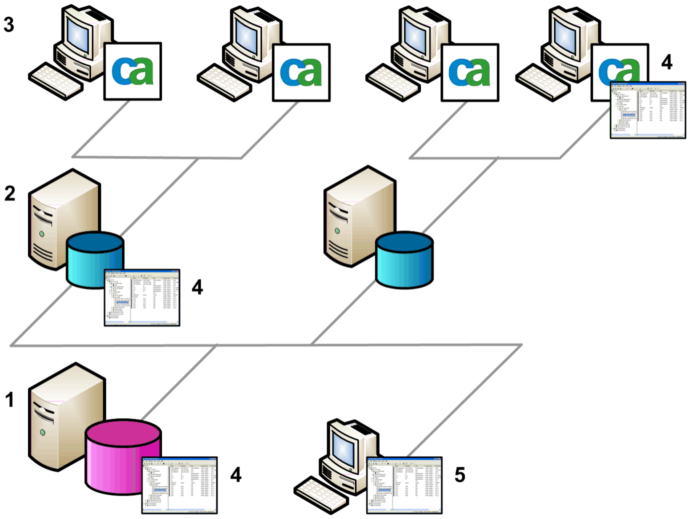

Example machine hierarchy

CA DataMinder machines are organized in a virtual hierarchy. This need not correspond to your actual network topology.

1 CMS.

2 Gateway servers.

3 Client machines.

4 Consoles on CA DataMinder machines.

5 Console-only machine.

Upgrading CA DataMinder

For details about upgrading CMSs, gateway servers, and client machines to the latest version of CA DataMinder, see the *Upgrade Guide*. This guide highlights the essential issues you need to be aware of when rolling out an upgrade across your organization and describes the necessary post-upgrade tasks.

Version Compatibility

Ideally, when upgrading individual CA DataMinder machines we recommend that you upgrade all your CA DataMinder servers (gateways, Event Import machines, policy engines, and so on) at the same time as the CMS. Do this during a period of minimal user activity, beginning with the CMS and then working down the machine hierarchy to each server successively. When this is complete, we also recommend that you upgrade your client machines as soon as possible to reduce the amount of event data that needs to be periodically upgraded.

In practice, this is not possible and an upgrade rollout can easily take several weeks. Likewise, unforeseen complications may force you to upgrade some child machines before their parent server has been upgraded. For these reasons, you will inevitably have different version machines operating alongside each other during the upgrade rollout. Before starting the rollout, you need to understand how CA DataMinder handles data transfers (policy changes and captured data) between machines running different versions of the product. Full details about version compatibility are provided in the *Upgrade Guide*.

Chapter 2: Requirements

This section gives the hardware and software requirements for the Central Management Server (CMS), gateways, endpoint machines and other CA DataMinder host servers.

This section contains the following topics:

CMS and Gateways

Note the following requirements for the CMS and gateway servers:

Operating System

The CMS and gateways are included in the server.msi and server_x64.msi installation packages.

Server.msi supports 32-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008

Server_x64.msi supports 64-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)
- Windows Server 2008 R2
- Windows Server 2012

Note 1: We have not tested these operating systems with the current versions of server.msi and server_x64.msi.

Database

CA DataMinder servers need sufficient memory and processing power to run your chosen database application. See your database documentation for details. The supported databases are:

- Oracle 10g (10.2.0.4) or 11g (11.1.0.7), 11g Release 2, or later

  Oracle 10g users may see improved search and report performance by applying the following Oracle fix 5765456:7. See the following My Oracle Support (formerly Oracle Metalink) notice: "Bug 5040753 Optimal index is not picked on simple query / Column group statistics."

- Microsoft SQL Server 2005

  We recommend the SP2 release.

- Microsoft SQL Server 2008
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2008 Express Edition

  Not recommended for CMSs or FastStart base machines.

- Microsoft SQL Server 2012

Note: SQL Server is supported on Windows servers only. Verify that the SQL Server Browser service has started.

Disk space

20 GB

This storage estimate covers the CA DataMinder infrastructure, consoles and captured data (based on typical usage rates). Note:

- The CMS needs sufficient free disk space to store data captured on all client machines.
- A gateway server needs enough free disk space to store data captured on the client machines that it serves. (Note that you can purge this captured data as soon as it has been replicated to the CMS.)

Consoles

To run the Administration console or Data Management console, you require:

Microsoft Internet Explorer 7, 8, or 9.

Note: If you are installing the Data Management console on Windows 2003 or XP, apply the hotfix described in the following Microsoft knowledge base article: http://support.microsoft.com/kb/950094.

More information:

Configure the Management of Free Disk Space on Servers (see page 76)

Endpoint Agents

The following are requirements for CA DataMinder endpoint computers:

Operating System

CA DataMinder endpoint agents are included in the client.msi and client_x64.msi installation packages.

Client.msi supports 32-bit versions of these operating systems:

- Windows XP (see note 1)
- Windows Vista (see note 2)
- Windows 7 (see note 2)
- Windows 8
- Windows Server 2003 (see notes 2 and 3)
- Windows Server 2008 (see notes 2 and 3)

Client_x64.msi supports 64-bit versions of these operating systems:

- Windows XP (see notes 1 and 2)
- Windows Vista (see note 2)
- Windows 7
- Windows 8
- Windows Server 2003 (see notes 2 and 3)
- Windows Server 2008 (see notes 2 and 3)
- Windows Server 2008 R2 (see note 3)
- Windows Server 2012 (see note 3)

Note 1: On XP computers, the Client File System Agent (CFSA) and Client Print System Agent (CPSA) require SP2 or later.

Note 2: We have not tested these operating systems with the current versions of client.msi and client_x64.msi.

Note 3: CA DataMinder endpoint agents support 'centralized applications' running on Windows Server. For these deployments, users access applications such as Outlook on a central server using, for example, Citrix or Remote Desktop Connection.

Important! For details about CA DataMinder and the Windows firewall, see the Known Issues (see page 592) section.

Memory

128 MB

Disk space

Allow approximately 45 MB for the CA DataMinder infrastructure plus an Administration console.

You also need sufficient free disk space to store captured data in a local database. (You can purge this captured data as soon as it has been replicated to the parent server.

Email integration

- Microsoft Outlook 2003, 2007, 2010, or 2013
- Lotus Notes 7, 8, or 8.5

Note: We recommend to use Outlook 2010 with "Cached Exchange Mode" enabled if the CA DataMinder client agent is installed. Disabling Cached Exchange Mode results in increased network traffic between the Outlook client and the Exchange Server.

Browser integration

- Google Chrome
- Mozilla Firefox
- Microsoft Internet Explorer 8, 9, or 10

File system integration

The CFSA requires Windows XP SP2 or later.

Print integration

Before you install the CPSA, close any applications that are running.

For details, see Before Installing on Client Machines.

Consoles

To run the Administration console or Data Management console, you require:

- Microsoft Internet Explorer 7, 8, or 9

Note: If you are installing the Data Management console on Windows 2003 or XP, apply the hotfix described in the following Microsoft knowledge base article: http://support.microsoft.com/kb/950094.

Limitation of 32-bit Client.msi on 64-bit OS

When deploying a client package to a machine with a 64-bit operating system, you typically use client_x64.msi.

But if you want to use a single deployment across hosts with both 32-bit and 64-bit operating systems, you can deploy the 32-bit client.msi package to a 64-bit operating system. However, the client.msi package does not contain any 64-bit binary files. If you choose this deployment option, be aware that the following endpoint features are not supported:

- Client File System Agent
- Client Print System Agent
- Integration with 64-bit version of Microsoft Outlook
- Integration with 64-bit version of Microsoft Internet Explorer
- Application Monitor

Email Server Agents

Note the requirements for the following email server agents:

Exchange Server Agent

Microsoft Exchange Server 2003, 2007, 2010, or 2013

Exchange 2013

Install the Exchange 2013 server agent on each mailbox server.

CA DataMinder can analyze and apply policy to emails passing through Microsoft Exchange Server 2013. But see the support limitations in the Announcements section of the *Release Notes*.

Exchange 2007 and 2010

Install the Exchange 2007 or 2010 server agent on each Hub Transport Server in the Active Directory site.

See also the 'MAPI client and CDO 1.2.1' requirement.

Domino Server Agent

Lotus Domino 7.0.2, 8, and 8.5

CA DataMinder only supports 32-bit versions of these releases. 64-bit versions are not currently supported.

Note: Integration has been tested using the Domino versions listed above. CA DataMinder may integrate with other Domino versions, but these have not been tested.

IIS SMTP Agent

The SMTP server hosting the IIS SMTP agent must be running IIS 6.0.

The IIS SMTP agent is only supported on 32-bit systems.

See also the 'MAPI client and CDO 1.2.1' requirement.

MAPI client and CDO 1.2.1

CA DataMinder integration with Exchange 2007 and 2010 and SMTP servers requires the Messaging API and Collaboration Data Objects 1.2.1 component.

This software is typically referred to as the 'MAPI client and CDO 1.2.1' component. Download the 'MAPI client and CDO 1.2.1' component from the Microsoft Web site. We recommend using the latest build, which is unrelated to the 1.2.1 version number.

PE domain user

The PE domain user must be a member of the local Administrators group on the machine hosting the policy engine hub and email server agent. Confirm that this is so before installing the policy engine hub.

Policy Engine Host Machines

The following are requirements for policy engine host machines:

Operating System

Policy engines are included in the server.msi and server_x64.msi installation packages.

Server.msi supports 32-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008

Server_x64.msi supports 64-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)
- Windows Server 2008 R2
- Windows Server 2012
- **Note 1:** We have not tested these operating systems with the current versions of server.msi and server_x64.msi.

Disk space

The host machine needs sufficient memory to cache all ‘effective’ user policies.

See Host Machine Memory (see page 339).

CA DataMinder server

The host machine must be the CMS or a gateway server.

For best performance, we recommend a gateway.

PE domain user

The PE domain user must belong to the local Administrators group on the policy engine host machine.

See PE domain user (see page 340).

iConsole Servers

The operating system and web service requirements apply to the iConsole front-end Web server and application servers. The dashboard and browser requirements apply to browser host machines.

Operating System

iConsole front-end Web servers and application servers are included in the web.msi installation package.

Web.msi supports 32-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)

Web.msi also supports 64-bit versions of these operating systems:

- Windows Server 2008 R2
- Windows Server 2012

Note 1: We have not tested these operating systems with the current versions of web.msi.

Web service

IIS

The host server requires Microsoft Internet Information Services (IIS).

IIS must be running in 32-bit mode. iConsole servers cannot run if IIS is hosting other applications that require it to be in 64-bit mode.

If the host server is running IIS 7.0, 7.5, or 8.0, you must install various Web Server Role Services (see page 113).

.NET Framework

The host server must be running .NET Framework 2.0.

Browser

The iConsole runs in the following browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Internet Explorer 8, 9, or 10

Note: The iConsole may run in other browsers, but these have not been tested.

BusinessObjects InfoView

(Applicable only if CA DataMinder integration with BusinessObjects Intelligence is enabled.)

InfoView is the BusinessObjects web portal. You can launch InfoView from the iConsole or you can browse to InfoView directly.

In BusinessObjects XI 3.1 SP5, InfoView is supported in:

- Mozilla Firefox
- Microsoft Internet Explorer 8 and 9 (in Compatibility View only)

More information:

Requirements for Application Servers and Front-End Web Servers (see page 110)

External Agent Machines

The following are requirements for external agent machines:

Operating System

The External Agent API is included in the integration.msi and integration_x64.msi installation packages.

Integration.msi supports a 32-bit version of:

- Windows Server 2003 (see note 1)
- Windows Server 2008

Integration_x64.msi supports 64-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)
- Windows Server 2008 R2
- Windows Server 2012
- **Note 1:** We have not tested these operating systems with the current versions of integration.msi and integration_x64.msi.

Email

Host machines must be running a Microsoft Exchange-compatible email application such as:

- Microsoft Outlook 2003, 2007, 2010, or 2013

Important! Outlook *must* be the default email application on the host machine.

Content Services

Note the following requirements for CA DataMinder Content Services.

IDOL Content Database and Connector Framework

The IDOL content database and IDOL connector framework that ship with CA DataMinder are available in the Content_IDOL.msi installation package.

Content_IDOL.msi supports 64-bit versions of these operating systems:

- Windows Server 2008 R2
- Windows Server 2012

Important! If you install the IDOL connector framework included with CA DataMinder, you must install it on the same server as the content indexer and content proxy server.

Content Indexer and Content Proxy Server

You can install these components on any CA DataMinder server.

These components are included in the server.msi and server_x64.msi installation packages.

Server.msi supports 32-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008

Server_x64.msi supports 64-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)
- Windows Server 2008 R2
- Windows Server 2012
- **Note 1:** We have not tested these operating systems with the current versions of server.msi and server_x64.msi.

Chapter 3: CA DataMinder Servers

This section describes how to install and uninstall CA DataMinder servers. These include the CMS, gateways and utility machines.

This section contains the following topics:

Server Types

This section summarizes the key characteristics of CMSs, gateways and utility machines. You manage these servers in the Administration console.

More information:

Central Management Server (CMS) (see page 35)
Gateway Servers (see page 36)
Utility Machines (see page 36)

Central Management Server (CMS)

The CMS holds the central database for your CA DataMinder installation. This database contains the policies for all your systems and users, plus all the captured data (such as emails, files, and Web pages) associated with these users.

Install the CMS before deploying CA DataMinder to client systems. Also install an Administration console. Use the console to perform various configuration tasks before you roll out CA DataMinder across your organization.

Gateway Servers

Gateways are optional data-routing servers, operating between the CMS and client systems. Each gateway can serve multiple client systems and is connected to a single parent server. The Event Import utility is also hosted on a gateway. The parent is either the CMS or another gateway. This hierarchical, distributed deployment provides resilience and network load balancing.

Note: Gateways are optional. You can connect any client system directly to the CMS if preferred.

Utility Machines

Utility machines act as Content Proxy servers or iConsole application servers. Utility machines let you run these components without overloading your existing CA DataMinder servers.

Install a utility machine separately *before* installing an iConsole application server.

Before Installing a CA DataMinder Server

Configure your database and verify that computer name resolution is operating correctly on your network before you install a CA DataMinder server.

Server Name Resolution

When you deploy CA DataMinder to client systems, identify its parent server (the CMS or gateway) by name or IP address. If you specify the CMS or gateway by its *name*, verify that the client systems can resolve this name. If they cannot do so, they cannot locate it. Choose a method of computer name resolution that suits the needs of your organization, for example, DNS or a WINS server.

Database Configuration

Before installing a CMS or gateway server, your chosen database engine must already be installed and correctly configured on the target computer. For Oracle and SQL Server configuration guidelines, see the *Database Guide*.

Disable 8.3 File Names

For NTFS file systems, we recommend that you disable creation of 8.3 file names on the system hosting the CA DataMinder data folder. If 8.3 file creation remains enabled, performance can be adversely affected. To disable 8.3 file creation, you set the NtfsDisable8dot3NameCreation registry value to 1 (one) and reboot; for further details, search for this registry value on www.support.microsoft.com.

Gateways

To simplify mass deployments, you can bulk create new gateway accounts and assign gateways to parent servers in advance of the CA DataMinder rollout. This enables you to deploy multiple gateways using a single source image (which identifies a single parent server) while ensuring that each gateway automatically connects to its 'correct' parent server immediately after installation. To bulk create new accounts, you use the Account Import feature.

Note: If you install Event Import or Import Policy on a gateway, the target server must meet the email archive integration requirements.

Utility Machines

Utility machines inherit the common client system policy. For Content Proxy servers and iConsole application servers, the Communications Encryption policy setting (in the Security folder) and settings in the Logging folder are relevant.

More information:

Command File Format - Users (see page 281)

Server Installation Features

When you install a CA DataMinder server, the following features are available:

CA DataMinder Server

Installs the CA DataMinder infrastructure, enabling the various components to run and communicate.

Enterprise Server

Enables the local computer to function as a CA DataMinder server. Exclude this feature if you want a console-only installation.

Policy Engine

Policy engines permit CA DataMinder to integrate with Microsoft Exchange and other external email sources. Policy Engines are normally installed on a gateway server. For details, see Policy engines.

Socket API

Enables you to call a policy engine from a remote location (via the External Agent API), including from a non-Windows system. For example the CA DataMinder Network Boundary Agent uses the Socket API to analyze traffic leaving or entering the corporate network from the internet. For details, see Socket API.

Event Import

Imports emails and IM conversations from external sources such as Microsoft Exchange mailboxes.

Remote Policy Engine Connector

Installs a modified policy engine hub that enables Event Import to connect to policy engines. It is part of the Import Policy feature.

Templates

Installs predefined import configuration files. These can be customized as required—see Template configuration files.

Content Services

This installs the content proxy server and content indexer components that interact with the CA DataMinder infrastructure—see Content Services deployment (see page 409).

Content Indexer Console

Lets you specify which events to index.

Content Indexer Server

Installs the indexing service to index events into a Content database.

Content Proxy Server

Installs the service that enables the iConsole or Data Management console to access a content database when running a content search.

Remote Data Manager

Enables the Data Management console to retrieve and display events archived in third party remote storage locations. This feature is normally installed on a utility machine. For details, see Remote Data Manager.

Quarantine Manager

Ensures that e-mails released from quarantine are sent on to their original recipients.

CMS Storage Connectors

Installs the connectors needed to implement third party object storage integration.

EMC Centera Connector

Enables EMC Corporation's Centera content addressed storage (CAS) solution to integrate with CA DataMinder.

IBM Content Manager

Enables IBM DB2 Content Manager to integrate with CA DataMinder. It allows storage management, archiving and retrieval of large volumes of captured data.

Management Console

This feature enables the host system to run these consoles:

Administration Console

The Administration console enables CA DataMinder administrators to manage user accounts, machine accounts, and policies. It also enables administrators to manage scanning jobs and content agents and to view log files.

Data Management Console

The Data Management console (DMC) allows reviewers to search for, view and update the audit status for captured events.

Note: By default, only the Administration console is installed with the Management Console feature.

More information:

Account Import (see page 245)
Quarantine Manager (see page 225)
Object Storage (see page 307)

Installing a CMS or Gateway

Note: To install a CMS or gateway, you need local administrator rights on the target server.

To install a CMS or gateway server

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose the CA DataMinder Platform and then click Install.

   The CA DataMinder server installation wizard starts in a separate window.

4. In the server installation wizard, browse to the Custom Setup screen.

5. In the Custom Setup screen, choose the components that you want to install.

   To install these components to a different location, click Change.

   To check whether the target volumes have sufficient free disk space for the selected components, click Disk Space.

   Important! If you install to a different location, the target path must not include folders whose names contain Far Eastern characters. The CA DataMinder infrastructure cannot handle these paths.

6. In the Server Type screen, install one of the following:

   Central Management Server

   > If you install a CMS, note that certain user accounts are created automatically.

   Gateway

   > You typically use a gateway to host policy engines or Event Import.

7. (Gateways only) In the Connectivity screen, enter the name or IP address of the parent server. This can be the CMS or a gateway server.

8. In the Data Location screen, specify the name and network location of the data folder.

   This folder contains all the configuration data and captured data associated with the current server. You can accept the default location or specify a different location.

   Note: See the recommendation to disable 8.3 file names (see page 36).

   Accept the default location

   Go to step 9.

   Specify a different location

   Click Change to specify a different location. Remote data folders are described below.

   Important! Do not install the data folder to an encrypted folder or file system! Also, do not compress the data folder if you are using SQL Server.

   Specify a remote data folder

   You can specify a network file share, using the universal naming convention (UNC) or mapped network drive. For example:

   ```
   \\MyMachine\share_name\target_folder
   ```

   ```
   \\NAS_device\share_name\target_folder
   ```

   Sharing and Security settings for the remote folder *must* allow Full Control for both the user running the installation wizard and the machine running the CA DataMinder Enterprise Server software.

   Note: For SQL Server users, you cannot specify a network file share. This is a known issue. See the *Database Guide*; search the index for ‘Data folder: remote, and SQL Server’.

9. The next step depends on whether you are reusing an existing CA DataMinder database:

   - If this is a wholly new installation, and the installation wizard does not detect any existing CA DataMinder database, go to step 10.

   - If an existing CA DataMinder database is detected, the installation wizard displays details about the database here in the Data Location screen. If you do not want to re-use this database *or* you do want to re-use the database but in step 8 you specified a new Data folder, go to step 11.

10. If an existing CA DataMinder database is detected, the installation wizard displays details about the database here, in the Data Location screen. If you want to re-use this database *and* the original Data folder, no further configuration is required; simply click Next. Go directly to step 23.

11. In the Database Type screen, select the database engine to use on the current server. Choose Oracle or Microsoft SQL Server.

12. In the Database Identification screen, enter details about your chosen database.

    a. You can select a local or remote database. In either case, specify the host server (enter localhost to specify a database on the local machine) and the TCP/IP port number used by the host server.

       Note: Regardless of the type of database engine, if the installation wizard is unable to validate the host server, for example because it is not switched on, the wizard adds a Bypass Database Validation check box to the screen. You can select this check box to skip the validation, but verify the spelling of the server name, otherwise the installation fails.

    b. Depending on your chosen database engine, supply further details.

       - For **Oracle***,* go to step 13.
       - For **SQL Server**, go to step 14.

13. **Oracle database**

    You must provide an appropriate service identifier (SID) to identify the correct database. The SID corresponds to the SID_NAME value in the listener.ora file on the Oracle host server. The installation wizard attempts to validate this SID if you chose the current machine as the host server.

    Now go to step 15.

14. **SQL Server database**

    Select the host server. The IP port and database name are set automatically, though you can change both.

    To set up your SQL Server database:

    a. Click Server to select the host server in the Database Server dialog. This dialog lists any servers found to be hosting SQL Server. For multiple SQL Server instances running concurrently on the same computer, the dialog identifies each instance as:

       ```
       <machine name>\<Instance name>
       ```

       Where <machine name> is the name of the server on which SQL Server is running, and <Instance name> is the name of the SQL Server instance. For example:

       ```
       MyDBServer\Instance_1
       ```

b. The IP port is set automatically when you choose the host server. You do not normally need to change this. But if necessary, click Port to manually set the port number.

c. Enter the database name. This defaults to:

```
WGN_<local machine name>
```

Where <local machine name> is the name of the server on which you are installing the CMS. For example, if this server is named CMS-HARDY:

```
WGN_CMS-HARDY
```

If you have already created a database for CA DataMinder, change the default name to the name used in SQL Server.

Note: For SQL Server 2005 databases, verify that the SQL Server Browser service has started. For details, see the *Database Guide*; search the index for 'SQL Server 2005'.

15. In the Database Accounts screen, specify the database accounts used by CA DataMinder to access the CMS database:

Primary User

This is the main CA DataMinder database account. The infrastructure uses this account to access the CMS database. For SQL Server databases, the primary user owns the schema.

Search User

CA DataMinder consoles use this database account when searching the CMS database for events. This is a secure account that is subject to row level security (RLS) when searching the database for events. This ensures that reviewers cannot see events that they are not permitted to see when they run a search. If multiple database security models are enabled on your CMS, specify a separate Search User database account for each security model.

The database account you specify now is automatically associated with the default database security model, Management Group (Standard).

Schema Owner

(Only available for Oracle CMSs) This optional account owns the database schema. Some organizations choose to have separate accounts for the primary user and the database owner. This is typically for security reasons, for example, to ensure that employees cannot connect to the CMS database as the primary user and delete sensitive data or drop the underlying database objects.

In all cases, you can specify existing database accounts or instruct the wizard to create new ones. If you specify existing accounts, verify that they have appropriate roles and privileges. The requirements for your Oracle users or SQL Server logins are provided in the *Database Guide*.

To specify the database user credentials:

a. For the Primary User account, click the ... button. In the resulting User Credentials dialog, specify the username and password for the database account. If this is a new account, select the Create User check box.

   Important! All SQL Server accounts—for the Primary User, Search User, and even the database administrator in step 16—**must** use SQL Server Authentication! You specify the authentication method in the Login Properties dialog in SQL Server Enterprise Manager.

b. Repeat the above steps for the Search User database account.

c. (Optional) For Oracle CMSs only, repeat the above steps for the Schema Owner database account.

   Note: The installation wizard does not try to validate these account details. Verify that you have entered them correctly, otherwise the installation fails.

d. If any of the specified database user accounts are new, specify an existing database account (the ‘Database Administrator user’) that the installation wizard can use to log in to SQL Server or Oracle to create the new accounts; go to step 16.

   If all the specified accounts are existing database user accounts, go directly to step 19.

16. Still in the Database Accounts screen, go to the Database Administrator User field and click the ... button. In the resulting User Credentials dialog, specify the username and password for the database administrator account (but see the SQL Server note in step 15.a).

    - For Oracle CMSs, you now need to define a tablespace for each new user; go to step 17.
    - For SQL Server CMSs, go directly to step 18.

17. For Oracle CMSs only. In the Database Tablespace Names screen, you must define the tablespace names for each new account.

    By default, Oracle creates all new databases in the ‘Users’ tablespace, but we strongly recommend that you create separate tablespaces for the CA DataMinder database accounts. This has several advantages. For example, it allows precise monitoring of the CMS database and also permits off-line maintenance of the tablespace without disrupting other products that may be sharing your Oracle server.

18. (CMS installations only) In the Data Warehouse Configuration screen, specify whether to enable data warehousing and whether to collect event participant data.

    Note: For full deployment details, see the Data Warehouse chapter later in this guide.

19. (CMS installations only. Applicable only if you enable data warehousing.) In the Data Warehouse Database Account screen, define the following database accounts:

    - Reporting User
    - Unrestricted Search User

    Note: For full details about these database users, see the Data Warehouse chapter later in this guide.

20. (CMS installations only) In the Administrator Credentials screen, enter details (name and password) for your Primary Administrator account. This is one of three accounts created automatically on the CMS.

    This account has full administrative privileges and full management group coverage. You use an administrator account to configure CA DataMinder after deployment.

21. In the Service Accounts screen, specify the logon accounts used by the various CA DataMinder services. For example, the infrastructure defaults to the LocalSystem account.

    Important! If you chose to specify a remote data folder in step 8, you will need to change the credentials of the infrastructure service.

    Important! If you change any service logon account to a named user, you must manually assign the ‘Log on as a service’ security privilege to this account. See the next section.

22. (Applicable only if you install the Remote Data Manager.) In the Remote Data Manager Configuration screen, specify which archives to integrate with.
23. In the final wizard screen, click Install to start the file transfer.

More information:

Before You Start using CA DataMinder (see page 71)
Before Installing a CA DataMinder Server (see page 36)

Assign the ‘Log on as a service’ Privilege

If you changed any service logon account to a named user, you must manually assign the ‘Log on as a service’ security privilege to this account:

1. Ensure that you are logged on to the target server with local administrator rights on.
2. On the target server, open the Local Security Policy applet or, if this server is a domain controller, open the Domain Controller Security Policy applet. Both applets are available in Administrative Tools.
3. Expand the Local Policies branch and select User Rights Assignment. This security area determines which users have logon privileges on the local computer.
4. Assign the Log on as a service privilege to the service logon account.

Installations using Terminal Services

We strongly recommend that you install the CMS server, gateway server, or CA DataMinder utility machines directly on the target server and do not user Terminal Services. But if this is not possible and you do need to use Terminal Services, be aware of the logon requirements on the target server or target machine.

Logging on to the server or utility machine

Specifically, when you log on to the target server or utility machine using a Terminal Services connection, we recommend you use an account with rights to restart it. This is because the installation wizard may sometimes prompt you (during the file copying stage) to restart the target server or utility machine.

If you do not have rights to restart the server or utility machine, ensure that the colleague who will restart the server or utility machine on your behalf is already logged on before you run the installation wizard. If your colleague attempts to log on to the server or utility machine after the wizard has started, the installation will fail.

Restarting the utility machine

If you do need to restart the server or utility machine, you will need to reopen your Terminal Services connection.

Important! When you do this, you **must** log on to the server or utility machine using the same account that you used to start the installation wizard!

This is because for server or utility machine installations using Terminal Services, it is essential that the first account to log on to the server or utility machine after it restarts is the same account that launched the installation wizard. If a different account is the first to log on after the server or utility machine restarts, the installation will fail.

Unattended Installations

CA DataMinder installs using the Microsoft Windows Installer service. This enables you to use command line options of msiexec.exe to install or uninstall. For example, you can install CMSs or gateways on multiple machines using logon scripts or other third party deployment software.

Installing CMSs or Gateways

When installing CMSs or gateways from a command line, you must reference the source image for CMSs and gateways that ships with CA DataMinder. To deploy multiple CMSs or gateways, you need to copy this source image to each target server. You can then install the CMSs or gateways automatically using a script based on standard command line options for Msiexec.exe, plus various proprietary variables unique to CA DataMinder.

Note: These instructions also apply to utility machines.

The basic syntax for an installation is:

```
msiexec /i <Path>\server.msi
```

Where server.msi is the CA DataMinder source image and <Path> points to the location of this image.

However, you need to supplement this basic syntax with parameters to configure the installation folder, server type, the primary administrator account, details about the database engine, and so on. These are discussed in the next section.

Note: Unlike command line deployment of client machines, there is no need to perform an administrative installation of the CA DataMinder source image for servers. Instead, you can directly reference the server source image that ships with CA DataMinder.

Msiexec.exe Parameters

CA DataMinder supports a range of variables that you can use as Msiexec.exe parameters to configure any server installation from a command line. Full details are given in Command line parameters for Msiexec.exe. The key variables are:

INSTALLDIR

Specifies the installation folder.

INSTALLDIR64

Specifies the installation folder for 64-bit files.

WGNADMINUSERNAME

For CMSs only. Specifies the primary administrator in CMS installations.

WGNADMINPASSWORD

For CMSs only. Specifies the primary administrator’s password in CMS installations.

WGNSERVERTYPE

Specifies whether to install a CMS or gateway.

WGNDATABASETYPE

Identifies the database engine that you will use on the CMS—Oracle or SQL Server.

WGNDATABASESERVER

Identifies the host server for your database.

WGNDATABASESERVICENAME

For Oracle databases, this specifies a service name to identify the correct database tables.

WGNDATABASENAME

For SQL Server databases, this specifies the name of the database.

WGNDATABASEUSERNAME

Specifies the database account used by CA DataMinder to access the CMS database.

WGNDATABASEPASSWORD

Specifies the password for the database account.

WGNDATA

Specifies the name and network location of the CMS data folder.

WGNPARENTSERVERNAME

For gateways and utility machines only. Identifies the parent server.

User Accounts Created Automatically on the CMS

When you install a CMS, the following user accounts are created automatically in the top-level Users group:

Primary Administrator

This is the CA DataMinder account you created when you installed the CMS. This account has full administrative privileges.

Use this account to configure CA DataMinder and create other user accounts, including other administrators and managers.

UnknownInternalSender

This account has no management or administrative privileges; its sole purpose is as a conduit to enable policy engines to apply policy to internal emails from unrecognized senders.

When you install a CMS, the Unknown Internal Sender setting in the machine policy defaults to this UnknownInternalSender user account.

ExternalSender

This account has no management or administrative privileges; its sole purpose is as a conduit to enable policy engines to apply policy to external emails.

When you install a CMS, the External Sender setting in the machine policy defaults to this ExternalSender user account.

DefaultFileUser

This account has no management or administrative privileges; its sole purpose is as a conduit to enable policy engines to apply policy to scanned, captured or imported files if no other means are available to determine the policy participant.

When you install a CMS, the Default Policy for Files setting in the machine policy defaults to this DefaultFileUser user account.

DefaultClientFileUser

This account is similar to the DefaultFileUser account. The account is used solely by the Client File System Agent (CFSA) when scanning local workstations. The CFSA uses this account to apply the same policy to scanned files across all workstations.

When you install a CMS, the Default Policy for Data At Rest setting in the machine policy defaults to this DefaultClientFileUser user account.

DefaultClassificationUser

This account is used solely by the Content Classification Service (CCS) when classifying documents forwarded to CA DataMinder by external applications. The CCS uses this account to apply the same policy to all documents that require classification.

When you install the CCS, the Default Policy for Classification setting in the machine policy defaults to this DefaultClassificationUser account.

NT AUTHORITY* accounts

CA DataMinder creates a series of CA DataMinder user accounts to match Windows system user. For example, the NT AUTHORITY\System account is used by policy engines running as LocalSystem when logging on to the CMS.

Note: All of these accounts are automatically exempt from CA DataMinder policy and do not count towards the licensed user limit.

Uninstalling a CMS, Gateway or Utility Machine

You can uninstall a CMS, gateway, or CA DataMinder utility machine using the Add/Remove Programs applet in the Windows Control Panel.

To uninstall using Add/Remove Programs

1. In the Add/Remove Programs dialog, launch the installation wizard:
2. Select 'CA DataMinder Server' and click Change:

   Note: If you click Remove instead, CA DataMinder is uninstalled but you do not get the option to keep the database (see step 5).
3. When the Installation Wizard starts, go to the Program Maintenance screen.
4. Choose Remove to uninstall CA DataMinder from the current machine.
5. In the Remove the Program screen, you can choose to keep or remove the database.
   - If you choose to keep the database, you can reconnect to it when you reinstall CA DataMinder. To do this, you simply re-enter the user name and password for a valid database account when you run the installation wizard.
   - If you choose to delete the database, note that:

     SQL Server: The database and its contents are removed completely.

     Oracle: The database contents are deleted but the database itself is not deleted.
6. The wizard now has all the information it needs and begins the uninstall immediately.

Installing a Utility Machine

To install a utility machine

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose the CA DataMinder Platform and then click Install.

   The CA DataMinder server installation wizard starts in a separate window.

4. In the server installation wizard, browse to the Custom Setup screen.

5. In the Custom Setup screen, choose the components that you want to install.

   To install these components to a different location, click Change.

   To check whether the target volumes have sufficient free disk space for the selected components, click Disk Space.

   Important! If you install to a different location, the target path must not include folders whose names contain Far Eastern characters. The CA DataMinder infrastructure cannot handle these paths.

6. In the Server Type screen, click Utility Machine.

7. In the Connectivity screen, enter the name or IP address of the parent server. This can be the CMS or a gateway server.

8. In the Data Location screen, specify the name and network location of the data folder.

   This folder contains all the configuration data and captured data associated with the current server. You can accept the default location or specify a different location.

   Note: See the recommendation to disable 8.3 file names (see page 36).

   Accept the default location

   Go to step 9.

   Specify a different location

   Click Change to specify a different location. Remote data folders are described below.

   Important! Do not install the data folder to an encrypted folder or file system! Also, do not compress the data folder if you are using SQL Server.

   Specify a remote data folder

   You can specify a network file share, using the universal naming convention (UNC) or mapped network drive. For example:

   ```
   \\MyMachine\share_name\target_folder
   ```

   ```
   \\NAS_device\share_name\target_folder
   ```

   Sharing and Security settings for the remote folder *must* allow Full Control for both the user running the installation wizard and the machine running the CA DataMinder Enterprise Server software.

   Note: For SQL Server users, you cannot specify a network file share. This is a known issue. See the *Database Guide*; search the index for ‘Data folder: remote, and SQL Server’.

9. The next step depends on whether you are reusing an existing CA DataMinder database:

   - If this is a wholly new installation, and the installation wizard does not detect any existing CA DataMinder database, go to step 10.
   - If an existing CA DataMinder database is detected, the installation wizard displays details about the database here in the Data Location screen. If you do not want to re-use this database *or* you do want to re-use the database but in step 8 you specified a new Data folder, go to step 11.

10. If an existing CA DataMinder database is detected, the installation wizard displays details about the database here, in the Data Location screen. If you want to re-use this database *and* the original Data folder, no further configuration is required; simply click Next. Go directly to step 18.

11. In the Database Type, select the database engine to use on the utility machine. The default is Microsoft Jet.

    If you select Microsoft Jet, no further database configuration is needed. Go directly to step 18.

    if you select Oracle or SQL Server, go to step 12.

12. In the Database Identification screen, enter details about your chosen database.

    a. You can select a local or remote database. In either case, specify the host server (enter localhost to specify a database on the local machine) and the TCP/IP port number used by the host server.

       Note: Regardless of the type of database engine, if the installation wizard is unable to validate the host server, for example because it is not switched on, the wizard adds a Bypass Database Validation check box to the screen. You can select this check box to skip the validation, but verify the spelling of the server name, otherwise the installation fails.

    b. Depending on your chosen database engine, supply further details.

       - For **Oracle***,* go to step 13.
       - For **SQL Server**, go to step 14.

13. **Oracle database**

    You must provide an appropriate service identifier (SID) to identify the correct database. The SID corresponds to the SID_NAME value in the listener.ora file on the Oracle host server. The installation wizard attempts to validate this SID if you chose the current machine as the host server.

    Now go to step 15.

14. **SQL Server database**

    Select the host server. The IP port and database name are set automatically, though you can change both.

To set up your SQL Server database:

a. Click Server to select the host server in the Database Server dialog. This dialog lists any servers found to be hosting SQL Server. For multiple SQL Server instances running concurrently on the same computer, the dialog identifies each instance as:

```
<machine name>\<Instance name>
```

Where <machine name> is the name of the server on which SQL Server is running, and <Instance name> is the name of the SQL Server instance. For example:

```
MyDBServer\Instance_1
```

b. The IP port is set automatically when you choose the host server. You do not normally need to change this. But if necessary, click Port to manually set the port number.

c. Enter the database name. This defaults to:

```
WGN_<local machine name>
```

Where <local machine name> is the name of the server on which you are installing the CMS. For example, if this server is named CMS-HARDY:

```
WGN_CMS-HARDY
```

If you have already created a database for CA DataMinder, change the default name to the name used in SQL Server.

Note: For SQL Server 2005 databases, verify that the SQL Server Browser service has started. For details, see the *Database Guide*; search the index for ‘SQL Server 2005’.

15. In the Database Accounts screen, specify the database accounts used by CA DataMinder to access the CMS database:

Primary User

This is the main CA DataMinder database account. The infrastructure uses this account to access the CMS database. For SQL Server databases, the primary user owns the schema.

Schema Owner

(Only available for Oracle CMSs) This optional account owns the database schema. Some organizations choose to have separate accounts for the primary user and the database owner. This is typically for security reasons, for example, to ensure that employees cannot connect to the CMS database as the primary user and delete sensitive data or drop the underlying database objects.

In all cases, you can specify existing database accounts or instruct the wizard to create new ones. If you specify existing accounts, verify that they have appropriate roles and privileges. The requirements for your Oracle users or SQL Server logins are provided in the *Database Guide*.

To specify the database user credentials:

a. For the Primary User account, click the ... button. In the resulting User Credentials dialog, specify the username and password for the database account. If this is a new account, select the Create User check box.

   Important! All SQL Server accounts—for the Primary User, Search User, and even the database administrator in step 16—**must** use SQL Server Authentication! You specify the authentication method in the Login Properties dialog in SQL Server Enterprise Manager.

b. (Optional) For Oracle CMSs only, repeat the above steps for the Schema Owner database account.

   Note: The installation wizard does not try to validate these account details. Verify that you have entered them correctly, otherwise the installation fails.

c. If any of the specified database user accounts are new, specify an existing database account (the 'Database Administrator user') that the installation wizard can use to log in to SQL Server or Oracle to create the new accounts; go to step 16.

   If all the specified accounts are existing database user accounts, go directly to step 18.

16. Still in the Database Accounts screen, go to the Database Administrator User field and click the ... button. In the resulting User Credentials dialog, specify the username and password for the database administrator account (but see the SQL Server warning in step 15.a).

    ■ For Oracle CMSs, you now need to define a tablespace for each new user; go to step 17.

    ■ For SQL Server CMSs, go directly to step 18.

17. For Oracle CMSs only. In the Database Tablespace Names screen, you must define the tablespace names for each new account.

    By default, Oracle creates all new databases in the 'Users' tablespace, but we strongly recommend that you create separate tablespaces for the CA DataMinder database accounts. This has several advantages. For example, it allows precise monitoring of the CMS database and also permits off-line maintenance of the tablespace without disrupting other products that may be sharing your Oracle server.

18. In the Service Accounts screen, specify the logon accounts used by the various CA DataMinder services. For example, the infrastructure defaults to the LocalSystem account.

    Important! If you chose to specify a remote data folder in step 8, you will need to change the credentials of the infrastructure service.

    Important! If you change any service logon account to a named user, you must manually assign the 'Log on as a service' security privilege to this account. See the next section.

19. In the final wizard screen, click Install to start the file transfer.

More information:

Before You Start using CA DataMinder (see page 71)
Before Installing a CA DataMinder Server (see page 36)

Installing an Administration Console

After deploying CA DataMinder to your CMS, further configuration is needed before you deploy across your organization. This mainly involves changes to the default policies, and some user administration.

More information:

Is the Administration Console Already Installed? (see page 57)
Console-only Installations (see page 58)

Is the Administration Console Already Installed?

To perform these tasks, you use an Administration console. When you install a CMS, this console is normally installed automatically on the target machine. But if you chose a Custom setup type and excluded the Management console feature (which incorporates the Administration console only), then you will need to install the console on another machine.

Console-only Installations

To install the Administration console

1. Choose the machine on which you want to run the Administration console.

   Note: You can install the Administration console on as many machines as you like.

2. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

3. Click Advanced Installation.

4. In the Advanced Install Options screen, choose Endpoint Agents and then click Install.

   This launches the CA DataMinder client installation wizard in a separate window.

5. In the client installation wizard, navigate to the Customer Setup screen.

6. In the Custom Setup screen, include the **Management Console** but exclude all other features.

7. Navigate to the final wizard screen.

8. Click Install to start the file transfer.

Chapter 4: Clustered CMS Deployments

It is possible to install a CMS on Microsoft Failover Clustering as part of a CA DataMinder deployment. The cluster functions as a virtual CMS, with client machines connecting to the virtual CMS (using the virtual server name or IP address), not to the actual CMS on the node servers.

This section assumes the reader has a working knowledge of Microsoft Failover Clustering and is familiar with the best practices for creating clusters and administering cluster nodes and resources.

This section contains the following topics:

Requirements

This section describes a clustered CMS deployment using:

- CA DataMinder 12.5 or later
- Windows 2008 Failover Clustering
- Independent database server

  Oracle 10.2.0.4 or 11.1.0.7, or Microsoft SQL Server 2005 or 2008, installed on a remote server which is not part of the cluster.

  Note: This technical note assumes that your chosen database engine is not clustered, but installed on a server outside of the cluster. However, SQL Server is cluster-aware; if you do intend to include SQL Server in the cluster, please refer to Microsoft documentation.

Cluster Configuration

This document assumes that you have already set up a two-node cluster, with a remote database deployed on a separate server, as shown in the architecture diagram (see page 61).

Note: To decrease the complication of additional service failovers, do not host additional components on the Clustered CMS setups, such as PE, Importer, or Content Services.

Component Names

This document uses the following component names:

MYCLUSTER

Stands for the cluster name.

V-CMS

Stands for the virtual CMS server

MACHINE1

Stands for the name of the server acting as node 1.

MACHINE2

Stands for the name of the server acting as node 2.

E:

Stands for the cluster disk. Both the primary and secondary nodes are connected to shared data disks dedicated to the cluster. The nodes connect to shared storage devices using a SCSI or fiber channel.

Architecture Diagram

The cluster functions as a virtual CMS, with client machines connecting to the virtual CMS (using the virtual server name or IP address), not to the actual CMS on the node servers.

The CA DataMinder infrastructure runs on the secondary node in stand-by mode. If the primary node fails, control passes to the secondary node.

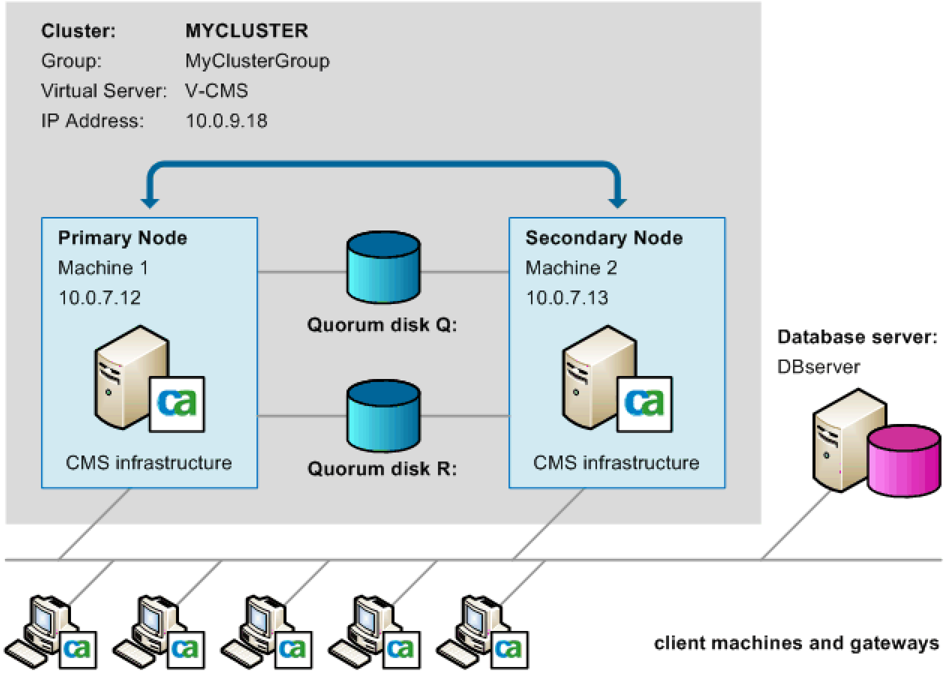

Example Clustered CMS Deployment

How to Deploy the CMS to the Cluster

To deploy the CMS to a cluster, install and configure the CA DataMinder CMS on each node. Use the Microsoft Failover Cluster Management utility as detailed in this chapter.

To deploy the CMS to the Cluster

1. Configure a client access point and storage.
2. Specify the preferred owners.
3. Install the CA DataMinder CMS on node 1.
4. Set up a resource for the CA DataMinder infrastructure.
5. Install the CA DataMinder CMS on node 2.
6. Test the cluster.

Configure Client Access Point and Storage

Use the High Availability wizard to configure the client access point and storage devices for the virtual CMS.

To configure the client access point and storage

1. Click 'Configure a service or application' under Actions.

   The High Availability wizard opens.

2. Click Other Server and click Next.

   The Client Access Point screen displays.

3. Enter Name and IP address of the V-CMS and click Next.

   The Select Storage screen displays.

4. Select cluster disk E: for shared storage and click Next.

   The confirmation screen displays.

5. Confirm the setup and finish the wizard at the summary screen.

   The client access point and storage for V-CMS are configured.

Specify Preferred Owners

If a node or resource fails, and Preferred Owners have been defined, the Cluster Service fails the group to the next available node in the node list. The node list is composed of the Preferred Owners list.

To specify the preferred owner

1. Right-click the V-CMS node under Services and Applications, and click Properties.

   The Properties window opens.

2. Select both nodes, MACHINE1 and MACHINE2, as preferred owners, and click OK.

   The preferred owner nodes are specified.

Install CA DataMinder CMS on Node 1

Use a command line transform to install the virtual CMS on node 1.

Note: Refer to the Requirements and CA DataMinder Servers chapters in the *CA DataMinder Platform Deployment Guide* for details on the CMS installation (see page 40) process.

To install CA DataMinder CMS on node 1

1. Run the server.msi installation from the command line with the V-CMS machine name as a parameter. For example:

   ```
   msiexec /I <DISTRIBUTION_MEDIA>\server.msi WGNINTERNALMACHINENAME=V-CMS
   ```

   The server installation wizard starts.

2. Accept the default options until the Change Current Data Folder screen displays.
3. Select an appropriate location, such as "E:\CA DataMinder DATA", on the shared disk for the data location and click OK
4. Continue with the installation wizard. Accept the defaults until the installation is complete.
5. Verify that a confirmation message similar to the following appears in the activity log:

   ```
   Machine name is: v-CMS. (OS Name=MACHINE1, Physical Node Name=MACHINE1, Override Name=V-CMS
   ```

   The CA DataMinder CMS is installed on node 1.

To install iConsole Standard Searches and Reports

1. Run setup.exe from the distribution media.

   Accept the license agreement and continue to the Advanced Install Options screen.

2. Select iConsole Standard Searches and Reports and click Install.
3. Accept all defaults and finish the wizard.

   The iConsole Standard Searches and Reports components are installed.

More information:

Performing an Administrative Installation (see page 523)

Set up a Resource for the CA DataMinder Infrastructure

Create the CA DataMinder infrastructure resource, add it to the server V-CMS, and configure its properties. In this step you specify the registry replication which is necessary so the installer automatically recognizes the database settings when you set up the virtual CMS on node 2.

To set up a resource for CA DataMinder infrastructure

1. Right-click *V-CMS* under Services and Applications, and click Add a Resource, Generic Service.

   The New Resource Wizard opens.

2. Select CA DataMinder Infrastructure and finish the wizard.

   A new resource is created.

3. Right-click *V-CMS* under Services and Applications and click Properties.

   The CA DataMinder Infrastructure properties window opens.

4. Fill in the dependencies tab:

   And/OR Resource:

   ```
           Name: V-CMS
   AND     IP Address: xxx.xxx.xxx.xxx
   AND     Cluster Disk 1
   ```

   Important: Click Apply before continuing to get the required dependency configured. Otherwise the next step will fail.

5. Fill in the General tab:

   Use network name for computer name

   Select this checkbox to define the computer name.

6. Fill in the Advanced Policies tab:

   Possible Owners

   Select both *MACHINE1* and *MACHINE2.*

7. Fill in the Registry Replication tab:

   Registry key (for 64-bit machines)

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
     \CurrentVersion\Properties
   ```

   Registry key (for 32-bit machines)

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
     \CurrentVersion\Properties
   ```

8. Fill in the Policies tab:

   If resource fails, do not restart

   Select this option.

   Click OK.

9. Right click the resource, and choose ‘bring this resource online’, then ‘Take this resource offline’, and then again, ‘bring this resource online’. This synchronizes the resource with the CA DataMinder Infrastructure service.

   The CA DataMinder infrastructure resource is configured.

Install CA DataMinder CMS on Node 2

After you have installed the CMS on node 1 and set up the resource, proceed with installing the CMS on node 2.

Note: Refer to the Requirements and CA DataMinder Servers chapters in the *CA DataMinder Platform Deployment Guide* for details on the CMS installation (see page 40) process.

To prepare the installation

1. Give node 2 control of the clustered server V-CMS. The CA DataMinder infrastructure resource is offline on node 2 while node 1 has control.

2. Run 'regedit' on node 2 and verify that the following registry settings are there.

   Registry key (for 64-bit machines)

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
     \CurrentVersion\Properties
   ```

   Registry key (for 32-bit machines)

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
     \CurrentVersion\Properties
   ```

   You have verified that the registry keys have been successfully replicated to the next node.

3. Right-click the V-CMS server, click Move this service or application to another node, and click MACHINE2.

   The Failover Clustering Management console displays the following in the summary for V-CMS:

   ```
   Current Owner: <MACHINE2>
   ```

Note: The CA DataMinder Infrastructure resource will fail to start on node 2. This is correct behavior as the CA DataMinder Infrastructure has not yet been installed to node 2.

To install the CA DataMinder CMS on node 2

1. Run the server.msi installation from the command line with the V-CMS machine name as parameter.

   Example:

   ```
   msiexec /I <DISTRIBUTION_MEDIA>\server.msi WGNINTERNALMACHINENAME=V-CMS
   ```

   The InstallShield Wizard opens.

2. Enter the same data path as for Node1. Accept the defaults and continue the Wizard.

   The Data Location screen opens.

3. Verify that the installer automatically recognizes the database settings from the replicated registry keys.

   Important: If the database settings are not detected, the registry has not replicated correctly. Do not continue! Verify that the CA DataMinder Infrastructure resource has the correct registry path set in the Registry Replication tab, and restart the resource. Then repeat the CMS installation process for node 2.

4. Finish the installation accepting the defaults.

   The CMS is installed on the secondary node.

5. Open the properties of the CA DataMinder Infrastructure resource and set the policies according to your failover strategy.

To install iConsole Standard Searches and Reports

1. Run setup.exe from the distribution media.

   Accept the license agreement and continue to the Advanced Install Options screen.

2. Select iConsole Standard Searches and Reports and click Install.

3. Accept all defaults and finish the wizard.

   The iConsole Standard Searches and Reports components are installed.

More information:

Performing an Administrative Installation (see page 523)

Test the Cluster

Test the cluster by connecting a remote Admin Console to the virtual CMS. Do some operations, fail the CMS over to the other node and see if the connection recovers from the remote Admin Console.

From a remote iConsole, perform a search and ensure the results are as expected. Failover the CMS over to the other node and when complete, perform the search again, verifying the results.

Upgrade a Clustered CMS

To upgrade a clustered CA DataMinder

1. Transfer control of the clustered CMS to node 1.
2. Select ‘Do not restart’ on the policies tab of the CA DataMinder Infrastructure resource.
3. (Only required if you are upgrading from 14.0 or earlier) On the primary node, migrate the existing product registry key to the new CA DataMinder registry key.

   To migrate the registry key, you must run the migreg.exe utility from a command line on the primary node. For details, see the 'Pre-Upgrade Tasks' chapter in the *Upgrade Guide*.

   Note: CA Technologies is renaming its security portfolio to follow the Minder product family. In particular, CA DLP has been renamed to CA DataMinder in the current release. In parallel with this product name change, the registry key used by CA DataMinder components has also been renamed.
4. Open a command prompt on the primary node.

   Run the CA DataMinder server installation wizard and specify the CMS Cluster Network Name in the WGNINTERNALMACHINENAME property. For example:

   ```
   msiexec /i c:\DLP_AdminInstall\server.msi WGNINTERNALMACHINENAME=V-CMS
   ```

   The installation wizard opens.
5. Step through the wizard screens. Restart the server if requested.

   The CMS on the primary node is upgraded.

6. Right-click the V-CMS node, select Move this service or application to another node, and click MACHINE2.
7. (Only required if you are upgrading from 14.0 or earlier) On node 2, migrate the existing product registry key to the new CA DataMinder registry key.

   See step 3 for details.
8. Open a command prompt on node 2.

   Run the CA DataMinder server installation wizard and specify the CMS Cluster Network Name in the WGNINTERNALMACHINENAME property. For example:

   ```
   msiexec /i c:\DLP_AdminInstall\server.msi WGNINTERNALMACHINENAME=V-CMS
   ```

   The InstallShield Wizard opens.
9. Step through the wizard screens. Restart the server if requested.

   The node 2 CMS is upgraded.
10. Open the advanced policies tab of the CA DataMinder Infrastructure resource and set the policies for the infrastructure resource in line with your failover strategy.

To test the upgraded clustered CMS

1. Verify the controlling node has brought all resources online and that the CA DataMinder infrastructure has started correctly.
2. Connect a remote Admin Console to the virtual CMS. Do some operations, fail the CMS over to the other node and see if the connection recovers from the remote Admin Console.
3. From a remote iConsole, verify that data can be retrieved successfully when either node has control.

Apply Hotfixes and Service Packs

To apply hotfixes and service packs to a Clustered CMS

1. Verify that you set the CA DataMinder infrastructure resource to 'If resource fails, do not restart' .
2. Apply the hotfix or service pack to node 1.
3. Move CA DataMinder to next cluster node.

   Important: During installation of Hotfixes or Service Packs, isolate both nodes from child machines in order to prevent data corruption.
4. Apply the hotfix or service pack to the next cluster node.
5. Open the properties of the CA DataMinder Infrastructure resource and set the policies according to your failover strategy.

Uninstall Clustered CMS Components

To uninstall CA DataMinder from the primary node

1. Remove CA DataMinder Standard Reports.
2. Remove CA DataMinder Server.

   Important: Verify that the 'Delete CA DataMinder Database (including data)' option is *not* selected.
3. Click Change, Modify
4. Move the server V-CMS to node 2 using the Failover Clustering Management utility.

To uninstall CA DataMinder from the secondary node

1. Remove CA DataMinder Standard Reports.
2. Remove CA DataMinder Server.
3. Click Change, Modify to remove.

   Important: If this is the last node in the cluster and you want to remove the database and all data, verify that the 'Delete CA DataMinder Database (including data)' option is selected. If you would like to keep the data, clear this checkbox.

To clean up the registry entries:

On the primary node, delete the 'CA DataMinder' registry entries.

Registry key (for 64-bit machines)

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
  \CurrentVersion\Properties
```

Registry key (for 32-bit machines)

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
  \CurrentVersion\Properties
```

Chapter 5: Before You Start using CA DataMinder

After deploying your CMS, some further configuration is needed before you start using CA DataMinder. This mainly involves changes to the default policies for key user groups and new machines. You also need to edit the account properties of any new administrators and managers (though you may prefer to leave this task until you have deployed CA DataMinder to your client machines). Finally, you may need to amend your browser security settings if you intend using any Web page control triggers. This section covers the complete range of post-deployment tasks.

Follow these steps:

1. Choose an appropriate account to configure CA DataMinder (see page 72)
2. Configure your CMS machine policy to handle new accounts (see page 73)
3. Configure event purging (see page 74)
4. Configure the Management of Free Disk Space on Servers (see page 76)
5. Configure the Common Client and Gateway Policies (see page 77)
6. Synchronize the Clocks on Your CA DataMinder Machines (see page 78)
7. Configure the Policy for the Default User Group (see page 79)
8. Create and Organize a Hierarchy of User Groups (see page 80)
9. Create Your Administrators and Managers (see page 81)
10. Set Up Support for Unicode Characters (see page 82)
11. Install iConsole Searches (see page 84)
12. Integrate with Third-Party Object Storage Solutions (see page 84)
13. Configure event auditing labels (see page 85)
14. Set Up Policy Engines (see page 83)

This section contains the following topics:

Choose an Appropriate Account to Configure CA DataMinder

To configure CA DataMinder as described here, you must log on to the Administration console using an account with adequate administrative privileges and management group coverage. The simplest way to ensure this is to log on using the Primary Administrator account. This account is created during the CMS installation and has full administrative privileges and full management group coverage.

Configure Your CMS Machine Policy to Handle New Accounts

A machine policy governs how CA DataMinder machines operate, communicate with each other, and protect confidential data. The machine policy on the CMS uses some extra settings to control account handling for unknown users or machines and database management.

When you use CA DataMinder for the first time after installation, you must edit the CMS policy to determine how it handles new machine accounts. It can either create new accounts automatically, or it can require an administrator to manually add new machines in the Administration console.

To configure CMS policy

1. In the Administration console, expand the Machine Administration branch.
2. Select the CMS and click Edit Policy or right-click and choose Edit Policy.
3. In the Machine Policy Editor, browse to the Central Management Server folder.

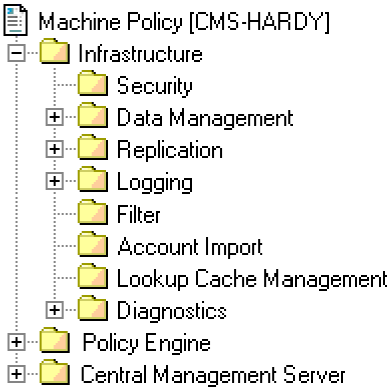

4. **New user accounts:** Configure how the CMS handles new users. In the right-hand pane, double-click the setting Account Handling for New Users and set it to:

   Create New Account

   CA DataMinder creates a new user account automatically. User names for the new accounts are generated automatically based on Microsoft Windows authentication and include a domain prefix, for example, UNIPRAXIS\srimmel.

   Ignore

   A new account is not created but the user is permitted to continue using their browser and email application.

   Disable Internet applications

   A new account is *not* created and the user is blocked from using their browser and email application.

5. **New machine accounts:** Configure how the CMS handles new client machines. In the right-hand pane, double-click the setting Account Handling for New Client Machines and set it to:

   Create New Account

   CA DataMinder creates a new machine account automatically.

   Ignore

   A new account is not created but users on the machine are permitted to continue using their browser and email application.

   Disable Internet applications

   A new account is not created and users on the machine are blocked from using their browser and email application.

Configure Event Purging

Important! Before using CA DataMinder for the first time after installation, we strongly recommend you turn on event purging in your common client machine and common gateway policies.

After installing CA DataMinder, you need to turn on event purging. You need a separate purging strategy for your CMS, which holds captured data for your entire organization, and one or more strategies for your gateways and client machines.

- **CMS purges:** A strategy that meets regulatory requirements on the retention of electronic communications typically requires scheduled event purges. These run at regular intervals and are configurable in the CMS machine policy.
- **Gateway and client machine purges:** Events stored on these machines are eventually replicated up to the CMS, but purging prevents free disk space falling to dangerously low levels (with the attendant risk of the CA DataMinder infrastructure being suspended—see task **5** Configure the management of free disk space on servers). By default, events are purged automatically after replication to the parent server, though you can disable this feature if required and run scheduled purges instead.

To configure the relevant machine policies

1. Expand the Machine Administration branch .
2. To configure event purging for:
   - **The CMS**, right-click the CMS and choose Edit Policy.
   - **All gateways**, right-click the CMS and choose Edit Common Gateway Policy.
   - **All client machines**, right-click the CMS and choose Edit Common Client Policy.
3. In the Machine Policy Editor, browse to the Data Management folder.

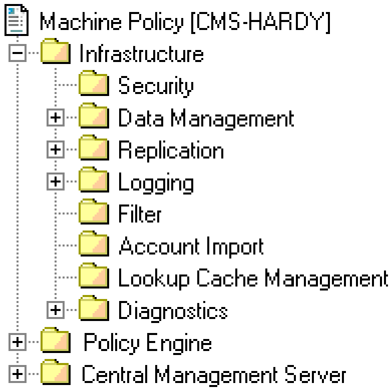

4. **Set the purge frequency:** You can configure purges to run immediately after the data has been replicated, or you can schedule purges to run at regular intervals.

   Scheduled purges

   To schedule regular purges, set Purge Events on Replication? to False (clear the check box). Then configure the settings for the minimum retention period (step 5) and the event purge frequency and time.

   Purging after replication

   For gateways and client machines, the Purge Events on Replication? setting defaults to True. This purges events stored locally as soon as they have been replicated to the parent server. This setting is ignored on CMSs; it is intended for use only with gateways and client machines.

5. **Set the minimum retention period:** This determines how long events are retained before they become eligible for purging. The default is 1095 days (three years). You may need to change this period on the CMS to meet regulatory requirements. Likewise, if you schedule purges on gateways and client machines, we recommend you reduce this period to prevent a shortage of free disk space.
6. **Configure purge performance:** Other settings in the Data Management folder provide further control over purge operations. For example, you can suspend the CA DataMinder infrastructure during purge operations or you can specify a purging timeout.
7. Save the policy. Event purging is turned on as soon as the new settings replicate to the target CA DataMinder machines.

Configure the Management of Free Disk Space on Servers

Important! Make these changes in both the CMS and common gateway policies.

CA DataMinder automatically monitors the level of free disk space on the drive hosting the data folder (for the CMS and gateways) or on the drive hosting the installation folder (for client machines).

Settings in the machine policy determine how low free disk space can fall on these drives before the CA DataMinder infrastructure is suspended, and the level at which the infrastructure automatically restarts. By default, these settings are optimized for client machines so you need to adjust these values in your CMS policy and common gateway policy.

Calculate Disk Space Values

The amount of free disk space you require essentially depends on the expected disk usage. You can calculate how much this will be using two factors:

- The rate at which CA DataMinder imports data.
- The Disk Space Check Interval. That is, how frequently CA DataMinder checks the amount of available free disk space. This setting is in the CMS policy—see next section.

For example, if CA DataMinder imports events at a rate of 50 MB per minute and the Disk Space Check Interval is set to 10 minutes, then you can potentially write 500 MB of data between disk space checks. This would mean that the Disk Space Error Level had to be at least 500 MB. To roughly calculate the Disk Space Warning Level value, multiply the Disk Space Error Level by five.

To set disk space levels

1. Expand the Machine Administration branch .
2. Right-click the CMS and choose Edit Policy or Edit Common Gateway Policy. This opens the Machine Policy Editor.
3. In the Machine Policy Editor, browse to the Data Management folder (see the policy tree in Configure event purging).
4. Specify how often free disk space is checked. Set Disk Space Check Interval to one minute.

5. Set the Disk Space Important Level, for example, to 2 GB. When CA DataMinder detects that free disk space has fallen below this level, it adds a warning to the Audit log file.

   Note: The warning level also represents the 'safe' level at which the infrastructure automatically restarts following a suspension and subsequent recovery in free disk space.

6. Set the Disk Space Error Level, for example, to 1 GB. When CA DataMinder detects that free disk space has fallen below this level, the CA DataMinder infrastructure is suspended.

7. Save the policy.

Configure the Common Client and Gateway Policies

New client machines automatically inherit the common client policy, while new gateway servers automatically inherit the common gateway policy. Configure these policies to suit your requirements.

Purging strategy was discussed in task 4 Configure event purging, but you will probably want to amend other policy areas. For example, the Infrastructure folder contains Security and Replication subfolders:

- Security settings define the encryption strategy for data sent across the network to the CMS from gateways and client machines.
- Replication settings determine how often CA DataMinder machines send notification of newly captured data or local infrastructure changes. These notification messages act as triggers for data replication.

To edit the common client or common gateway policies

1. Expand the Machine Administration branch .

2. Right-click the CMS and choose Edit Common Client Policy or Edit Common Gateway Policy. This opens the Machine Policy Editor.

3. Edit the policy settings required.

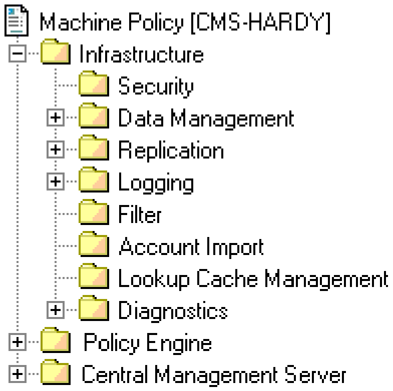

Machine policy folders

Synchronize the Clocks on Your CA DataMinder Machines

When an event is captured on a client machine, it is time stamped. This time stamp is preserved when the event is replicated up to the CMS. If a machine is set to the wrong date, incorrect capture dates are saved with any events captured on that machine. For this reason, it is important that all of your CA DataMinder machines are set to the correct date and local time.

This normally happens automatically. If your CA DataMinder machines are all members of a domain, their clocks are probably synchronized automatically by a network time server. Machines that are not members of a domain are probably synchronized by an Internet time server.

But if a CA DataMinder machine does not have a continuous Internet connection, this automatic synchronization may not occur, for example, because a firewall blocks time synchronization. Similarly, if a machine is set to the wrong date set, clock synchronization is blocked.

Therefore, if CA DataMinder is installed on any machines that are not members of a domain, make sure that the owners of these machines understand the need to keep their machine set to the correct date and time.

Note the following:

- Such problems cannot affect event auditing. When a reviewer updates an event's audit trail in the Data Management console, the audit entry is always time stamped on the CMS.
- When displaying event capture times, CA DataMinder automatically takes account of time zone differences. For details about time zone handling, see the Data Management Console online help; search the index for 'time zones'.

Configure the Policy for the Default User Group

Important! Before using CA DataMinder for the first time after installation, we strongly recommend you choose a **new** default group and define a restrictive policy for this group.

A user group is a collection of associated users that share a common policy. Each group has its own customizable policy, providing you with a centralized but highly flexible method of user administration. When new users add themselves to CA DataMinder, they are automatically assigned to the **default group**. You make any user group the default group.

Why is this a problem? The default group is effectively a holding group until you can move new users into more appropriate groups. But when you use CA DataMinder for the first time, there is only one existing group. This is the 'Users' group and so it is automatically set to be the default group. Of necessity, 'Users' has—and must have—a non-restrictive policy: no settings are disabled, enforced or hidden.

This means any new user who inherits this policy has complete freedom to change any setting in their policy. In other words, they could potentially define their own policy to dodge the rules in your organization governing acceptable Web and email usage. But you can easily prevent this by choosing a default group with a restrictive policy. That is, key policy settings are enforced, hidden or disabled. This ensures that new users adhere to the rules governing acceptable Web and email usage.

Predefine the Default Group

For deployment operations based on Msiexec.exe, you can use a variable to customize the parent group for newly created users. This allows different teams or departments to install CA DataMinder from separate source images so that their respective users are added automatically to separate groups. For details, see general variable WGNDEFAULTUSERGROUPPATH in Command line parameters for Msiexec.exe (see page 531).

To create a new default user group

1. In the Administration console, expand the User Administration branch.
2. Click Edit, New Group.
3. Select the new group and click Edit, Set As Default.

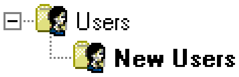

Example default group: All self-enrolled new users are added to this group.

To edit the default group policy

1. Select the default group and click Edit, Edit Policy.

   The User Policy Editor opens.

2. Edit the policy to suit your requirements. For example, if you want to:

   - **Disable** a folder, click Disable. When you disable a folder, CA DataMinder ignores all settings in the folder itself and its subfolders.
   - **Enforce** a current folder plus its subfolders, right-click the folder then click Enforce Branch.
   - **Disable Web and email applications** when the CA DataMinder infrastructure is not running, edit the Infrastructure Failure setting. Find this in the Initialization subfolder. (Click Edit, Find to locate this setting).

Create and Organize a Hierarchy of User Groups

Note: CA DataMinder integrates directly with CA Identity Manager. This integration allows you to use CA Identity Manager to maintain your CA DataMinder user accounts. For more information, see the Technical Information (see page 582) section of this guide.

You need to create a set of user groups and configure the policies for these groups. You can create as many groups as you need and arrange them in any way you want. For example, you can organize users into groups based on location, job, or purchasing permissions:

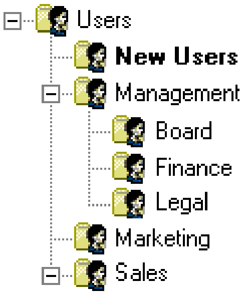

- Users
  - **New Users**
  - Management
    - Board
    - Finance
    - Legal
  - Marketing
  - Sales

Use the Account Import feature to create a groups hierarchy

To simplify mass deployments, Account Import enables administrators to import user details into CA DataMinder from an external Lightweight Directory Access Protocol (LDAP) directory or a source file. Account Import can import new users and groups into the existing CA DataMinder user hierarchy, or it can reorganize existing users to synchronize them with an external hierarchy. It can also import user attributes such as email addresses and employee IDs. For example, you can import your existing user hierarchy from Microsoft Active Directory or Domino Server.

Account Imoprt also enables you to synchronize users' email addresses in the CMS database with addresses in an external source such as Active Directory. Such synchronization is essential for CA DataMinder features that rely on email address mapping (see page 485), especially policy engines.

For full details, see the Account Import (see page 245) section of this guide.

Manually create a group hierarchy

You can also use the Administration console to manually create user groups and organize these groups into a hierarchy. For details, see the User Groups section of the *Adminstration Guide*.

More information:

Account Import (see page 245)

Create Your Administrators and Managers

Note: You may prefer to leave this task until you have deployed CA DataMinder to your client machines, in order to allow users to enroll themselves as CA DataMinder users. This entails setting up the CMS machine policy to permit self-enrollment.

For CA DataMinder installations with large numbers of client machines, you will almost certainly need to share the administrative workload with selected colleagues. You will also need to provide senior executives with the ability to manage users and extract information held in CA DataMinder databases.

You can promote ordinary users into administrators or managers by granting them administrative privileges. To limit the scope of their authority, you can withhold specific privileges and assign a management group (this limits which groups they can manage).

To assign administrative privileges

1. Right-click a user and choose Properties.
2. In the Properties dialog, click the Privileges tab.
3. Select the administrative privileges as required.

To assign a management group

1. Right-click a user and choose Properties.
2. In the Properties dialog, click the Details tab.
3. Click the Browse button and choose which group you want as the management group.

Note: You will need to provide administrators and managers with details of their CA DataMinder account (their user name and password) before they can run the Administration console.

More information:

Configure Your CMS Machine Policy to Handle New Accounts (see page 73)

Set Up Support for Unicode Characters

CA DataMinder consoles support Unicode character sets. For example, you can search in the iConsole emails associated with trigger names that were defined using Korean characters in the user policy.

CMS and Gateways

- **Oracle:** To implement Unicode support on Oracle CMSs or gateways, set up the database for CA DataMinder to use UTF-8 encoding for the DBMS code page. For details, see the *Database Guide*; search for ‘UTF-8’.
- **SQL Server:** There is no equivalent requirement for SQL Server CMSs or gateways. SQL Server databases automatically support Unicode characters.

Endpoint Computers

Implement Unicode support on all CA DataMinder endpoint computers running non-double byte character set versions of Windows (that is, non-DBCS Windows) and which are likely to capture events containing Unicode characters. For example, this applies to Japanese email trigger details captured on a computer running an English version of Windows.

To implement Unicode support:

1. Stop the 'CA DataMinder infrastructure' service. From a command prompt, run:
   ```
   net stop wgninfra
   ```
2. Edit the startup.properties file. Find this file in the \system subfolder of the CA DataMinder installation folder.

3. Open this file and add the following line to the [Database] section:

   ```
   db.charset=UTF-8
   ```

4. Restart the CA DataMinder infrastructure service. From a command prompt, run:

   ```
   net start wgninfra
   ```

Important! Do *not* make this change to startup.properties on CA DataMinder computers running Japanese, Korean or Chinese versions of Windows.

Set up Policy Engines

Setting up policy engines requires a number of pre- and post-installation tasks:

1. Set up the required user accounts.

   Before you can deploy your policy engines, you must specify a Windows domain user that allows the policy engines and hub to communicate. You must also create a new CA DataMinder user account.

2. Edit the machine policy on each policy engine host server.

   After installing a policy engines, you must edit its machine policy. In particular, you must configure these settings:

   - Unknown Internal Sender
   - External Sender
   - Internal Email Address Pattern
   - Default Policy for Files

3. Edit the registry on each policy engine host server to enable user lookup operations.
4. (Optional) You can configure CA DataMinder to detect, analyze, and apply policy to emails encrypted by Voltage SecureMail. As part of the integration setup for SecureMail, you must edit registry values on your policy engines to identify the SecureMail server.

More information:

Deploying Policy Engines (see page 338)
Specify PE User Accounts (see page 340)
Configure the Local Machine Policy (see page 343)
Configure the Policy Engine Registry Values (see page 347)

Integrate with Third-Party Object Storage Solutions

CA DataMinder can integrate with third party object storage solutions to allow storage management, archiving and retrieval of large volumes of captured data. Specifically, the following third party solutions can integrate with CA DataMinder to work as an alternative object store:

- EMC Centera
- IBM DB2 Content Manager
- NetApp SnapLock.

In each case, only events captured by CA DataMinder **after** integration has been set up will be migrated to the third party object store. Events captured and replicated to the CMS before integration has been set up will not be migrated. For this reason, we recommend that you set up integration as soon as possible after deploying CA DataMinder.

More information:

Integrating with EMC Centera (see page 309)
Integrating with IBM DB2 Content Manager (see page 319)
Integrating with NetApp SnapLock (see page 327)

Install iConsole Searches

After deploying the iConsole, you must install and publish the default iConsole event searches.

More information:

Installing Searches, Reports, and the Dashboard (see page 175)

Configure Event Auditing Labels

Full CA DataMinder auditing features are available in the iConsole, but administrators must first configure audit status labels and the contents of the auditing dialogs. To do this, they must use the Administration console. If these audit features are not fully configured, then they will not be available in the iConsole. That is, reviewers will not be able to audit events. We therefore recommend that you configure these event auditing features before your reviewers start using CA DataMinder to audit captured events.

Full details are available in the Administration Console online help; search the index for ‘event auditing, configuring audit features’.

Chapter 6: Advanced Encryption Mode

This section contains the following topics:

Overview

You can deploy CA DataMinder in Advanced Encryption mode. When deployed in this mode, CA DataMinder uses Transport Layer Security (TLS) and certificates to enable FIPS 140-2 compliant data transfers between CA DataMinder machines.

CA DataMinder machines use a single enterprise certificate across the CA DataMinder enterprise. There is no authentication of individual machines. Any machine possessing the enterprise certificate and its associated private key can communicate with any CA DataMinder machine that uses the same certificate.

More information:

What Is FIPS 140-2? (see page 87)
Which FIPS Certified Cryptographic Modules Are Used? (see page 88)
Which Encryption Algorithms Are Used? (see page 88)
What Data Is Encrypted? (see page 89)
Can I Convert My Existing CA DataMinder Deployment To Be FIPS Compatible? (see page 89)
Key Points (see page 89)

What Is FIPS 140-2?

The Federal Information Processing Standards (FIPS) 140-2 publication is a security standard for the cryptographic libraries and algorithms that a product should use for encryption.

On Federal networks, FIPS 140-2 encryption affects the communication of all sensitive data between components of CA products and between CA products and third-party products. FIPS 140-2 specifies the requirements for using cryptographic algorithms within a security system protecting sensitive, unclassified data.

Which FIPS Certified Cryptographic Modules Are Used?

When deployed in Advanced Encryption Mode, CA DataMinder uses the Advanced Encryption Standard (AES) adopted by the US government. Specifically, it uses the TLS protocol to transfer sensitive data between machines.

To allow this, CA has licensed the RSA BSAFE Crypto-J 4.0 and SSL-J 5.1.1 cryptographic libraries. These libraries have been validated as meeting the FIPS 140-2 Security Requirements for Cryptographic Modules. The FIPS 140-2 Certificate number for these libraries is 1291. CA DataMinder uses these libraries to encrypt sensitive data being transferred between machines.

Note: The term ‘FIPS 140-2 compliant’ and the FIPS 140-2 standard relate to the requirements of a cryptographic module (that is, one actually implementing cryptographic algorithms) and not to an application's use of cryptography. An application's use of cryptography is guided by what a FIPS 140-2 compliant cryptographic module will provide. CA DataMinder will not be a FIPS 140-2 compliant product. Instead, it will use a FIPS 140-2 compliant cryptographic module (RSA BSAFE Crypto-J 4.0) and will only use cryptographic algorithms, for purposes such as symmetric encryption, that are approved by FIPS 140-2. We therefore use the term ‘FIPS compatible’ to describe the CA DataMinder support for FIPS 140-2.

Which Encryption Algorithms Are Used?

In Advanced Encryption Mode, CA DataMinder uses these encryption algorithms:

- **Data Transfers:** Sensitive data sent across the network between CA DataMinder machines is encrypted with TLS, using AES 128-bit as the symmetric cipher algorithm.
- **Captured data:** Blob files (binary large objects) containing captured data are encrypted using AES 128-bit as the symmetric cipher algorithm. They are saved in the CMS data store.
- **Local encryption keys:** These keys, used to encrypt captured data and policy data, are themselves encrypted with a master key using the 3DES (Triple Data Encryption Standard) algorithm.

What Data Is Encrypted?

When two CA DataMinder machines transfer data using the Java RMI service, the data is encrypted with TLS.

In practical terms, this means that any potentially sensitive data is encrypted. The cryptographic modules are used to encrypt communications between machines running the CA DataMinder Infrastructure, plus data stored by the infrastructure such as encryption keys and Binary Large Object files (blobs) containing captured data.

In terms of its cryptographic boundary, CA DataMinder is self-contained. It has no dependency on an external Public Key Infrastructure (PKI).

Note: A blob file contains the text content of a captured file, e-mail or Web page, stored in CA DataMinder format.

Can I Convert My Existing CA DataMinder Deployment To Be FIPS Compatible?

No. Such migrations are not supported.

Note: In theory, it is possible to convert your existing CA DataMinder machines to run in Advanced Encryption Mode. But in practice, this requires you to take all CA DataMinder machines offline and reconfigure them before restarting CA DataMinder. Any machines not changed at this point would cease to communicate with other machines in the CA DataMinder enterprise. For a typical CA DataMinder enterprise, with hundreds or thousands of protected machines, this is unlikely to be practicable.

Key Points

Key Points

- CA DataMinder uses TLS and certificates to enable FIPS 140-2 compliant data transfers between machines.
- CA DataMinder machines use a single enterprise certificate and private key across the CA DataMinder enterprise.
- There is no authentication of individual machines. Any machine possessing the enterprise certificate and its associated private key can communicate with any CA DataMinder machine that uses the same certificate.

Certificates and Key Store

- A self-signed root certificate and a single enterprise certificate, and associated key pairs, are generated before installing CA DataMinder.
- A Key Store containing the root certificate, the enterprise certificate, and the private key for the enterprise certificate key pair is deployed throughout the CA DataMinder enterprise.
- Possession of the Key Store is enough to permit any CA DataMinder machine to communicate with other CA DataMinder machines.
- The critical files (keystore.dat, revocation.properties, and wigan.java.security) are stored in the CA DataMinder \data and \system folders. You must secure these file locations as part of the general machine hardening process after deployment.

Deployment

- Advanced Encryption Mode must be enabled at install time, and if enabled must be enabled on every CA DataMinder machine. There is no backward compatibility with existing CA DataMinder installations.
- There is no automatic integration with third-party Public Key Infrastructures (PKIs).
- Mechanisms to replace the enterprise certificate and its key-pair are not built into CA DataMinder. Instead, you must use a manual process, or a third-party software distribution mechanism, in conjunction with the OpenSSL.exe utility (provided by CA).

Advanced Encryption Certificates

This section contains the following topics:

More information:

How Is FIPS Implemented?

In Advanced Encryption Mode, each CA DataMinder machine holds a copy of the Key Store. This contains the root certificate, the enterprise certificate, plus the private key for the enterprise certificate key pair.

All network communications with the potential to transmit sensitive data are protected by TLS, using AES 128 Bit.

Deployment Architecture

The following diagram summarizes the FIPS 140-2 implementation.

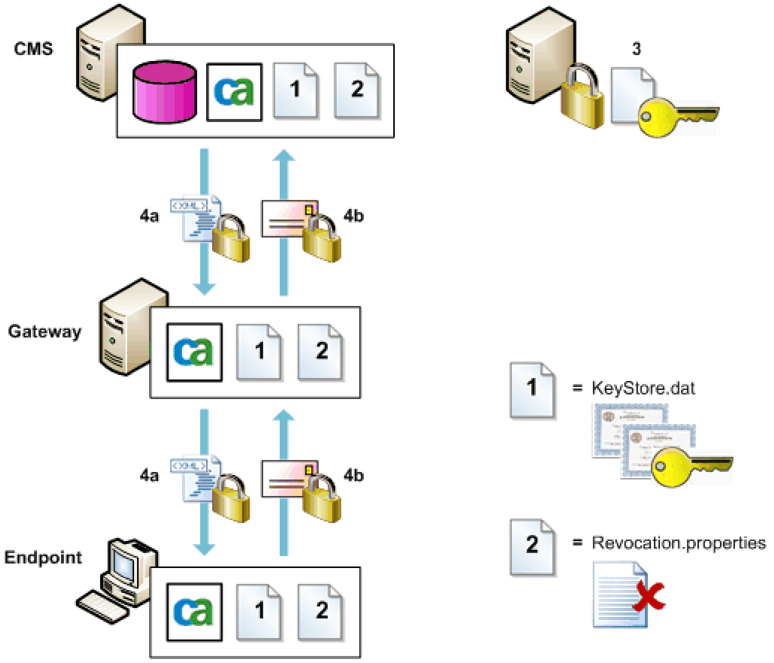

Deployment Architecture: Advanced Encryption Mode

1. **KeyStore.dat.** This is the Key Store file. It contains the root certificate, the enterprise certificate, and the private key for the enterprise certificate key pair. A copy is held on each machine in your CA DataMinder enterprise.
2. **Revocation.properties.** This is the Revocation List file. It contains a list of all revoked enterprise certificates. A copy is held on each machine in your CA DataMinder enterprise.
3. **Root certificate private key.** This private key must be kept separate from your CA DataMinder enterprise on a secure server. It is used when you create replacement enterprise certificates.
4. **Encrypted sensitive data.** This includes infrastructure changes such as policy edits or user account updates (**4a**) replicated from the CMS to gateway servers and endpoint machines. It also includes captured data, such as emails, files or Web activity (**4b**) replicated from endpoint machines and gateway servers to the CMS.

Root and Enterprise Certificates

CA DataMinder uses certificates with a two-level hierarchy:

- **Root certificate:** This serves as the trusted root certificate. It is used to sign the Enterprise certificate. The root certificate enables CA DataMinder machines to authenticate each other before transferring sensitive data. You must create a self-signed root certificate on a secure server.
- **Enterprise certificate:** This certificate is signed by the trusted root certificate. CA DataMinder machines use this certificate to encrypt data transfers between machines.

  When you update the enterprise certificate, its serial number is incremented by 1 and the previous serial number is added to the Revocation List (see below).

The root certificate, plus the enterprise certificate and the private key from its associated key pair, are then added to the Key Store and distributed to all CA DataMinder machines. This enables any machine in the CA DataMinder enterprise to use TLS to communicate with any other CA DataMinder machine.

Certificate Security

Because every CA DataMinder machine has a copy of the same enterprise certificate, the security of any data transfer is at risk if the enterprise private key is compromised. If this happens, you will need to distribute a new enterprise certificate and private key to all CA DataMinder machines. As with any PKI, we recommend that you regularly replace the enterprise certificate (that is, revoke the existing certificate and issue a new one). The CA DataMinder scheme has been designed to make this as simple as possible.

Revocation List

The Revocation List identifies certificates that have been marked as revoked. It holds the serial numbers of revoked certificates.

The Revocation List is a Java properties file named revocation.properties. It is stored in the CA DataMinder \data folder. You must protect this file with the same level of operating system protection as the Key Store file.

Folders Used By Certificate Scripts

When you run the certificate generation scripts, GenerateRootCert.bat and GenerateKeyStore.bat, three subfolders are created: \tmp, \persist, and \output.

For example, if the scripts are stored in an \AdvancedEncryption folder, they will create subfolders such as \AdvancedEncryption\persist.

\tmp subfolder

This holds temporary files while the script is running. When the script completes, this subfolder should be empty. If it is not, you can safely delete its contents.

\persist subfolder

This subfolder is critical. It contains files needed to update the certificates and Key Store at a later date. It contains: a script log file; a text file with the serial number of the most recent enterprise certificate; the self-signed root certificate containing its public key; the root key pair, encrypted; and the enterprise certificate.

The file containing the encrypted root key pair must be kept secure because it is needed to sign every enterprise certificate generated. If this critical file is lost, the Key Store will need to be regenerated and redeployed to every machine in the CA DataMinder enterprise.

Important! Never delete any files in this folder!

\output subfolder

This subfolder contains keystore.dat and revocation.properties. Whenever you update your enterprise certificate after the initial CA DataMinder deployment, you will need to deploy these files to the \data folder on each CA DataMinder server and client machine using a secure software delivery mechanism.

Note: The \data folder holds all the configuration data and captured data used by your CA DataMinder enterprise. By default, when you install a CA DataMinder server or client machine this folder is added as a subfolder in the CA DataMinder installation folder. But you can rename it and locate it anywhere suitable on your network.

How Do I Deploy CA DataMinder In Advanced Encryption Mode?

For CA DataMinder to be compatible with FIPS 140-2, you deploy it in Advanced Encryption Mode. This section describes the deployment procedure.

Follow these steps:

1. Designate a secure server that is separate from your intended CA DataMinder enterprise.
2. Generate the self-signed root certificate.
3. Generate the Key Store and Revocation List.
4. Deploy your CA DataMinder servers and client machines.
   a. Create new administrative installation source images.
   b. Customize the new source images.
   c. Install the servers and client machines from the appropriate source image.
5. Confirm that encryption is correctly configured in the machine policy for all your CA DataMinder servers and client machines.
6. Secure the critical Advanced Encryption files on your CA DataMinder servers and client machines so that they can only be accessed by the CA DataMinder infrastructure.

More information:

Designate a Secure Server (see page 95)
Generate the Root Certificate (see page 95)
Generate the Key Store and Revocation List (see page 96)
Deploy CA DataMinder Machines (see page 97)
Ensure Machine Policy Is Correctly Configured (see page 99)
Secure the Critical Advanced Encryption Files (see page 100)

Designate a Secure Server

Important! It is essential that the root certificate's private key is kept secure.

The CA DataMinder Advanced Encryption Mode Enhancement solution package contains the files and utilities that you will need to deploy CA DataMinder in Advanced Encryption mode.

We recommend that you create an **\AdvancedEncryption** folder on a secure server that is separate from your intended CA DataMinder enterprise and then copy the required files and utilities to this folder. This ensures that, when you generate the root certificate and Key Store, these files are saved to a location that is secure.

Generate the Root Certificate

To generate the root certificate, run the batch file supplied with the CA DataMinder distribution media.

To generate the root certificate

1. From a command prompt on your designated secure server, change to the \AdvancedEncryption folder.
2. From a command prompt, run GenerateRootCert.bat.
3. When prompted, enter and confirm a *strong* passphrase to secure the root key pair.

   You will need to supply this passphrase later, when you self-sign the root certificate, and when you sign the enterprise-wide certificate.

   Important! This passphrase will not be stored anywhere. If you forget or lose it, you will need to regenerate all certificates and key stores!
4. GenerateRootCert.bat generates the root certificate and a key pair (root.crt and root.key respectively).

   These files are saved in the \AdvancedEncryption\Persist subfolder on your secure server. They will be used to generate the enterprise-wide certificate, the Key Store file, and the Revocation List file.

   Important! You must retain the contents of the \Persist subfolder for the lifetime of the CA DataMinder deployment. These contents are needed each time you update the enterprise wide certificate.

More information:

Folders Used By Certificate Scripts (see page 93)

Generate the Key Store and Revocation List

To generate the Key Store and Revocation List files, you run a batch file supplied with the CA DataMinder distribution media.

To generate the Key Store and Revocation List files

1. On your designated secure server, browse to the \AdvancedEncryption folder.
2. Run GenerateKeyStore.bat.
3. When prompted, enter the root certificate passphrase.

   The batch file now generates keystore.dat and revocation.properties. These files are saved in the \AdvancedEncryption\output subfolder on your secure server.

   The enterprise-wide certificate is stored in keystore.dat and has a serial number of 1.

   The certificate itself is saved as server1.crt. It is saved in the \AdvancedEncryption\persist subfolder on your secure server.

More information:

Folders Used By Certificate Scripts (see page 93)

Deploy CA DataMinder Machines

To deploy CA DataMinder in Advanced Encryption Mode, first perform an administrative installation to your network of each server and client msi you intend to deploy. The administrative installation extracts the contents of the original CA DataMinder Windows Installer packages to a network folder specified by you, and in a format that can be patched to support Advanced Encryption mode.

By performing the administrative installation, you create the basic source images that you use to install CA DataMinder servers and client machines.

After you create your CA DataMinder source images, patch and customize them, deploy CA DataMinder in Advanced Encryption mode.

Finally, you can install CA DataMinder servers and client machines directly from the patched and customized source images.

1. **Create new administrative installation source images.**

   To create the source images for your CA DataMinder servers and client machines, perform an administrative installation.

   Run the following commands to create administrative installation source images for CA DataMinder servers and client machines. These commands launch the installation wizard, which prompts for a target folder for the source images:

   ```
   msiexec /a <Path_source>\server.msi
   ```

   ```
   msiexec /a <Path_source>\client.msi
   ```

   ```
   msiexec /a <Path_source>\client_x64.msi
   ```

   <Path_source>\server.msi

   Identifies the Windows Installer package for servers on your CA DataMinder distribution media.

   <Path_source>\client.msi

   Identifies the Windows Installer package for client machines on your CA DataMinder distribution media.

   <Path_source>\client_x64.msi

   Identifies the Windows Installer package for 64-bit client machines on your CA DataMinder distribution media.

2. **Customize the administrative installation source images.**

   Customize the administrative installation source images so that they install the Key Store (and associated components) on all CA DataMinder servers and client machines. Run a script supplied with the CA DataMinder Advanced Encryption Mode Enhancement solution package.

   From a command prompt on your designated secure server, change to the \AdvancedEncryption folder and run the following commands:

```
EnableAdvancedEncryption.vbs /package:<Path_admin>\server.msi
/files:<path_keystore>
```

```
EnableAdvancedEncryption.vbs /package:<Path_admin>\client.msi
/files:<path_keystore>
```

```
EnableAdvancedEncryption.vbs /package:<Path_admin>\client_x64.msi
/files:<path_keystore>
```

/package:<Path_admin>\server.msi

Identifies the server source images that you created in step 1.

/package:<Path_admin>\client.msi

Identifies the client machine source images that you created in step 1.

/package:<Path_admin>\client_x64.msi

Identifies the 64-bit client machine source images that you created in step 1.

/files:<path_keystore>

Identifies the path to the folder containing the Key Store and Revocation List files, keystore.dat, and revocation.properties. When you generated these files (see the previous section), they were saved in the \AdvancedEncryption\output subfolder on your secure server.

3. **Deploy your CA DataMinder servers and client machines.**

   Important! Deploy your CMS before deploying the other servers and client machines!

   After you customize the administrative installation source images, you can deploy CA DataMinder servers and client machines using your preferred deployment methods. For example, use the following command syntax to deploy client machines as part of a managed deployment:

   ```
   msiexec /i <Path_admin>\client.msi WGNPARENTSERVERNAME=<Server>
   ```

   <Path_admin>\client.msi

   Identifies the client machine source image that you patched in step 2 and customized in step 3.

   <WGNPARENTSERVERNAME>=<Server>

   Identifies the parent gateway or the CMS.

   During the installation, the following critical files are installed: keystore.dat, revocation.properties, and wigan.java.security. As the final step in overall deployment, restrict access to these files.

Ensure Machine Policy Is Correctly Configured

The two settings in CA DataMinder machine policy that control data encryption are **Communications Encryption** and **Encrypt Stored Data?**. Find these settings in the Security folder of the machine policy.

When CA DataMinder runs in Advanced Encryption Mode, Encrypt Stored Data? must be set to True (this is its default value), while Communications Encryption is not used. Consequently, you do not normally need to change these settings after deploying CA DataMinder.

Machine Policy Setting 'Communications Encryption'

This setting covers encryption for network communications. It specifies the level of network encryption (none, low, medium, or high) for data sent between CA DataMinder machines. However, CA DataMinder ignores this setting when it runs in Advanced Encryption Mode. This is because network encryption using TLS is an integral part of Advanced Encryption Mode and cannot be disabled. Instead, the infrastructure logs an entry in the CA DataMinder Activity Log file indicating that it is running in this mode.

Machine Policy Setting 'Encrypt Stored Data?'

This setting covers stored data encryption. The machine policy setting specifies whether to encrypt Binary Large Object files (blobs) containing captured data. This setting remains active and must be set to True (the default) when CA DataMinder runs in Advanced Encryption Mode. This is because FIPS 140-2 states that all sensitive data must be encrypted with an approved algorithm.

Important! CA DataMinder administrators must therefore ensure that this setting is never set to False!

Note: If Encrypt Stored Data is inadvertently set to False, you will need to reset it to True across all machines in your CA DataMinder enterprise. To do this, you will need to edit this setting in the CMS machine policy, the common gateway policy and the common client policy. All gateway servers inherit the common gateway policy, and all client machines inherit the common client policy. For details about editing machine policies, see the online help for the CA DataMinder Administration console.

Secure the Critical Advanced Encryption Files

When you deploy your CA DataMinder servers and client machines in Advanced Encryption Mode, three critical files are installed. These are keystore.dat and revocation.properties, stored in the CA DataMinder \data folder, and wigan.java.security, stored in the CA DataMinder \system folder. You must secure these files to prevent unauthorized access.

To secure the critical Advanced Encryption files

On each CA DataMinder server and client machine, you must configure Windows security for each of these critical files so that:

- Each file can be accessed *only* by the Windows logon account used by the CA DataMinder infrastructure service. The infrastructure service only requires Read access to these files.
- No other process on the system is permitted to access these files.

How Do I Replace Enterprise Certificates?

Because every CA DataMinder machine has copies of the same enterprise certificate, if the enterprise private key is compromised then the security of any data transfer is at risk. As a security precaution, and as with any PKI, we therefore recommend that you periodically replace the enterprise certificate. The CA DataMinder scheme has been designed to make this as simple as possible.

The main steps are:

1. Create three machine searches for use in the CA DataMinder Administration console. You will use these searches to monitor progress across your CA DataMinder enterprise when you update your enterprise certificate.
2. Update the Key Store and Revocation List. You will do this on your secure server using the CA-supplied script, GenerateKeyStore.bat.
3. Deploy the new Key Store and Revocation List. This is a multi-step procedure designed to minimize disruption to your CA DataMinder enterprise.

Create Custom Machine Searches

Before you generate the replacement certificate, you need to create three machine searches that you run in the CA DataMinder Administration console. When you replace the enterprise certificate on your CA DataMinder servers and client machines, you will use these searches to monitor progress across your CA DataMinder enterprise.

To create these machine searches, you will copy SQL snippets (database search queries) from a CA-supplied file and save them as the following three custom searches in the Administration:

- All servers with an out-of-date Key Store
- All client machines with an out-of-date Key Store
- All machines with out-of-date Revocation List

To create custom machine searches

1. Create an Administration console search for any servers with an out-of-date Key Store.

    a. In the CA DataMinder Advanced Encryption Mode Enhancement solution package, open the CertificateSearches.txt file and copy the ‘All servers with an out-of-date Key Store’ SQL snippet.

    b. In the CA DataMinder Administration console, create a new administrative search.

    c. In the Administration Search dialog, go to the SQL tab and paste in the 'All servers with an out-of-date Key Store’ SQL snippet.

    d. Save the new search as ‘All servers with an out-of-date Key Store’.

    e. The new search is added to the Custom Searches folder in the Administration console.

    The new search is saved on the machine hosting the Administration console and is only available to you when you run the console on that machine.

    Note: For assistance with making the search available when colleagues run the Administration console on other machines, please contact CA Technical Support: http://ca.com/support (see page 21)

2. Create an Administration console search any for client machines with an out-of-date Key Store.

    Repeat step 1, but using the ‘All client machines with an out-of-date Key Store’ SQL snippet and saving the search as ‘All client machines with an out-of-date Key Store’.

3. Create an Administration console search for any servers and client machines with an out-of-date Revocation List.

    Repeat step 1, but using the ‘All machines with out-of-date Revocation List’ SQL snippet and saving the search as ‘All machines with an out-of-date Revocation List’.

Update the Key Store and Revocation List

Generate the replacement Key Store and Revocation List.

1. From a command prompt on your designated secure server, change to the \AdvancedEncryption folder.
2. Run GenerateKeyStore.bat.
3. When prompted, enter the root certificate passphrase.

   The batch file now generates keystore.dat and revocation.properties and saves these files in the \AdvancedEncryption\output subfolder on your secure server.

   The serial number for the enterprise certificate is incremented by 1. The certificate is saved in the new Key Store file, keystore.dat.

   The old serial number for the enterprise certificate is appended to the Revocation List in revocation.properties.

   The new enterprise certificate is saved as server<n+1>.crt, where <n> is the number used by the most recent certificate file. It is saved in the \AdvancedEncryption\persist subfolder on your secure server. For example, if the \perist folder already contains server1.crt and server2.crt, the newest replacement certificate will be saved as server3.crt.

More information:

Folders Used By Certificate Scripts (see page 93)

Deploy the New Key Store and Revocation List

Deploy the updated Key Store and Revocation List. You must follow the steps below to enable your CA DataMinder enterprise to continue with minimal disruption during the certificate deployment.

Important! Do not try to optimize the following procedure. For example, do not try to combine steps 2 and 6. The procedure below is explicitly designed to minimize the steps needed to replace certificates on your client machines while retaining a functioning CA DataMinder enterprise.

1. Distribute keystore.dat to the CA DataMinder \data folder on the CMS only. Then restart the CA DataMinder infrastructure service on the CMS, or reboot the server.
2. Distribute keystore.dat to the CA DataMinder \data folder on all gateway servers. Then restart the CA DataMinder infrastructure service on each server, or reboot them.
3. In the Administration console, run the 'All servers with out-of-date Key Store' custom search.

   When you can confirm that the CMS and all gateway servers have the new Key Store (that is, when this search returns zero results), continue to the next step.
4. Distribute keystore.dat and revocation.properties to the CA DataMinder \data folder on all client machines. Then restart the CA DataMinder infrastructure service on each machine, or reboot them.
5. In the Administration console, run 'All client machines with out-of-date Key Store' custom search.

   When you can confirm that all client machines have the new Key Store, continue to the next step.
6. Distribute revocation.properties to the CA DataMinder \data folder on the CMS and all gateway servers. Then restart the CA DataMinder infrastructure service.
7. Finally, in the Administration console run the 'All machines with out-of-date Revocation List' custom search to confirm that the CMS and all CA DataMinder gateway servers and client machines have the new list.

Chapter 7: iConsole Deployment

This section describes how to deploy the iConsole and its supporting Web services.

Note: iConsole searches and reports are covered in the following chapter.

This section contains the following topics:

Overview

The CA DataMinder iConsole comprises the following components:

iConsole

The iConsole is a lightweight, browser-based application providing event searching and auditing features. It is primarily aimed at auditors and reviewers.

Front-**E**nd Web Server

iConsole users must direct their browser to the front-end web server. That is, they must direct their browsers to a URL based on the name or IP address of the machine hosting the front-end web server.

The front end submits all event searches generated in the iConsole to the application server (see below) and renders the matching events returned from the application server as HTML search results screens. The iConsole screens are generated in Javascript, based on data retrieved using AJAX calls.

Application Server

This component provides the web service that connects to the CMS. It enables all event search and auditing activity conducted in the iConsole to be written to the CMS. It enables iConsole users to search for and retrieve events stored on the CMS and to update audit details for these events.

iConsole Architecture

The iConsole front-end Web server and application server are provided as separate components to allow maximum flexibility when deploying the iConsole. For example, you may prefer to install the front-end Web server on your existing corporate Web server while installing the application server on an existing CA DataMinder server. For installation instructions.

Each front-end Web server can only connect to a single, specific application server, but it is possible to connect multiple front-end Web servers to a single application server. For example, this may be preferable if your iConsole users are based in various offices around the world; this configuration enables each user to connect to a local front-end Web server, with each front-end Web server connected to a central application server.

You can also have multiple application servers, each serving different front-end Web servers but all connected to the same CMS. Larger organizations may choose this configuration for load-sharing purposes.

Finally, each application server has a default parent CMS, but can be configured to connect to multiple CMSs if required.

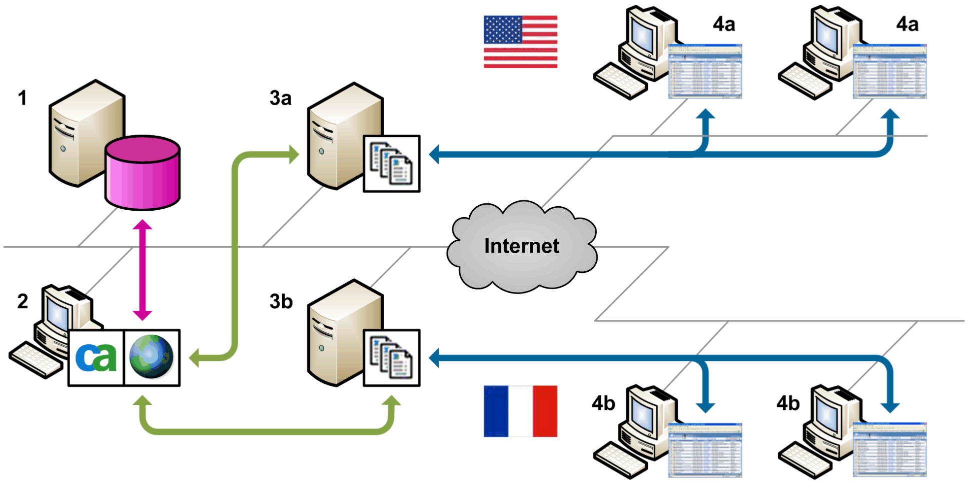

iConsole example deployment

1 CMS. This services all search requests submitted by the application server. All search SPs are stored in the CMS database; all XML search definition files are stored in the CMS file system.

2 iConsole application server. This submits iConsole event searches and audit updates to the CMS and returns search results to the front-end web server. Each application server is parented to a single CMS. If required, multiple application servers can connect to a single CMS.

3a, **3b iConsole front-end Web servers**. These generate the HTML content for the various screens in the iConsole. In this example, two front-end Web servers each serve separate groups of iConsole users (for example users based in New York and Paris), but connect to a single shared application server.

4a, 4b Browser-based iConsoles. Reviewers and administrators use the iConsole to search for and retrieve events stored on the CMS and to update audit details for these events. The iConsole URL incorporates the name or address of the front-end Web server host. In this example, the New York iConsoles (**4a**) connect to front-end Web server **3a**, while the Paris iConsoles (**4b**) connect to front-end Web server **3b**.

More information:

Deployment Procedure (see page 110)
Connect iConsoles to Multiple CMSs (see page 120)

iConsole Registry Values

The following sections describe the available registry values for the iConsole servers. If you need to configure the iConsole or change its default behavior, you can find these registry values in the following registry keys:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\Web
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\Web\Logging
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\WebService
```

This key contains the following registry values:

- AdditionalCMSList (see page 120)
- AddressBookSearchLDAPFilter (see page 145)
- AdvancedSecurity (see page 128)
- AllowConcurrentSearches (see page 141)
- AllowUnlimitedSearch (see page 133)
- AllowURLLogon (see page 124)
- AlwaysLogSearchCriteria (see page 150)
- ApplyPolicyToAuditEmails (see page 142)
- AuditEmailAddress (see page 146)
- AuditEmailAddressMask (see page 147)
- BulkAuditBatchSize (see page 133)
- CaptureEventReviewTime (see page 142)
- EnableAddressBookDropDownMenu (see page 145)
- EnforceEncryptedLogon (see page 126)
- EnforceLogonTimestamp (see page 127)
- EventCachePeriodMinutes (see page 130)
- EventDisplayMaxIMParticipants (see page 136)
- EventDisplayActiveIMParticipantsOnly (see page 136)
- EventDisplayMaxMailAddresses (see page 135)
- EventDisplayMaxSummaryParticipants (see page 135)
- EventMailDownloadFormat (see page 138)
- FriendlyNameLDAPAttribute (see page 150)
- LogLevel (see page 150)
- LogMaxNumFiles (see page 150)
- LogMaxSizeBytes (see page 150)
- LogonTimestampInterval (see page 127)
- LookupLDAPServers (see page 147)

- MaximumResultSetSize (see page 133)
- PasswordExcludeChar (see page 123)
- PolicyCachePeriodMinutes (see page 132)
- PreAuthenticate (see page 121)
- RequestPollPeriodMSec (see page 140)
- SearchParamsCachePeriodMinutes (see page 133)
- SearchResultsCachePeriodMinutes (see page 133)
- SearchResultsConverterTimeoutSecs (see page 132)
- SearchResultsDownloadConvertTo<MyFormat> (see page 139)
- SearchResultsDownloadFormat (see page 137)
- SearchResultsDownloadExt (see page 137)
- SearchResultsPageSize (see page 133)
- SearchResultsUpdateAuditStatus (see page 133)
- SessionTimeoutMinutes (see page 129)
- SessionTimeoutWarningSeconds (see page 129)
- SMTPClientTimeout (see page 146)
- SMTPServer (see page 143)
- WebServiceMachine (see page 123)
- WebServicePort (see page 123)
- WebServiceTimeoutSeconds (see page 131)
- WebServiceUseSSL (see page 123)

Deployment Procedure

Deploying the iConsole across your organization is a four-step procedure. First, before installing an iConsole application server or front-end Web server, you must ensure that the target machines have .NET Framework and Microsoft IIS correctly installed and configured (this allows the necessary Web services to run).

Next, you must install an iConsole application server and one or more front-end Web servers. To do this, you run the iConsole installation wizard on an existing CA DataMinder utility machine.

Then you need to perform certain post-installation tasks. In particular, you need to install the default event searches. You must also ensure that the iConsole servers are correctly configured to communicate with your chosen SMTP server. Other optional tasks include configuring a global sender for audit emails and renaming the iConsole virtual directory.

Finally, to permit your reviewers to start using the iConsole, you must provide them with the correct URL.

These steps are described in the following sections.

More information:

iConsole Requirements (see page 110)
Deploy the iConsole (see page 117)
Post-deployment Tasks (see page 118)

iConsole Requirements

Requirements for Application Servers and Front-End Web Servers

Before installing an iConsole application server or front-end web server, verify that the host server has the necessary software installed.

Briefly, the host server requires an appropriate operating system and Microsoft Outlook. It must be a CA DataMinder server. It needs appropriate versions of .NET Framework and Microsoft IIS. Finally, Kerberos must also be correctly configured.

Operating System

The host server has the following system requirements:

Supported Versions

Web.msi supports 32-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)

Web.msi also supports 64-bit versions of these operating systems:

- Windows Server 2008 R2
- Windows Server 2012

Note 1: We have not tested these operating systems with the current versions of web.msi.

Note 2: If running Windows Server 2008, see the special requirement for Microsoft IIS 6 WMI Compatibility.

More information:

Microsoft IIS (see page 113)

CA DataMinder Server

Applies to iConsole application servers only.

The host server must be a CA DataMinder server. This automatically connects the application server to a parent CMS. We recommend that you install iConsole application servers on CA DataMinder **utility** machines.

Note: iConsole front-end Web servers do **not** require any other CA DataMinder software on the host machine.

More information:

Utility Machines (see page 36)

iConsole Version Compatibility

iConsole application servers and iConsole front-end Web servers **must** be running the same version of CA DataMinder!

Important! We recommend that you upgrade all your iConsole servers at the same time. We do not support iConsoles where the application servers and front-end Web servers are running different versions of CA DataMinder.

Microsoft Outlook

(Applies to iConsole application servers only.)

> The host server must be running Microsoft Outlook and Outlook must be the default email application on the host server.

This ensures that CA DataMinder email events are correctly converted to MSG files when downloading e-mails in the iConsole Search Results screen or when attaching original messages to audit emails.

Note: iConsole front-end Web servers do *not* require any Microsoft Outlook on the host machine.

.NET Framework

The host server must be running .NET Framework 2.0.

To check the installed version of ASP.NET, run the CA DataMinder Version Check utility, wgncheck.exe. Search the output for the Microsoft .NET Framework section.

About .NET Framework

> .NET Framework provides various IIS Web Service extensions required by iConsole application servers and front-end Web servers. We recommend that .NET Framework is installed on the target server *before* Microsoft IIS; if this was not the case, or is not practical, you may need to manually register the necessary ASP.NET Web Service extensions—see 'Registering IIS extensions' below.

Special Requirement for Windows Server 2003 x64 Edition

> If you want to deploy an iConsole server on a machine running Windows Server 2003 x64 edition, you must first install the **32-bit version** of ASP.NET 2.0.
>
> Instructions for switching to the 32-bit ASP.NET 2.0 are available in Microsoft Knowledge Base article 894435. This article describes how to enable 32-bit mode and then install ASP.NET 2.0 (32-bit) and the script maps. You must also ensure that the status of ASP.NET 2.0 is set to ‘Allowed’ in the Web service extension list in IIS Manager.

More information:

Version Check Utility: Wgncheck.exe (see page 116)

Microsoft IIS

The host server requires Microsoft Internet Information Services (IIS). Note also the special requirements for IIS 7.x and 8.0 (see below).

To check the installed version of ASP.NET, run the CA DataMinder Version Check utility, wgncheck.exe. Search the output for the Microsoft IIS section. When you check the IIS version, you also need to confirm that the ASP.NET Web Service extension is registered with IIS. If it is not, you must register it manually (see below).

Registering IIS Extensions

iConsole application servers and front-end Web servers require IIS and various Web Service extensions provided with .NET Framework. We strongly recommend that you install or enable IIS *before* installing .NET Framework.

If this is not possible or practical, you can install .NET Framework before you install IIS, but you must then manually register the Web Service extensions before running the iConsole installation wizard:

To register the Web Service extensions

a. On the target server, go to the following folder:

   %windir%\Microsoft.NET\Framework\v2.0.50727

b. In this folder, run the following command:

```
aspnet_regiis /i
```

c. Restart the target server.

IIS must be running in 32-bit mode

iConsole servers cannot run if IIS is hosting other applications that require it to be in 64-bit mode.

IIS subcomponent 'SMTP Service' must be enabled

The IIS subcomponent, SMTP Service, must be running on the target server to allow the local machine to deliver SMTP emails, including audit emails sent by iConsole reviewers; see Set Up SMTP Email (see page 143).

Special Requirements for IIS 7.0, 7.5, and 8.0

(Applies to iConsole servers on Windows Server 2008, 2008 R2, and 2012)

If the host server is running IIS 7.0, 7.5, or 8.0, you must install various Web Server Role Services (see page 113).

To install the required Web Server Role Services

a. Log on to Windows Server.

b. Click Start and choose Server Manager.

c. Browse to Roles, Web Server (IIS) and click Add Role Services.

d. In the Select Role Services screen, ensure that the following role services are installed:

   Web Server, Management Tools

   Select all IIS 6 Management Compatibility items.

   Web Server, Application Development

   Select ASP.NET. All associated role services are installed automatically.

   Web Server, Common HTTP Features

   Select Static Content.

   Web Server, Security

   Select Windows Authentication.

More information:

Version Check Utility: Wgncheck.exe (see page 116)
Unable to Download or Forward Original .msg File (see page 603)

Kerberos Authentication

Applicable if the application server and front-end Web server are on separate machines.

The iConsole uses Microsoft's Kerberos Authentication to allow the credentials of the user accessing the iConsole to be passed to the CMS for logon (either for direct use if using CA DataMinder single sign-on functionality, or to record the native user name being used to access the CMS), using Windows Delegation. For this process to work if the iConsole front-end server and application server are on separate machines, you mst adhere to the following requirements:

1. The iConsole servers must be in the same Active Directory domain. If the value for WebServiceMachine is not the Fully Qualified Domain Name (FQDN), then the front-end machine must be trusted for Delegation. For details, see the Microsoft TechNet article ‘Allow a computer to be trusted for delegation’. The URL for this article is:

   http://www.microsoft.com/technet/prodtechnol /windowsserver2003/library/ServerHelp /b207ee9c-a055-43f7-b9be-20599b694a31.mspx

2. You must configure the Microsoft Internet Information Services (IIS) version 6 Application Pool to run as the Network Service account. This is the default configuration.

3. Kerberos must be correctly configured. Check the Windows System Event log for errors. See the next section for details.

4. Internet Explorer on the user's machine must have the Enable Integrated Windows Authentication (requires restart) setting enabled. This is the default setting in most configurations of Internet Explorer.

If you do not adhere to these requirements, this can result in the error ‘You are not authorized to connect to the CA DataMinder iConsole’, with a 401 error code.

Is Kerberos Active?

To check whether Kerberos is active on an iConsole server, run a netdom command:

Syntax

```
netdom verify /d:<domain> <server>
```

Example

```
netdom verify /d:unipraxis.com ux-hardy-as
```

Note: netdom is not installed by default, but is available from support.cab in the \Support\Tools folder on your Windows distribution media.

If Kerberos is active, this command generates a confirmation, such as:

```
The secure channel from UX-HARDY-AS to the domain UNIPRAXIS.COM has been verified.
The connection is with the machine \\UX-SRVR.UNIPRAXIS.COM. The command completed
successfully.
```

if Kerberos is not active, check for Kerberos entries in the Security event log in Windows Event Viewer. The most common local problem is timing; the server clock must be within five minutes of the domain controller clock. Other Kerberos problems typically affect the entire domain or require domain administrator permissions. For example, if Kerberos cannot authenticate a user because their account has become corrupt in Active Directory, the account must be reset on the domain controller.

More information:

About Single Sign-On (see page 164)

Requirements for Browser Host Computers

The iConsole is a browser-based application. These requirements apply to the browser host machine.

Browser

The iConsole runs in the following browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Internet Explorer 8, 9, or 10

Note: The iConsole may run in other browsers, but these have not been tested.

BusinessObjects InfoView

(Applicable only if CA DataMinder integration with BusinessObjects Intelligence is enabled.)

InfoView is the BusinessObjects web portal. You can launch InfoView from the iConsole or you can browse to InfoView directly.

In BusinessObjects XI 3.1 SP5, InfoView is supported in:

- Mozilla Firefox
- Microsoft Internet Explorer 8 and 9 (in Compatibility View only)

Microsoft Outlook

When browsing search results, if a reviewer wants to view a copy of an actual email (that is, if they want to open a downloaded .msg file), the browser host machine must be running Microsoft Outlook.

Note: If Outlook is not available, the reviewer can still save the downloaded .msg file.

Version Check Utility: Wgncheck.exe

Wgncheck.exe gathers diagnostic data, including CA DataMinder log files, that can be forwarded to CA technical support staff. The output is sent to a zip file on your desktop.

To run the utility, double-click Wgncheck.exe or run this command:

```
wgncheck
```

To see the full usage, run this command:

```
wgncheck -?
```

Find this utility in the \Support folder on the CA DataMinder distribution media.

Deploy the iConsole

This section describes the iConsole deployment procedure.

Before Installing

Before you launch the iConsole installation wizard:

- **IIS:** Note the IIS requirements.
- **Install on a utility machine:** If you plan to install an iConsole **application** server, the host machine must already have a CA DataMinder server installed. If this is not so, install a server now. We recommend installing the application server on a CA DataMinder **utility** machine.

More information:

CA DataMinder Servers (see page 35)

Install the iConsole

You install the iConsole servers using the CA DataMinder iConsole installation wizard.

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose iConsole Web Server and then click Install.

   This launches the CA DataMinder iConsole installation wizard in a separate window.

4. In the iConsole installation wizard, navigate to the Custom Setup screen.

5. In the Custom Setup screen, choose the Front-End Web Server and/or the Application Server.

   If you only install the front-end Web server, go to step 6.

   If you install both the application server and front-end Web server on the same machine, the front-end Web automatically connects to the local application server. Go to step 7.

6. In the iConsole Configuration screen, connect the front-end Web server to an application server:

   Server

   > Specify the name or IP address of the machine hosting the application server. Type 'localhost' to specify the local machine.

   Port

   > Specify the TCP port used for communication between front-end Web server and the application server. This defaults to port 80. If you specify a non-default port, you must also update the application server to use the same port.

   Use SSL

   > If you intend to use SSL to communicate over a secure port (for example, 443), select this check box. This ensures that the port used for communication between the front-end Web server and the application server will use SSL.

7. In the final wizard screen, click Install to start the file transfer.

   But note there are several optional post-installation tasks.

Post-deployment Tasks

After installing the iConsole, there are various optional post-deployment tasks. These are listed below.

More information:

Install iConsole Searches and Reports (see page 119)
Set Up the Review Queue (see page 119)
Set Up iConsole Connectivity (see page 119)
Hide User Logon Credentials in the iConsole (see page 124)
Protect Against ClickJacking Attacks (see page 128)
Set Up iConsole Timeouts (see page 128)
Set Up Search Results Handling (see page 132)
Set Up Event Auditing (see page 141)
Configure Audit Emails (see page 142)
General Setup Tasks (see page 150)

Install iConsole Searches and Reports

You install the iConsole standard searches, reports, and the Review Queue by running setup.exe. You must repeat the installation on your CMS, your iConsole application servers and your iConsole front-end Web servers. This is described in the following chapter.

Set Up the Review Queue

The Review Queue (RQ) feature enables reviewers to generate lists of events that they need to review or audit. Specifically, it provides an iConsole search that reviewers can run to retrieve all unreviewed events in their review queue (unreviewed events are assigned to reviewers based on their management groups). It also includes various reports that provide administrators with technical information about RQ database searches.

The required RQ database components are installed automatically when you install or upgrade your CMS. You install the RQ search and administrative reports when you install the iConsole standard searches and reports; see the following chapter.

However, before your reviewers can run the RQ search, your DBAs must populate the queue with events that need to be reviewed. For details of how to configure and manage the Review Queue, see the *iConsole Review Queue Configuration Guide*.

Set Up iConsole Connectivity

After installing the iConsole, you may need to perform the following tasks.

Rename IIS Virtual Directory for Front-end Web Server

(Optional) The iConsole installation wizard creates separate virtual directories for the front-end Web server and application server. Both are installed onto the default Web site for IIS (also called the home directory).

Front-end Web Server

The installation wizard creates this virtual directory:

```
CADATAMINDER
```

This virtual directory is incorporated into the target URL for iConsole users—see Start the iConsole. If required, you can rename this virtual directory in IIS. For example, you may want to rename it to "Compliance".

Application Server

The installation wizard creates this virtual directory:

```
CAWebService
```

Important! Do **not** rename this virtual directory!

Connect iConsoles to Multiple CMSs

To allow iConsole users to connect to multiple CMSs

1. On the application server host machine, go to the WebService registry key.
2. Within this key, modify this registry value:

   ```
   AdditionalCMSList
   ```

   Type: REG_SZ

   Data: Specifies a comma-separated list of CMSs that the local application server can connect to.

   Now, when users next use the iConsole they can browse to the Logon page and choose which CMS to connect to.

Enable Pre-authentication

This task is optional. Authentication between the front-end Web server and application server must be correctly configured for best performance, for example, when displaying individual events or paging between screens. If you experience poor performance and you suspect that this is due to slow authentication, you can modify the registry on the iConsole front-end Web server to use pre-authentication.

Note: This improvement does not apply to search performance, which is dependent on how the CMS database is configured.

After making this registry change, you will also need to implement a Microsoft workaround to accommodate security fixes introduced for Windows XP SP2 and Windows Server 2003 SP1.

1. First, configure the front-end Web server to use pre-authentication. This ensures that the logon credentials are always passed to the application server, rather than using 'anonymous access'.
   a. Locate the Web registry key on the front-end Web server.
   b. Within this registry key, set the following value to True:

      ```
      PreAuthenticate
      ```

2. Changing this registry value can lead to instability (for example, HTTP 401.1 errors). To prevent this instability, you must now implement a Microsoft workaround. The required workaround is described in MS Knowledge Base article Q896861. This article describes two alternative workarounds: disabling the loopback check; and specifying host names. We recommend that you implement the first method and disable the loopback check:
   a. As described in article Q896861, locate the following registry key:

      ```
      HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa
      ```

   b. In this registry key, create the following registry value and set it to a DWORD value of 1

      ```
      DisableLoopbackCheck
      ```

3. Restart the machine hosting the front-end Web server.

More information:

Set Up iConsole Timeouts (see page 128)

Enable Anonymous Access

If your iConsole application servers and front-end Web servers are deployed in an environment where they cannot communicate using Microsoft Windows authentication, you must enable anonymous access for both of the iConsole virtual directories.

For example, you may need to do this if the application server and front-end Web server are in separate domains and where, for security reasons, the front-end Web server's domain is not trusted by the application server's domain.

To enable anonymous access, you need to configure the authentication methods for both of the iConsole virtual directories. By default, these are named CA and CAWebService.

To enable anonymous access

1. In IIS, edit the properties of the application server virtual directory, CAWebService.
    a. In the Properties dialog, go to the Directory Security tab and click the Edit button under Anonymous Access and Authentication Control (IIS 5.x) or Authentication and Access Control (IIS 6.x) to open the Authentication Methods dialog.
    b. In the Authentication Methods dialog, enable anonymous access **only**. This method ensures that when the front-end Web server accesses the CAWebService virtual directory, no authentication credentials are required.
2. Still in IIS, now edit the properties of the front-end Web server virtual directory (its default name is cadataminder).
    a. As before, in the Properties dialog, go to the Directory Security tab and click the Edit button under Anonymous Access and Authentication Control (IIS 5.x) or Authentication and Access Control (IIS 6.x) to open the Authentication Methods dialog.
    b. In the authentication dialog, enable anonymous access **only**.

Specify Disallowed Characters in Logon Passwords

The iConsole provides CA DataMinder users with the ability to change their logon password. You can configure the iConsole to disallow specific characters in these passwords. To do this:

1. On the iConsole front-end Web server, locate the Web registry key on the front-end Web server.
2. Within this registry key, create or modify the following value:

   PasswordExcludeChar

   Type: REG_SZ

   Data: Defaults to empty. Specifies which characters are not allowed as part of a password. For example, to disallow:

   ```
   < > ( ) " ' % ; & + - =
   ```

   Set this registry value to:

   ```
   <>()"'%;&+-=
   ```

   The new password rules become effective for all subsequent logon sessions.

   Note: We recommend that you specify this list globally using the machine policy setting ‘Prohibit password characters’ on the CMS.

Specify a Non-default TCP Port

When installing the iConsole application server and front-end Web server separately using the installation wizard, you must specify the **TCP port** used for communication between them. This defaults to port 80, but if required you can specify a non-default port at installation.

This section provides instructions for changing the application server or TCP port *after* installation.

To edit the registry on the front-end Web server

1. On the host machine for the front-end Web server, locate the Web registry key.
2. In this registry key, you need to check the following registry value:

   `WebServiceMachine`

   This REG_SZ value specifies the name of the machine hosting the application server.

   - If this value is set to either LocalHost, or the name of the machine hosting the front-end Web server, the iConsole assumes that both components are hosted on the same machine.
   - If this value is not set to Localhost or the host machine for the front-end Web server, you need to specify the name of the machine hosting the application server.

3. In the same registry key located in step 1, you need to configure the following registry value:

   `WebServicePort`

   This REG_SZ value defaults to 80. It specifies the TCP port used for communication between the front-end Web server and the application server.

   a. Change this value to the TCP port number you want to use.

   b. If you do change the TCP port, you must ensure that the application server is communicating on the same port. This can be checked using Internet Information Services (IIS).

   c. If you are using SSL to communicate over a secure port (for example, 443), you need to configure the following registry value:

      `WebServiceUseSSL`

      This REG_DWORD registry value defaults to zero. Set this to 1 if you want the port used for communication between the front-end Web server and the application server to use SSL.

Hide User Logon Credentials in the iConsole

If integrating the iConsole with your own Web application, there are two supported methods for automatically passing user credentials to the CMS, so avoiding the need for users to manually log on. These methods are: Web form logon, which uses HTTP POST, and URL query string logon, which uses HTTP GET.

Note: Both methods are supported by SSL connections.

To enable Web form or URL query string logon

1. Locate this Web key on the front-end Web server.
2. Within this registry key, set the following REG_DWORD registry value to 1: `AllowURLLogon`

Note: For full implementation details, please contact CA Support at http://ca.com/support.

More information:

Web Form Logon Method (see page 125)
URL Query String Logon Method (see page 126)
Enforce Encrypted Logons (see page 126)
Enforcing a Logon Timestamp and Timeout (see page 127)

Web Form Logon Method

Using this HTTP POST method, your Web application passes the user's CA DataMinder account credentials to an HTML form. The form then submits these credentials to the iConsole front-end Web server. This ensures that user credentials are not exposed in the browser Address bar.

Example HTML syntax is shown below. The example assumes your CMS is not using single sign-on.

```
<form id="friConsole" method="post" runat="server">
  <!-- If your CMS is using single sign-on, you can use "http" -->
    <action="https://localhost/cadataminder/default.aspx">
  <!-- If your CMS is using single sign-on, you can omit the "username"
  and "password" fields -->
    <input type="hidden" name="formAction" value="login" />
    <input type="hidden" name="username" value="unipraxis\srimmel" />
    <input type="hidden" name="password" value="password" />
    <input type="submit" value="Logon" />
</form>
```

The formAction name is case-sensitive.

This example also assumes that:

- The iConsole front-end Web server is running on the local machine.
- The virtual directory for the iConsole front-end Web server is ‘cadataminder’.
- Your CMS is **not** using single sign-on, so user credentials are submitted using HTTPS.

If your CMS **is** using single sign-on, you must:

- Include the formAction and submit input fields. You do not need the username and password fields.
- Set the action to use HTTP when submitting the form if no user credentials are included.

More information:

About Single Sign-On (see page 164)

URL Query String Logon Method

This logon method uses http get and passes a user's CA DataMinder account credentials to the iConsole in the form of a URL query string.

However, this logon method is only suitable for iConsole deployments that hide the browser Address bar. This is because the user's name and password are appended to the iConsole URL in the Address bar and are therefore potentially visible to other users.

Enforce Encrypted Logons

In addition to SSL support, CA DataMinder also enables you to encrypt HTML POST form variables using Triple DES (Data Encryption Standard) or AES (Advanced Encryption Standard).

Use this additional encryption where SSL sessions can potentially be intercepted.
For example, use this encryption to prevent ‘man-in-the-middle’ attacks where users are working remotely and creating traffic across the internet.

To enforce encrypted logons, edit this registry value in the Web registry key on the front-end Web server:

EnforceEncryptedLogon

Type: REG_DWORD

Data: Set this value to 1 to enforce encrypte logons.

Important! If you enforce logon encryption, the POST form variable supplied with the Web form logon method must be encrypted. If the POST form variable is not encrypted, the logon will fail.

Enforcing a Logon Timestamp and Timeout

For added security, you can set up a timestamp variable, which determines if a logon request occurs within a configurable time period. This can help to prevent 'replay attacks'.

To use this optional timestamp, the time of the request (in UTC time) must be submitted via the ‘timestamp’ POST variable (embedded inside the encrypted ‘agentID’ variable), formatted in US date format (that is, ‘mm/dd/yyyy hh:mm:ss’).

To enable and configure the allowed range of timestamps, edit these registry values on the front-end Web server:

EnforceLogonTimestamp

> **Type:** REG_DWORD
>
> **Data:** Defaults to 0. Set to 1 to enforce the timestamp in the POST form variables.
>
> **Important!** If it is enabled, then the POST variable supplied with the Web form logon method may include a timestamp which is checked to see if it is within a specified time range. If it is not within this range, then the logon fails.

LogonTimestampInterval

> **Type:** REG_DWORD
>
> **Data:** Defaults to 5 (minutes). Specifies a valid time interval. If set to 5, the logon timestamp must be within 5 minutes either side of the current time. For example, if the timestamp is 12:00, then the timestamp is valid from 11:55 to 12:05.

Note: For full details on implementing these additional security measures, please contact CA Support at http://ca.com/support.

Protect Against ClickJacking Attacks

You can protect the iConsole from frame-based clickjacking (or UI redressing) attacks. Use the AdvancedSecurity registry value to add the X-Frame-Options HTTP header to the iConsole. The header is automatically set to 'DENY'. This prevents the iConsole from being hosted in an iFrame.

To enable protection against clickjacking, modify the AdvancedSecurity value in the Web registry key on your front-end Web servers:

AdvancedSecurity

Type: REG_DWORD

Data: Defaults to True. Set this value to 1 or True to add an X-Frame-Options header to the iConsole web page and prevent the iConsole from being hosted in an iFrame.

If you set this value to 0 or False, an X-Frame-Options header is not added to the iConsole web page. This allows the iConsole to be hosted in an iFrame.

Note: In a frame-based clickjacking attack on the iConsole, an attacker hosts the iConsole in an iFrame. Unsuspecting users are unaware of the danger and the attacker is able to intercept any information submitted to the iConsole, such as user credentials or passwords.

Set Up iConsole Timeouts

The section describes how to modify the default timeouts for the iConsole.

Modify the Session Timeout

By default, if the iConsole detects no user activity (such as running a search or auditing an event) for 20 minutes, it displays a warning that the current session is about to expire. Clicking Cancel enables you to reset the session timeout for a further 20 minutes. The warning message itself will timeout after 20 seconds if no action is taken and the user is redirected to the Reconnect screen.

The automatic session timeout ensures that disconnections are handled efficiently and ensures that hung sessions do not remain on the iConsole application server, where they can consume system resources.

For example, a session will fail to terminate correctly if a user quits the iConsole by clicking the browser Close button instead of clicking the Logoff button in the iConsole. If this happens, the session will persist on the application server. The automatic timeout ensures that the residual session left after clicking the Close button is eventually terminated when the timeout expires.

Note: The iConsole session timeout overrides any specified Microsoft IIS timeout.

To lengthen or shorten the automatic session timeout, or to adjust how long the session timeout warning is displayed, you need to modify values in the Web registry key on the front-end Web server. Within this registry key, modify the following values:

SessionTimeoutWarningSeconds

Type: REG_DWORD

Data: Defaults to 20. Specifies how long (in seconds) the session timeout warning is displayed before the user is redirected to the Reconnect screen.

SessionTimeoutMinutes

Type: REG_DWORD

Data: Defaults to 20. Specifies the session timeout (in minutes) for the iConsole. The timeout countdown begins as soon as the user logs on to a CMS and restarts each time CA DataMinder detects user activity. If the timeout expires and no user activity was detected, CA DataMinder terminates the current session and displays the iConsole Logon screen.

Modify the Event Timeout

When an event is viewed in the iConsole, its details are cached by the iConsole application server to enable faster viewing of the same event in future. This is especially noticeable when viewing large events.

If multiple users perform large searches and then view multiple events, these lists of cached events can potentially use up a large amount of resources. You can therefore configure how long these events are retained in the cache list, after which they are released.

To configure this timeout, you need to modify a value in the WebService registry key on the iConsole application server. Within this registry key, modify the following value:

EventCachePeriodMinutes

Type: REG_DWORD

Data: Defaults to 5. Specifies the cache timeout (in minutes) for viewed events in the iConsole. The timeout begins when the event is added to the cache. When the timeout expires, that event is removed from the cache.

Modify the Web Service Timeout

Note: It is highly unlikely that you need to modify this registry value.

Each time an individual event is viewed in the iConsole, the front-end web server calls the web service on the application server. By default, this call times out after 100 seconds. That is, the front-end web server has 100 seconds to download and display the event in the Search Results screen.

But for extremely large events (whose size is measured in tens of megabytes), or where retrieval of the stored event from a third-party archive is slow, this timeout may expire before the event has fully downloaded. In these exceptional conditions, increase the web service timeout.

To increase the timeout for large files

1. Modify a value in the web registry key on the iConsole front-end web server:

   WebServiceTimeoutSeconds

   Type: REG_DWORD

   Data: Specifies the timeout (in seconds) for the web service running on the application server. Defaults to 100 seconds. Increase this timeout if extremely large events are failing to display in the iConsole Search Results screen.

2. If you do increase the timeout, also increase the iConsole request timeout:
   1. Browse to the CA DataMinder installation folder.
   2. Edit the web.config file on the front-end web server.
   3. Find the httpRuntime executionTimeout attribute and set this timeout to be approximately 10 seconds longer than the WebServiceTimeoutSeconds timeout.

Modify the Results Conversion Timeout

By default, the iConsole displays reports in XML Spreadsheet format, but some reports output results in other formats such as PDF. Such reports typically require a post-processing application (for example, Apache FOP) to convert the output. To safeguard against the conversion process failing or taking too long, you can specify the maximum time allowed for the conversion.

If this timeout expires before the conversion is complete, the report is effectively canceled. An on-screen warning notifies the reviewer that the report has failed and a corresponding entry is written to the iConsole log file (see Configure iConsole log files).

To configure the results conversion timeout, you need to modify a value in the WebService registry key on the iConsole application server. Within this registry key, modify the following value:

SearchResultsConverterTimeoutSecs

Type: REG_DWORD

Data: Defaults to 300. Specifies the maximum time allowed (in seconds) for the results conversion. The timeout begins when the conversion process starts; it does not apply to the time taken to retrieve the report results from the database. When the timeout expires, the report is canceled.

Modify the Policy Cache Period

Important: Only change this registry setting if instructed to do so by CA technical staff.

If you often open and close the same large policy, it may take a long time to load. Use this setting to increase the period how long the policy remains cached (at the expense of memory used).

You find the following key in the WebService key:

PolicyCachePeriodMinutes

Type: DWORD

Default: 10 minutes

Specifies the policy cache timeout. The policy document is automatically cached the first time a number of settings are requested. Only the document is cached, individual policy settings are not cached. The policy cache is unloaded when the user's editing session is finished, and after the specified time period after the user closes the browser.

Set Up Search Results Handling

When you run a search and view its results, the iConsole displays all matching events in the Search Results screen. You can configure how the iConsole handles these results.

Configure the Search Results Cache

When you run a search, the iConsole caches the results to support faster paging through the results. You can configure how the iConsole handles the search results cache.

To do this, you need to modify values in the WebService registry key on the iConsole application server. Within this registry key, modify the following values:

AllowUnlimitedSearch

Type: REG_SZ

Data: Defaults to False. Determines whether reviewers are permitted to run unlimited (or ‘uncapped’) event searches. That is, the iConsole will return all events that match the search criteria, even if the total exceeds the limit specified by the MaximumResultSetSize value (see below).

To enable unlimited event searches, set this registry value to True.

Note: This registry value simply configures the iConsole application server to support unlimited searches; individual reviewers require the ‘Events: Allow searches of unlimited size’ privilege before they run these searches. See the Administration console online help for details; search for ‘privileges’.

MaximumResultSetSize

Type: REG_DWORD

Data: Defaults to 1,000. Specifies the maximum number of results that can be returned by a search for events. Because search results are cached in memory on the iConsole application server, setting a maximum result limit prevents the cache consuming excessive memory and adversely affecting system performance.

SearchParamsCachePeriodMinutes

Type: REG_DWORD

Data: Defaults to 5. Specifies how long search parameters (obtained through database stored procedures) are cached for. This enables the iConsole to display the Customize Search screen without making frequent calls to the CA DataMinder database.

SearchResultsCachePeriodMinutes

Type: REG_DWORD

Data: Defaults to 40. Specifies the result retention period (in minutes) for the event cache on the iConsole application server. Events are flushed from the cache if it is not accessed before this period elapses. The cache period countdown begins as soon as the search results are displayed and restarts each time the iConsole detects a ‘cache access’, such as the user browsing to a different results page or selecting an individual event.

For example, the retention period is set to 30 minutes. If a user runs a search but then does not use the iConsole for 30 minutes, the search results are flushed. If the user subsequently wants to view the results again, they will need to rerun the search.

Next, you modify values in the Web registry key on the iConsole application server. Within this registry key, modify the following values:

BulkAuditBatchSize

Type: REG_DWORD

Data: Defaults to 50. Specifies the number of results that can be processed as part of a bulk review. That is, how many events you can select and then use one of the audit buttons.

You may want to adjust this setting if you are displaying a large number of events on the same page. In this case, also set the registry value SearchResultsPageSize to, for example, 100 or more—see above.

SearchResultsUpdateAuditStatus

Type: REG_SZ

Data: Defaults to True. Determines whether the iConsole updates the Audit Status column in the list of search results after an event has been reviewed. Setting this to False can speed up switching between events.

Configure How Many Participants Are Displayed

When you view an email or IM event in the Search Results screen, the iConsole lists all the ‘participants’ (email recipients and senders or IM participants) in the Mail and Information tabs. If the event has a large number of participants, the process of retrieving them can take a long time and sometimes causes timeout problems.

You can configure how many participants are displayed in these tabs by configuring values in the WebService registry key on the iConsole application server.

Configure the Number of Event Participants

You can configure how many participants the iConsole displays in the Information tab of the Search Results screen. To do this, modify the following registry value in the WebService registry key:

EventDisplayMaxSummaryParticipants

Type: REG_DWORD

Data: Defaults to 0. Specifies the maximum number of participants that the iConsole will process in order to display them in the Information tab of the Search Results screen. If this registry value is set to zero or does not exist, the list of participants is not capped and all are displayed.

Configure the Number of Email Recipients

You can configure how many recipients the iConsole displays in the Mail tab of the Search Results screen. To do this, modify the following value in the WebService registry key:

EventDisplayMaxMailAddresses

Type: REG_DWORD

Data: Defaults to 1000. Specifies the maximum number of email addresses that the iConsole will process in order to display email event recipients in the Mail tab of the Search Results screen.

If this number is exceeded, the list of e-mail addresses in the Mail tab is truncated. The full list of recipients for outgoing emails is available in the Summary section of the Information tab. For incoming emails, CA DataMinder only stores the details for a single sender and recipient.

Configure the Number of IM Participants

You can configure how many participants the iConsole displays in the IM Message tab of the Search Results screen. To do this, modify the following values in the WebService registry key:

EventDisplayMaxIMParticipants

Type: REG_DWORD

Data: Defaults to 0. Specifies the maximum number of participants from the original instant message to display in the IM Message tab of the Search Results screen. If this registry value is set to zero or does not exist, then all participants are displayed.

Note: Any participant names that have not been retrieved due to capping but appear in the message itself are displayed as <unknown>.

EventDisplayIMActiveParticipantsOnly

Type: REG_DWORD

Data: Defaults to 0. Set this to 1 to specify that only the active participants of the original instant message are listed in the Message tab of the Search Results screen. That is, the participants that appear in the message lines of the IM event. If this registry value is set to zero or does not exist, then all participants are displayed.

Note: If EventDisplayMaxIMParticipants is set to a lower value than the actual number of active participants, then it overrides the value of EventDisplayIMActiveParticipantsOnly and only a subset of the active participants is displayed.

Specify the Default Format for Downloaded Search Results

Note: For full details, see the *iConsole Search Definition Guide*; search the index for ‘downloads, custom file formats.’ This guide is available to download from CA Technical Support (see page 21).

You can download the entire set of search results to formats such as XLS and CSV. By default, the ‘Download all results’ button downloads results to XLS.

To set this to a different download format, you need to configure the Web registry key on the iConsole application server. Within this registry key, add these registry values:

SearchResultsDownloadFormat

Type: REG_SZ

Data: Specifies the file type for downloaded search results. CA DataMinder provides built-in support for xls and csv format. For example, set this value to csv to download search results to a comma-separated value file, or xls to download to an Excel spreadsheet.

SearchResultsDownloadExt

Type: REG_SZ

Data: Specifies the file extension the iConsole will use when downloading search results. This registry value then ensures that the downloaded file is associated with correct application. For example, set this value to xml to download search results to a spreadsheet file compatible with Excel 2007.

More information:

Specify the Format for Downloaded Emails (see page 138)

Specify the Format for Downloaded Emails

You can specify the supported file formats for emails downloaded via the iConsole Search Results screen. E-mail events from the following sources are automatically downloaded in .MSG file format:

- Microsoft Outlook
- Exchange Server
- Exchange Journal mailbox
- Symantec Enterprise Vault

If the email is from a different source, then the iConsole checks the event for a ‘MAPI-enabled’ flag. If this flag does not exist, or cannot be read, then the iConsole checks a registry value on the front-end Web server to see which file format to use.

To configure this registry value, you need to add it to the Web registry key on the front-end Web server. Within this registry key, add the following value:

EventMailDownloadFormat

Type: REG_SZ

Data: Specifies the default file format for any email downloaded using the iConsole which is *not* automatically downloaded in .MSG format. Set this to MSG, ZIP, or TXT.

If this setting does not exist or is left blank, and the email does not satisfy the criteria listed above, then the email can only be downloaded in .MSG file format if the mail body is in .RTF format. Otherwise, it will be downloaded as a .ZIP file containing an RTF representation of the email message, plus any attachments.

Note: To enable .MSG email download, the iConsole host server must be running Microsoft Outlook.

Setting the Default Format for Displaying Search Results

By default, the iConsole displays search results in XML Spreadsheet format. These can then be exported into Microsoft Excel.

For a more detailed description of supported functionality, see the *iConsole Search Definition Guide*; search the index for ‘results, downloading’.

Custom File Formats

CA DataMinder provides built-in support for XLS and CSV download formats (see above), but you can also implement custom formats using Extensible Stylesheet Language Transformations (XSLT).

Note: For full details, see the *iConsole Search Definition Guide*.

To configure the search definition file and the custom format stylesheet

1. In the custom format stylesheet results-formats-custom.xsl, you need to:
   - Define the custom format you want using an xsl:choose tag.
   - Write a transformation to render the search results to the custom format. Writing such transformations is beyond the scope of this Deployment guide.

   Note: This file is located on the iConsole application server in the \Web\transformations folder. If the file does not already exist, you will need to create it.

2. In the search definition file, reference the custom format using the format attribute in the <results> element.

   ```
   <results ... format="$doc:results.txt">
   ```

Post-processing Downloaded Search Results

To configure the iConsole to download search results to non-binary formats such as PDF, you need to define a custom download format that uses a post-processing step to make the conversion.

This post-processor is specified in the WebService registry key on the iConsole application server. In this key, you must create the following registry value:

SearchResultsDownloadConvertTo<MyFormat>

Type: REG_SZ

Name: When you create this registy value, replace <MyFormat> with the download format created by the post-processor. For example, to convert downloaded search results to PDF format:

```
SearchResultsDownloadConvertToPDF
```

Data: Specifies the full path and file name of the post-processor. This must be a command line application, accessible on the iConsole application server, that takes the names of an input file and output file as parameters and performs the conversion from one to the other. For example, set the data to:

```
C\Myfolder\MyConvertor.bat.
```

Creating an Additional Download Button on the Toolbar

You can configure the iConsole to add a toolbar button to the Search Results screen that enables reviewers to download search results to a different format. To do this, you need to configure the <tool> element and DownloadResults function in the iConsole search definition. See the example below:

```
<tool name="TXT_download"
  tooltip="Download file in TXT format"
  icon="download-TXT.gif"
  function="DownloadTXT"/>
<script>
  <![CDATA[
    <script language="javascript">
    <!--
      function DownloadTXT()
      {
        DownloadResults('txt');
      }
    -->
    </script>
  ]]>
</script>
…
```

Set the Request Poll Period

Important: Only change this registry setting if instructed to do so by CA technical staff.

Increase this value to reduce network traffic when you access the iConsole over a slow connection.

You find the following key in the Web registry:

RequestPollPeriodMSec

Type: DWORD

Default: 500 milliseconds

Defines the period (in milliseconds) that is used for polling AJAX requests in the iConsole, for example, when running a search.

Disallow Concurrent Searches

Important: Only change this registry setting if instructed to do so by CA technical staff.

Set the AllowConcurrentSearches key to false when you run the iConsole on a single-processor server to prevent concurrent search threads from thrashing.

You find the following key in the WebService key:

AllowConcurrentSearches

Type: REG_SZ

Default: true

Specifies whether concurrent searches are allowed.

Set Up Event Auditing

After installing the iConsole, you may need to perform the following tasks to enable event auditing in the iConsole.

Set Up the Auditing Feature

By default, CA DataMinder does not enable the auditing feature in the iConsole.

To enable your reviewers to audit events using the iConsole

1. In the Administration console, choose Tools, Options. The following message is displayed:

   ‘The audit strings have not been configured yet. Would you like to create the default settings?

   Note: If you select 'No', the audit tab will not be available on this options dialog, and you will not be able to perform auditing operations.’

2. Select Yes to enable the audit functionality in the iConsole and its subsequent configuration in the Administration console.

Note: For details on configuring the iConsole audit strings, see the Administration Console online help; search the index for ‘iConsole, configuring event audit features’.

Enable Time Recording for Reviewed Events

You can optionally configure the iConsole to record how long reviewers spend viewing or auditing individual events. The associated metrics (including the average review time per event for each reviewer) can be optionally included in ‘View Time’ columns in the results screen of the Compliance Audit Report; for details, see the online help for the report’s Customize and Results screens.

To enable time recording, you need to edit a registry value on the machine hosting the iConsole front-end Web server. On the host machine, locate the Web registry key. In this registry key, you need to edit this value:

CaptureEventReviewTime

Type: REG_DWORD

Data: Defaults to 0. Specifies whether to record the time spent reviewing events. By default, time recording is *not* enabled. Set this registry value to 1 to enable time recording.

Configure Audit Emails

From the iConsole reviewers can send emails to colleagues, alerting them to messages or issues that require their attention. These emails are referred to as ‘audit emails’. When a colleague receives an audit e-mail, the From: field indicates the sender.

Does CA DataMinder Apply Policy to Audit Emails?

By default, CA DataMinder excludes audit emails from policy. That is, these emails are ignored by the policy engine. To change this behavior, you need to edit a registry value on the machine hosting the iConsole front-end Web server. On the host machine, locate the Web registry key. In this registry key, you need to edit this value:

ApplyPolicyToAuditEmails

Type: REG_DWORD

Data: Defaults to 0. Specifies whether policy is applied to audit email. By default, policy is *not* applied. Set this registry value to 1 to ensure that policy *is* applied to audit emails.

How Does the iConsole Set the From: Field in Audit Emails?

The sender identity is generated automatically, unless you have defined a global sender. If the global sender has not been defined, the iConsole retrieves the SMTP addresses associated with the currently logged-on CA DataMinder user account.

- If AuditEmailAddressMask *is* configured: the iConsole uses the first valid SMTP e-mail address that matches the address mask.
- If AuditEmailAddressMask is *not* configured: the iConsole uses the first valid SMTP email address.

If there are no valid email addresses, or none match the mask, then the iConsole will attempt an LDAP lookup (see the next section) on the Windows account of the logged-on user:

- If this lookup succeeds, it then extracts the display name configured for this Windows user in the LDAP directory and writes this in the From: field (for example, 'Lynda Steel').
- If this lookup fails (that is, there is still no match), it discards any domain ('UNIPRAXIS\') and simply writes the remaining name directly into the From: field. For example, a reviewer logs onto the CMS using a bespoke CA DataMinder account, 'Unipraxis Compliance Officer'. This account has no corresponding LDAP entry, so the From: field is simply set to 'Unipraxis Compliance Officer'.

  Note: This lookup is only possible if Single-Sign-on is enabled on the CMS, or the user logged on manually with a valid Windows account.

More information:

Specify an Address Mask for Audit Emails (see page 147)

Set Up SMTP Email

To allow iConsole users to send audit e-mails, you must configure the front-end Web server so it can connect to an SMTP server (that is, a machine that can deliver SMTP emails).

More information:

Does CA DataMinder Apply Policy to Audit Emails? (see page 142)

Configure the SMTPServer Registry Value

If the SMTP service is not being used on the local server, you must edit the registry on the front-end Web server to point to a remote SMTP server. To do this, locate the Web registry key. Within this registry key, edit the following value:

SMTPServer

> **Type:** REG_SZ
>
> **Data:** This defaults to localhost. If the SMTP service is not being used on the local server, set this value to a remote SMTP server, that is, the DNS name of the server hosting the SMTP service. For example, you can set this value to your Exchange server (if it is configured to relay SMTP messages).
>
> **Note:** The SMTP server must be configured for pass-through authentication.

Enable the SMTP Service

The SMTP Service allows the local machine to deliver SMTP emails. Typically, the front-end Web server points to an existing, remote SMTP server, but if required you can enable the SMTP service locally.

1. On the host machine for the front-end Web server, ensure that you are logged on with local administrator rights.
2. Open the Windows Components Wizard. This is available from the Add or Remove Programs applet.
3. Select Internet Information Services (IIS) and click Details to display the IIS subcomponents.
4. Select the SMTP Service subcomponent and click OK to return to the Windows Components Wizard. This service supports the transfer of emails.

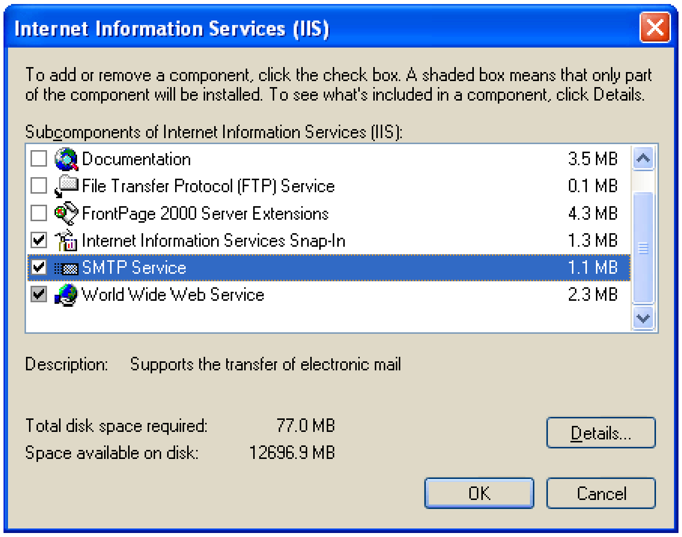

To add or remove a component, click the check box. A shaded box means that only part of the component will be installed. To see what's included in a component, click Details.

Subcomponents of Internet Information Services (IIS):

☐	Documentation	3.5 MB
☐	File Transfer Protocol (FTP) Service	0.1 MB
☐	FrontPage 2000 Server Extensions	4.3 MB
☑	Internet Information Services Snap-In	1.3 MB
☑	SMTP Service	1.1 MB
☑	World Wide Web Service	2.3 MB

Description: Supports the transfer of electronic mail

Total disk space required: 77.0 MB

Space available on disk: 12696.9 MB

Windows Components Wizard: IIS subcomponents

5. Click Next to build the new configuration settings, then click Finish to close the wizard.

Configure Address Book Lookups

When composing audit emails in the iConsole Compose Mail dialog, reviewers can fill in the To, Cc or Bcc address fields by searching for matching entries in an address book such as Active Directory.

Registry Changes

Address book searching is automatically enabled when you install a front-end Web server. But you can configure this by editing values in the Web registry key. Within this registry key, edit the following values:

EnableAddressBookDropDownMenu

Type: REG_DWORD

Data: This defaults to 1. Set this value to 1 to enable autocomplete in the To, Cc or Bcc fields of the iConsole Compose Mail dialog. That is, as a reviewer types the first letters of a recipient’s name, the iConsole automatically shows a list of matching entries in the address book.

If set to zero, autocomplete is disabled. But reviewers can still search for matching address book entries using the iConsole Address Selector dialog—see the screenshot below .

AddressBookSearchLDAPFilter

Type: REG_SZ

Data: This value defines the filter used to query the LDAP directory for matching address book entries. The default query is:

```
(&(|(givenname={0})
(samAccountName={0}))
(|(objectCategory=person)
(objectCategory=group)))
```

You do not normally need to change this value. You only need to do so if, for example, you only want to return users, but not groups, from the address book.

Note: The LDAP directory is specified by registry value LookupLDAPServers.

More information:

Select a Specific LDAP Directory (see page 147)

Modify the Audit Email SMTP Timeout

Important: Only change this registry setting if instructed to do so by CA technical staff.

You find the following key in the Web registry:

SMTPClientTimeout

Type: DWORD

Default: 20 seconds

Specifies the timeout limit (in seconds) when sending iConsole audit email over SMTP. A change to this value is picked up only after you restart IIS.

Specify a Global Sender

You can define a 'global sender' for iConsole audit emails. That is, you can configure the iConsole so that the From: field in audit emails is always set to the same sender, such as 'compliance@unipraxis.com'. This is useful if you want to anonymize audit emails (that is, you want to conceal the reviewer's identity).

To set up a global sender, you need to add a value to the Web registry key on the front-end Web server. In this registry key, add the following value:

AuditEmailAddress

Type: REG_SZ

Data: Specifies a text string representing the global sender. This text appears in the From: box in all audit emails.

This registry value overrides the default sender and also the AuditEmailAddressMask registry value. The text string must be in the form of a valid email address and must not contain spaces.

More information:

Specify an Address Mask for Audit Emails (see page 147)

Select a Specific LDAP Directory

Note: In the current release, the iConsole only supports lookup operations for Active Directory.

If the iConsole front-end Web server is running in a domain under Windows 2003 or XP, the front-end Web server will automatically detect an Active Directory server. However, you can configure the front-end Web server to use a specific, non-default LDAP directory.

To do this, you need to edit a registry value on the machine hosting the iConsole front-end Web server. On the host machine, locate the Web registry key. Within this registry key, you need to edit this value:

LookupLDAPServers

Type: REG_MULTI_SZ

Data: CA DataMinder policy engines permit you to define multiple LDAP servers for this registry value. Specify a list of servers hosting the LDAP directory you want to use.

Server names can be ‘plain’ or include a domain suffix (UNI-EXCH or UNI-EXCH.UNIPRAXIS.COM). If the LDAP port number is not 389, you can add it after the server name; prefix the port number with a colon (UNI-EXCH:319). You can also prefix the server name with the account credentials used to access the LDAP database. The syntax is:

```
<username>:<password>@<server name>
```

Specify an Address Mask for Audit Emails

CA DataMinder users may have multiple email addresses associated with their user accounts. Consequently, when sending an audit email or notification message, the iConsole may need to choose the most appropriate addresses when populating the From: and To: fields.

Specifically, this situation arises if audit emails do **not** use a global sender or if the recipients were chosen using the iConsole Address Selector dialog (that is, the reviewer sent the email to members of a specific user group or dynamic address list).

In this situation, you can define an email address mask (that is, an address pattern such as @unipraxis.co*) to enable the iConsole to select the most appropriate address. To do this, you need to add a value to the Web registry key on the front-end Web server. In this registry key, add the following value:

AuditEmailAddressMask

Type: REG_SZ

Data: Specifies a text string that represents an SMTP address pattern. The iConsole populates the From: or To: fields with addresses that fit this ‘mask’.

More information:

Fitting a Mask to Recipient Addresses (see page 148)
Fitting a Mask to the Sender's Address (see page 148)
Address Mask Example (see page 149)
Address Mask Notes (see page 149)

Fitting a Mask to Recipient Addresses

If a reviewer uses the iConsole Address Selector dialog to send an email to all members of a user group or dynamic address list, the iConsole uses the address mask to pick the most appropriate recipient addresses. It then writes these addresses to the To: (or Cc: or Bcc:) field.

How does it pick recipient addresses? When a reviewer clicks the Send button, the CMS returns a list of matching user accounts. For each user account, the iConsole then examines the email addresses associated with that account and chooses the first address that fits the address mask. Finally, it writes that address to the To: field.

If no SMTP addresses fit the mask, the iConsole attempts an LDAP lookup on the recipient's Windows account in order to find an address. If successful, that address is written to the To: field.

More information:

How Does the iConsole Set the From: Field in Audit Emails? (see page 143)

Fitting a Mask to the Sender's Address

If audit emails do **not** use a global sender, you can set an address mask to ensure that the From: field is always populated with a sender address that matches a designated SMTP address pattern.

When a reviewer sends an audit email, the iConsole compares all SMTP email addresses associated with the reviewer's CA DataMinder account. It then chooses the first address that fits the address mask and writes that address to the From: field.

If no SMTP addresses fit the mask, the iConsole attempts an LDAP lookup on the reviewer's Windows account in order to find an address. If successful, that address is written to the From: field.

More information:

How Does the iConsole Set the From: Field in Audit Emails? (see page 143)

Address Mask Example

If the address mask is .@unipraxis.com and the reviewer is Lynda Steel, then these two addresses fit the mask:

```
LyndaSteel@unipraxis.com
lsteel@unipraxis.com
```

But these addresses do not:

```
LyndaSteel@unipraxis.co.uk
lsteel@unipraxis.co.uk
```

In this example, the iConsole chooses the first matching address it finds (LyndaSteel@unipraxis.com in this case) and writes that address to the From: field in the audit email.

Address Mask Notes

Note the following:

Sender address masks only

- If both AuditEmailAddress and AuditEmailAddressMask are specified, then AuditEmailAddress takes precedence.
- If AuditEmailAddress is present but empty, the iConsole uses the first address that matches AuditEmailAddressMask.
- If a CA DataMinder user has no associated email addresses, then the sender identity is generated automatically, depending on whether single sign-on is enabled on the CMS.

Sender and recipient address masks

- If a template is used to compose the audit email, any sender or recipient addresses specified in that template will take precedence.
- An empty value in AuditEmailAddressMask is equal to a .* regex (regular expression). That is, the mask fits all email addresses associated with the user. If the mask fits more than one email address, then the iConsole uses the first one in the list.

More information:

How Does the iConsole Set the From: Field in Audit Emails? (see page 143)

Specify the LDAP attribute used to populate To, Cc and Bcc lists

In the iConsole Compose Mail dialog, when a reviewer starts typing an address in the To, Cc, or Bcc fields, the iConsole displays a drop-down list of all matching users in Active Directory.

By default, this list is based on an LDAP lookup operation that queries the displayName attribute in Active Directory. But if required, you can specify which Active Directory attribute is used to populate the drop-down lists. To do this, you need to edit a value in the Web registry key on the front-end Web server. Within this registry key, edit the following value:

FriendlyNameLDAPAttribute

Type: REG_SZ

Data: Defaults to displayName. This specifies the Active Directory attribute that the iConsole uses to generate drop-down lists of matching users in the To, Cc, or Bcc fields of the Compose Mail dialog.

For example, to revert to version 6.0 behavior you can set this registry value to cn. This will force the iConsole to query the ‘common name’ Active Directory attribute.

General Setup Tasks

After installing the iConsole, you may need to perform the following tasks.

Configure iConsole Log Files

To configure iConsole logging, modify values in the \WebService\Logging registry key. Within this registry key, edit the following values:

LogLevel

Type: REG_DWORD

Data: Defaults to 2. This value determines the logging level. For example, you can configure the iConsole to log only errors or warning system messages.

Log entries are written to the iConsole_<date>.log file, where <date> is the date and time when the log file was created. The file is located in CA's \data\log subfolder of the Windows All Users profile. The supported logging levels are:

1 Errors only

2 Errors and warnings

3 As 2, plus informational and status messages

Note: If you set LogLevel=3, log files grow extremely rapidly. This level of logging is provided for testing and diagnostic purposes. For example, it shows storage and retrieval on every resource item.

LogMaxNumFiles

Type: REG_DWORD

Data: Defaults to 10. This value specifies the maximum number of log files. When the maximum number of log files exists and the maximum size of the latest is reached (see below), the oldest log file is deleted.

LogMaxSizeBytes

Type: REG_DWORD

Data: Defaults to 1,000,000. This value specifies the maximum size for each log file. When the current log file reaches its maximum size, the iConsole creates a new log file. Log entries are written to an iConsole_<date>.log file, where <date> is the date and time when the log file was created.

AlwaysLogSearchCriteria

Type: REG_SZ

Data: Defaults to True on an Enterprise installation, defaults to false on Express installations. If set to True, the iConsole automatically adds information messages containing search criteria and parameters to the iConsole log file regardless of the LogLevel entry. Search criteria are logged before the search is run and include the name and version of the associated stored procedure file, plus the parameters used. The search outcome indicates the success or failure of the search, plus:

- **If the search succeeds,** log entries also show how many search results were returned and how long the search took.
- **If the search fails,** it gives a reason for the failure.

If set to False, the iConsole uses the value of LogLevel.

Set Up iConsole Servers for Network Load Balancing

If required, you can deploy your iConsole servers in a cluster (or Web farm) using the Windows Network Load Balancing technology. If you do implement an iConsole clustered deployment, note the following:

- You can cluster both your application servers and front-end Web servers (both must be installed on each node) or just the front-end Web servers.
- The software environment (operating system and node configuration) must be identical on each node.
- On each node, the cluster operation mode must be configured to use the **IGMP multicast** protocol. Do this in the Cluster Parameters tab of the Network Load Balancing Manager.

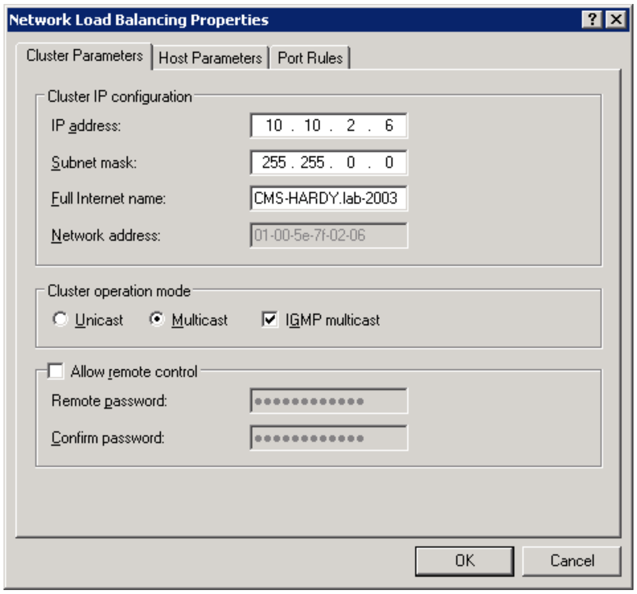

Cluster IP configuration

IP address: 10 . 10 . 2 . 6

Subnet mask: 255 . 255 . 0 . 0

Full Internet name: CMS-HARDY.lab-2003

Network address: 01-00-5e-7f-02-06

Cluster operation mode

○ Unicast ● Multicast ☑ IGMP multicast

□ Allow remote control

Remote password: ●●●●●●●●●●●●

Confirm password: ●●●●●●●●●●●●

Network Load Balancing Properties dialog: Cluster Parameters tab

- When defining the port rules on each node, the **filtering mode** must be set to ‘Multiple host’ and, crucially, the **Affinity** setting must be set to ‘Single’. Set these options in the Add/Edit Port Rule dialog of the Network Load Balancing Manager. For details about assigning ports, see the next section.

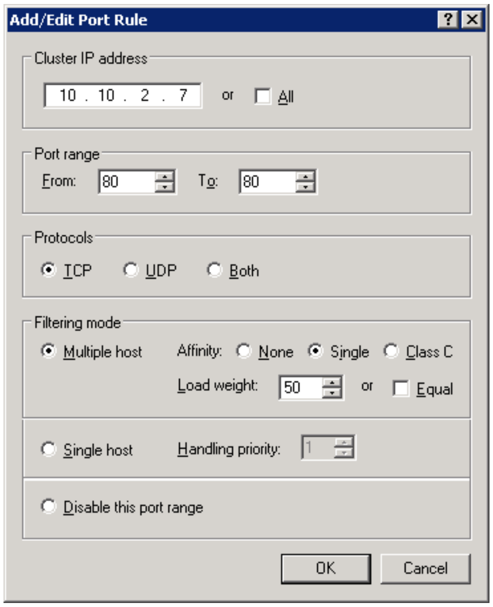

Add/Edit Port Rule

Cluster IP address

10 . 10 . 2 . 7 or □ All

Port range

From: 80 To: 80

Protocols

● TCP ○ UDP ○ Both

Filtering mode

● Multiple host Affinity: ○ None ● Single ○ Class C

Load weight: 50 or □ Equal

○ Single host Handling priority: 1

○ Disable this port range

Add/Edit Port Properties dialog:

Port range: These must both be set to 80.

Affinity: This must be set to Single. See the Important below.

Load weight: Defaults to 50 (ideally suits a two-machine cluster).

Important! The Single Affinity setting is essential. It ensures that iConsole users cannot be switched to an alternative node midway through a session. This configuration is also known as ‘sticky sessions’.

- On each node, configure a common encryption key for the ViewState encryption.

More information:

Port Rules (see page 154)
ViewState Encryption (see page 154)

Port Rules

When you define the port rules for your cluster nodes, we recommend that you filter on specific ports. We also recommend separate rules for HTTP and SSL for each port, across all nodes. Filtering on a specific port rather than a port range eliminates the risk of a third party application triggering a node switch if it takes up resources on a port covered by your iConsole port rule. For example:

- **HTTP:** Set the start port and end port to 80.
- **SSL:** Set the start port and end port to 443.

These ports must match those TCP and SSL ports defined in Microsoft IIS in the Advanced Web Site Identification dialog (this is where you define multiple identities for the iConsole Web site).

ViewState Encryption

If you deploy your iConsole front-end Web servers in a cluster, you need to use a common encryption key for the ViewState encryption. By default, the cluster nodes each use an auto-generated encryption key, but this can cause problems if a node switch occurs. Specifically, the iConsole browser can lose its connection to the CMS after the iConsole session times out.

To specify a common encryption key on each node in the cluster

1. On each node, you need to edit the .NET file machine.config. Find it in this folder:

   %windir%\Microsoft.NET\Framework\v2.0.50727
   \CONFIG

2. In this XML file, locate the machineKey:

   ```
   <machineKey
     validationKey="AutoGenerate...
     decryptionKey="AutoGenerate...
     validation=<encryption_algorithm>"
   />
   ```

   Note that the validation parameter can be set to any encryption algorithm, such as SHA1 or 3-DES.

3. Now change the machineKey parameters to:

   ```
   <machineKey
     validation=<encryption_algorithm>"
     validationKey=<hex_key>"
   />
   ```

   Where hex_key is an encryption key (in hexadecimal format). You can use any length key, but be aware that there is a trade-off between security and response times. Longer keys, (say, 128-bit) provide stronger security but also mean that data requests take longer to service. Conversely, shorter keys mean data requests are serviced more quickly but provide weaker security.

4. Make this encryption change on all nodes (that is, iConsole servers) in the cluster.

5. Finally, you need to restart Microsoft IIS on each node:

   a. In Cluster Administrator, take the cluster offline.

   b. Restart IIS on all nodes.

   c. Bring the cluster back online.

Manage iConsole Settings

You can configure the iConsole to suit your personal audit, print, search and screen preferences.

Administrators with the Admin: Manage iConsole privilege can configure global iConsole settings for all users. For example, an administrator can specify a default home page or enforce auditing settings. Administrators can also specify which settings a user is permitted to personalize.

In addition, administrators can define a default home page that appears every time a user logs on. The home page contains configurable portlets showing, for example, reports, searches, or messages to users. Again, administrators can specify which portlets a user is permitted to add, remove, or personalize.

Administrators should work together with a senior security auditor to identify which portlets are relevant to the security auditors at this site.

Configure Global Settings

If you have the 'Admin: Manage iConsole' privilege, you can configure global settings for all iConsole users.

To configure global settings

1. Click the ⚙ Settings link in the top right of the iConsole screen.

   The Settings dialog appears.

2. Click the Global tab.
3. Configure the global settings as required. Available settings are shown on the following subpages:

   Audit

   These settings control iConsole audit behavior. For example, you can specify whether focus automatically moves to the next event after auditing and events are removed from the results screen after being audited.

   Printing

   These settings determine whether printed search results include the event text content, summary details, and extended information.

   Search

   These settings control how search results are displayed. For example, you can specify the number of results per page and whether result rows can display multiple lines of text.

   If your iConsole supports content searches, you can specify which content proxy server to use.

   Home Page

   These settings configure the default behaviour for iConsole home pages.

   Show Home Page

   Specifies whether a home page is displayed when a user logs in to the iConsole.

   Users can change their personal settings if they do not want a homepage (unless you enforce this global setting).

   Non-Admin Users Can Define Portlets

   Specifies whether users are permitted to define their own portlets.

   Note: 'Non-Admin Users' are users who do not have the Admin: Manage iConsole privilege.

   Allow Portlet Auto Refresh

   Specifies whether portlets on the home page refresh automatically.

   Excluded Portlet Types

Specifies which portlet types are available for inclusion on the home page.

4. (Optional) Select the Enforce checkboxes for individual settings to override personal preferences defined by users, or to fix preferences so that users cannot modify them.

   Changes made to global settings become effective when a user next logs on to the iConsole. Enforced settings are grayed out in the user's personal Settings dialog.

   Example: If you always want to remove an event after a user has updated the event audit details, select the associated Enforce check box to fix this audit setting for all users.

5. (Optional) Click Change Password to reset your CA DataMinder password. You must enter this password when you log onto any CA DataMinder console.
6. **Note:** The Change Password button does not appear if you use Single-Sign On.
7. Click the OK button.

   The global settings are active.

Create a Global Portlet

Administrators maintain common portlets that any user can select and configure. You can pin mandatory portlets to every home page.

Note: Administrators must create common portlets before creating a default homepage.

To create a common portlet

1. Select the Home tab, and click Customize,

   The Customize Home Page wizard opens.

2. Click New Portlet and click a portlet type.

   The Home Page Portlet wizard opens.

3. Define required parameters.

   a. Specify settings that depend on the portlet type.

   b. Select the check box to specify that anyone can use this portlet.

   c. (Optional) Select the check box to pin this portlet to every home page.

   You have defined the basic details.

4. Click Close.

   The portlet is saved and is available to all users.

Note: Click Settings on the portlet to customize it further.

Edit a Portlet

Use the Home Page Portlet Wizard to edit portlet settings.

Note: Click the question mark in the portlet title bar to get details about the portlet.

To edit a portlet

1. Click ⚙ Settings in the portlet title bar.
2. Click the ✎ Portlet Definition button to launch the Home Page Portlet Wizard.
3. Specify the portlet settings.
4. Click Finish to close the wizard.

Admin

Click Settings on an Admin portlet to edit a policy pack, such as the Standard Policy Pack or the Classification Policy Pack.

BOE Report

Click the portlet title bar to view the full report in a new window and save the report as PDF or HTML file.

Important: The BOE Report portlet stores the unique ID of the BOE report. If you copy a report, or create a new one, and delete the old report, you change the report ID. Without the ID, the BOE portlet can no longer identify the report, and fails to load. To maintain the report ID, edit the original report in place, or export the report as a BIAR file, edit, and reimport it.

The wizard contains the following fields:

BOE Report

BOE Report

Specifies the BOE report that this portlet displays.

Output Format

Display Type

Specifies the default output format as PDF or HTML.

Clock Portlet

The wizard contains the following fields:

Timezone

Specifies the timezone.

Default: local

Label

Defines custom text to display together with the clock.

Example: Use the label to display the names of contacts or offices in this timezone.

Dashboard Chart Portlet

Hover the mouse over a chart for more details. The wizard contains the following fields:

Dashboard Chart

Chart

Specifies the dashboard that is used to plot the chart.

Output Settings

Custom Date Range

Specifies the time interval of the chart data. You can specify start and end values of a custom time period, or can select a predefined time period.

Examples: "This month", "today", "last year"

External

Click Settings on an External portlet to specify external web content to include on the home page. You can use this portlet to display custom content, such as internal web pages. The wizard contains the following fields:

External Portlet URL

Defines the URL of the web content to display in the portlet.

Favorites Portlet

A Favorites portlet shows a list of direct links to the searches and reports that you use most often. The wizard contains the following field:

Favorites

Select favorites items from a list of available BusinessObjects (if available), Reports, and Searches.

Report Portlet

The wizard contains the following fields:

iConsole Report

Report

Specifies the report that this portlet displays.

Output Settings

Automatically refresh the portlet?

Specifies whether this portlet refreshes automatically to display time-critical information. If you clear this check box, update the portlet by reloading the page in the browser. Administrators can deactivate this option globally.

Refresh Rate

Specifies the time interval after which the portlet refreshes. To improve overall performance, set the refresh rate to a value higher than 60 seconds.

Default: 60 seconds

iConsole Search

The wizard contains the following fields:

iConsole Search

Search

Specifies the search that this portlet displays.

Results Columns

Available Items

Displays possible columns in the search results. Select items and click the right arrow buttons to add them to the results.

Selected Items

Displays the columns that you have selected. Select items and click the left arrow buttons to remove them from the results.

Output Settings

Automatically refresh the portlet?

Specifies whether this portlet refreshes automatically to display time-critical information. If you clear this check box, update the portlet by reloading the page in the browser. Administrators can deactivate this option globally.

Refresh Rate

Specifies the time interval after which the portlet refreshes. To improve overall performance, set the refresh rate to a value higher than 60 seconds.

Default: 60 seconds

Number of results

Specifies that the portlet shows the top n results.

Default: 5

Message

Use Message portlets to display reminders to yourself or announcements for everyone. The wizard contains the following field:

Message

Defines the message content. Use standard HTML syntax if you want to reference files, for example to embed images. Store the files in the web directory of the front-end server.

Format: Plain text or HTML

Example: <h3>Reminder</h3><p><img src="warn.gif" />The server will be down for maintenance on 2011-03-30 between 22.00 and 23.00.</p><p>All users <b>must</b> log off during this period. </p>

RSS Feed

An RSS Feed portlet specifies an RSS news feed. The Wizard contains the following fields:

RSS Feed

Defines the URL of the RSS feed. For example:

`http://feeds.bbci.co.uk/news/rss.xml`

Number of results

Specifies that the portlet shows the top n results.

Default: 5

Create a Default Home Page

You can create a default home page for users that have not defined a personal home page.

Note: Administrators must create common portlets before creating a default homepage.

To specify a default home page

1. Select the Home tab, and click Default Layout.

   The default homepage layout displays.

2. Click Customize, Global Portlets.
3. Select the checkboxes to specify global portlets for the default home page.

   The portlets are added to the default home page.

4. Click the Layout tab to specify a column layout and click close.
5. (Optional) Drag portlets to the required location in a column. You can collapse portlets so they use less space.
6. (Optional) Click Settings in the portlet title bars to configure specific parameters for each.

   The default home page is saved automatically. It is now available to all users.

Note: Instructions describing how users define personal home pages are included in the *iConsole User Guide*.

Make Portlets Mandatory

You can specify which portlets are mandatory and appear on every user's homepage. Users cannot close mandatory portlets. You can specify that a common portlet is mandatory while you create the portlet, or you can modify this behavior later.

To make a portlet mandatory

1. Select the Home tab, and click Default Layout.

   The default homepage displays.

2. Click Global Portlets, and click Edit on a global portlet.

   The Home Page Portlet Wizard opens.

3. Select the check box to pin this portlet to every home page. Click Ok.

   Changes are saved and you return to the default home page.

4. Click Save Layout.

   The portlet appears on all home pages.

Disable Personal Home Pages

By default, users have the option of customizing the home page. If you disable personal home pages, you should create a default home page.

To disable or enable personal home pages

1. Click Settings, Global

   The Global Settings window opens.

2. Select Show Home Page under Home Page Preferences to enable personal home pages. Clear the checkbox to disable them.
3. (Optional) Select Allow Portlet Definition to allow users to configure personal portlets. Clear the checkbox to restrict all users to common portlets.
4. (Optional) Select Enforce to make this setting mandatory. This overrides personal settings of users who disabled or enabled the home page for themselves.
5. Click OK.

   The global settings are applied.

Disallow Portlet Types

Many iConsole portlet types are available, but your site may not require all of them. To increase usability and performance of the home page, Administrators have the option of disallowing individual portlet types globally. Exclusion settings are always enforced globally.

For example, if the administrator has disallowed RSS Feed portlets, users can no longer use or create RSS Feed portlets.

To exclude portlet types

1. Click Settings, Global

   The Global Settings window opens.

2. Select the portlet types that you want to exlude in the Excluded Portlet Types section under Home Page Preferences.
3. Click Apply and Close.

   The global settings are applied.

Start the iConsole

Users simply browse to a specified URL to start using the iConsole. If single sign-on is enabled on the CMS (see next section), they will not need to log on, but can begin searching for events immediately. The iConsole URL is:

```
http://<FE_Server>/<virtual dir>
```

where:

<FE_Server>

Is the name or IP address of the host machine for the front-end Web server.

<virtual dir>

Is the virtual directory for the front-end Web server. This defaults to 'cadataminder' but you can rename it—see the next section.

For example, if the front-end Web server is hosted on the server UX-WebSvr-01, the iConsole URL is:

```
http://UX-WebSvr-01/cadataminder
```

About Single Sign-On

If single sign-on (or 'SSO') is enabled for the parent CMS of the application server, users skip the logon dialog when they start up the iConsole. Instead of the user supplying credentials to access the console, CA DataMinder relies on the fact that the user has successfully logged into Windows as sufficient authorization to allow them to log on to the CA DataMinder account of the same name.

(To log on using a different account, a user must first log out of the iConsole, then log back on from the Logon screen.)

To configure CA DataMinder to use single sign-on, you must edit the CMS machine policy. You can also grant the administrate privilege Admin: Use single sign-on to individual users (this overrides the CMS policy). Note that account names for CA DataMinder users must be the same as their native Windows user name (sometimes referred to as the user logon name). That is, an account name prefixed with the user's domain, for example, unipraxis\lsteel.

For full details, see the Administration console online help; search for 'single sign-on'.

Chapter 8: iConsole Standard Searches, Reports and Policies

This section contains the following topics:

Standard Searches, Reports and Dashboard

The iConsole standard searches, reports and dashboard comprise a set of predefined event searches and reports.

Reviewers can run standard searches and reports to retrieve captured events from the CMS database and gain insight into user activity across your organization. Reviewers can also customize their own versions of these searches and reports. Results are typically shown in a tabular layout.

iConsole dashboards provide quick graphical breakdowns of user activity. A dashboard shows charts and metrics of captured events, broken down by type, channel, severity, and so on. Each chart or metric occupies a pane in the dashboard display. Reviewers can configure and rearrange panes to focus on the areas that most interest them.

Note: CA DataMinder can also integrate with BusinessObjects Enterprise, allowing you to run BusinessObjects reports for CA DataMinder in the iConsole. For details, see the *Reports Integration Guide*.

More information:

Requirements (see page 165)
Available Searches (see page 170)
Available Reports (see page 171)
Installing Searches, Reports, and the Dashboard (see page 175)
Defining New iConsole Searches (see page 177)
Customizing the iConsole Stylesheet (see page 178)

Requirements

Note the following requirements for the iConsole searches, reports and dashboard.

CMS and iConsole Requirements

Before you installing CA DataMinder standard searches and reports, verify that the following components are already installed:

CMS

CA DataMinder 12.0 or later

iConsole servers

CA DataMinder 12.0 or later

Browser

The iConsole homepage now uses Sencha Ext JS to create charts. In previous releases, the iConsole used Adobe Flash.

Database Requirements

The reports and dashboard reports.msi support the following databases:

Databases

CA DataMinder servers need sufficient memory and processing power to run your chosen database application. See your database documentation for details. The supported databases are:

- Oracle 10g (10.2.0.4) or 11g (11.1.0.7), 11g Release 2, or later

  Oracle 10g users may see improved search and report performance by applying the following Oracle fix 5765456:7. See the following My Oracle Support (formerly Oracle Metalink) notice: "Bug 5040753 Optimal index is not picked on simple query / Column group statistics."

- Microsoft SQL Server 2005

  We recommend the SP2 release.

- Microsoft SQL Server 2008
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2008 Express Edition

  Not recommended for CMSs or FastStart base machines.

- Microsoft SQL Server 2012

  Note: SQL Server is supported on Windows servers only. Verify that the SQL Server Browser service has started.

Primary User Requirements for Dashboard

To ensure that the dashboard is correctly populated with data, your CA DataMinder primary user (or login) must be correctly configured to run data warehousing jobs. Your primary user is the account that you specify in the Database Accounts screen of the CMS installation wizard. It is the main account that CA DataMinder uses to access the CMS database.

Note: By default, the CA DataMinder primary login in is WGNUSER.

Verify that your primary database account has the required privilege or role before you install the dashboard.

Oracle CMSs

Your primary user needs the CREATE JOB privilege. To verify that your primary user has this privilege, run this command:

```
SELECT * FROM DBA_SYS_PRIVS WHERE PRIVILEGE ='CREATE JOB';
```

If you primary user does not have this privilege, run this command line as an Oracle SYS user:

```
GRANT CREATE JOB TO WGNUSER; COMMIT;
```

Note: The CREATE JOB privilege is granted automatically when you install a version 12.0 or later CMS.

SQL Server CMSs

Verify that your primary login (such as WGNUSER) has been added as a user in the msdb system database. Verify that that this user has been granted the msdb database role SQLAgentUserRole.

Note: The SQLAgentUserRole role is granted automatically when you install a version 12.0 or later CMS.

Important! Do not confuse this WGNUSER **database user** in the msdb system database with the WGNUSER **instance-level login**!

SQL Server Requirements for Dashboard

Note: These sections do not apply if your CMS uses Oracle or SQL Server 2005 Express.

To successfully install the dashboard, your CA DataMinder primary login (such as WGNUSER) must have a database user in the msdb database. This database user must already been granted the SQLAgentUserRole role. This allows it to create and schedule SQL Server Agent jobs. Specifically, your primary login needs to create and run scheduled aggregation jobs.

New CMSs

The SQLAgentUserRole role in the msdb database is granted automatically to your primary login when you install a new SQL Server CMS. However, if the Standard Searches and Reports installation failed on a new CMS, you must first clean up the database before attempting to reinstall.

Cleaning up After a Failed Installation

If the reports.msi installation fails (for example, because your primary user did not have the correct database roles), you will need to clean up your CMS database before trying to reinstall. Specifically, you will need to remove any orphaned database users (that is, users no longer associated with a login), plus any dashboard aggregation jobs and related schedules.

For example, in SQL Server 2005 Management Studio:

1. Delete any orphaned users in the msdb database. In the Object Explorer pane:
   a. Select the msdb database in the \Databases\Systems Databases branch.
   b. Navigate to the \msdb\Security\Users branch.
   c. Delete any msdb database users with the same name as the CA DataMinder primary login.

   Important! Do not delete the primary login from the \Security\Logins branch!
2. Delete any orphaned users in the master database.
   a. Select the master database in the \Databases\Systems Databases branch.
   b. Navigate to the \master\Security\Users branch.
   c. Delete any master database users associated with the CA DataMinder primary login.
3. Delete any aggregation jobs and schedules:
   a. Navigate to the \SQL Server Agent\Jobs branch and delete any DLP_Aggregation_<DB_name> jobs.
   b. Right-click the \SQL Server Agent\Jobs folder and choose Manage Schedules.

c. Confirm that all schedules named DLP_Aggregation_Schedule_<DB_name> have been deleted.

Additional Notes

For all reports and the dashboard, note the following:

DBMS Testing

Reports have been tested using the supported versions of Oracle and SQL Server.

Maximum Results Rows

For performance reasons, the reports default to a maximum of 1000 rows of results.

To increase this limit, configure the $ROWLIMIT parameter in the report definition XML and republish the reports. For details, see the *iConsole Search Definition Guide*; search for '$ROWLIMIT parameter'.

More information:

Database Requirements (see page 166)

Policy Security Models Not Compatible With Some Reports or Review Queue

Certain reports, particularly the compliance reports such the Repeat Offender report and Compliance Audit Report, are not designed for use with Policy security models. This is also true for the Review Queue feature and the associated Reviewer search.

These reports and the Review Queue are explicitly designed to be run in conjunction with the Management Group security models. That is, they return data about users in specific user groups.

Important! We recommend that any users who need to run these reports or the Reviewer search are assigned to a Management Group security model, not to a Policy security model.

Available Searches

The following standard searches are available:

Content Search

Retrieves events based on their text content. Content searches can identify clusters of related documents, defined by their characteristic text patterns that reveal a shared subject or theme.

Content searches use intelligent pattern-matching technology to analyze the text content of indexed events in a content database and retrieve events that match the search criteria.

Note: Content searches are only available if explicitly included in your license agreement.

Data At Rest Standard Search

Retrieves Data At Rest events that match specific criteria. These events typically include files and other items scanned by the File Scanning Agent (FSA) and Client File System Agent (CFSA), plus files stored in an archive.

Data In Motion Standard Search

Retrieves Data In Motion events that match specific criteria. These events include: emails; network events entering or leaving your corporate network, including webmail and IM attachments and FTP file transfers; IM conversations captured by CA DataMinder Network; and web events, including file uploads.

Data In Use Standard Search

Retrieves Data In Use events that match specific criteria. These events include: files copied or saved to removable devices (such as USB flash drives), writable CD and DVD drives, and network locations; files sent to a printer; and instances when a user runs a specific application.

Quarantine Search

Retrieves quarantined emails.

Reviewer Search

(Available only with the Review Queue) Retrieves events in a user's personal review queue. These are events waiting to be reviewed.

Recent Incidents

Retrieves all events for a specified date range.

Standard Search

Retrieves events that match specific criteria. It covers all event types and includes all the commonly used search filters.

Available Reports

The available standard reports are summarized in the following sections.

Compliance Reports

These include:

Compliance Audit Report

Shows workload statistics for each reviewer in your management group(s). It shows the number of events allocated to each reviewer and the number already reviewed over a specified period. Optionally, it also indicates how long reviewers spend reviewing individual events.

Employees Not Reviewed Report

Lists users who have not had any of their associated events reviewed over a specified period. A 'reviewed event' is an event with one or more issues.

Proof of Supervision Report

Shows reviewed events as a percentage of all captured events, by user or group.

Repeat Offender Report

Shows the number of incidents associated with each 'offender'. An offender is any user associated with an issue. They can be the sender or a recipient of the message.

Review Latency

Shows the level of events that are still unreviewed. The report identifies the reviewers and the user groups associated with the unreviewed events. It also indicates how many days these events have been waiting to be reviewed. Percentage scores allow you to compare review rates by reviewer or/and user group.

Reviewer Activity

Shows the activity of individual reviewers in terms of the number of events viewed and audited. This report enables managers to monitor reviewer activity and, if necessary, confirm that events are being audited correctly.

Note: These reports are not designed for use with Policy security models. See the reference below for details.

More information:

Policy Security Models Not Compatible With Some Reports or Review Queue (see page 169)

Incident Reports

These include:

Incident Rate By Policy Report

Shows a breakdown of all the policies that have been triggered over a specified period, showing their counts as a percentage of all events processed. By default, it includes all policies, but it can be refined to focus on just a subset. It calculates percentages based on all events in the database captured over the specified period.

Incident Summary Report

Shows how many times a user (employee) has caused a trigger to fire during a specified period. You can use this report to search for details of captured events associated with a specific user, group or trigger, incoming or outgoing events, or events captured over a specific period.

Incidents By Location Report

Provides a summary of 'Data At Rest' trigger violations by file location (that is, the host machine and folder containing the file). It shows the number of incidents for each file location.

Incidents By Policy and Action Report

Shows how many file policies were triggered and the subsequent actions (for example, capture, delete, move, or copy a file).

Incidents By Policy and Channel Report

Shows the number of violations by policy or class for specific users or groups, broken down by channel. For example, it shows how many violations were caused by emails, FTP transfers or instant messages. Reviewers can analyze the results to see which policies trigger most often and on which communication channels.

Incidents By Policy and Time Period Report

Shows the number of violations over time broken down by policy or class, for specific users or groups. For example, it can show changing violation counts on a weekly or monthly basis. Reviewers can analyze the results at a policy or class level to understand the overall trend over time.

Incident Cause Frequency Report

Shows incident counts for specific policy conditions: Parameter 7 words detected by a Document Classifier trigger; e-mail senders; and e-mail subject lines. This diagnostic report is designed for policy administrators. It allows them to understand how policy is being applied to real data and, if necessary, to fine tune those policies to reduce false positive results. Note that you can also drill down into the results to see the individual events.

Issue Reports

These include:

Detailed Issue Report

Lists current details for individual issues, including audit status and resolution, reviewers, and associated users.

Issues By Status or Resolution Report

Shows the number of issues for specific users or groups, broken down by audit status or resolution. For example, reviewers can see the number of issues with an audit status of 'Escalate' or 'Pending' for events captured in the last month.

Review Queue

The Review Queue feature is optional. It includes the following administrative reports:

Review Queue Configuration

Shows the event selection rules for your management groups when a review queue job runs.

Review Queue Diagnostics

Provides details about the most recent review queue run. For each step, it shows the execution time, the actual query statement and the status.

Review Queue History

Shows summary details for previous database review queue runs, including run times, status and event counts.

Note: For instructions on how to customize and manage the Review Queue, see the *Review Queue Implementation Guide*. This is available to download from CA Technical Support.

More information:

Policy Security Models Not Compatible With Some Reports or Review Queue (see page 169)

BusinessObjects Intelligence Reports for CA DataMinder

(Available only if CA DataMinder integration with BusinessObjects Intelligence is enabled.)

The BusinessObjects page of the Review tab shows the following items:

InfoView link

InfoView is the BusinessObjects web portal. Click the link to open an InfoView window.

CA DataMinder Reports

The reports are versions of the following standard CA DataMinder reports redesigned for BusinessObjects:

- Compliance Reports (see page 171)
- Incident Reports (see page 172)
- Issue Reports (see page 173)

This folder maps directly to the CA DataMinder folder in the InfoView Document List.

Inbox

The inbox includes scheduled BusinessObjects reports for CA DataMinder that ran successfully.

My Favorites

This folder includes your customized versions of BusinessObjects reports for CA DataMinder.

This folder maps directly to the following folder in InfoView: \All\My Favorites.

Note: If you want your customized reports to be listed in this 'My Favorites' folder in the iConsole, you must save them to the \All\My Favorites in InfoView.

The iConsole dashboard shows charts and metrics of incidents and violations, broken down by policy, channel or impact. Data is shown in separate panes which you can configure and rearrange to focus on the areas of most interest.

You can also drill down into a pane to view a list of the underlying events or, for pie charts, a further chart showing a breakdown of the data in an individual pie slice. Likewise, in a 'table' pane you can click individual cell values to view the underlying incidents.

Note: Dashboard charts and metrics are based on aggregated data.

Installing Searches, Reports, and the Dashboard

This section describes how to install the iConsole standard searches, reports, and the dashboard. You must repeat the installation on your CMS, your iConsole application servers and your iConsole front-end Web servers.

To install on the CMS

Before you start, make sure that your database engine is properly configured.

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose iConsole Standard Searches and Reports and then click Install.

   This launches the Standard Reports installation wizard in a separate window.

4. In the Standard Reports installation wizard, navigate to the Custom Setup screen.

5. In the Custom Setup screen, choose the reports you want to install. You must choose the same combination of reports when you run reports.msi on your iConsole servers (see the following sections).

6. In the final wizard screen, click Install to start the file transfer.

To install on iConsole application servers

1. Follow steps 1 through 4 of the CMS instructions above.

2. In the Custom Setup screen, choose the reports you want to install. You must choose the same combination of reports that you installed on your CMS.

3. In the Administrator Credentials screen, enter account details (name and password) for the Primary Administrator.

   Note: CA DataMinder creates the Primary Administrator account during installation. This account has full administrative privileges and full management group coverage.

4. The installation wizard now has all the information it needs. Click Install to start the file transfer.

To install on iConsole front-end Web servers

1. Follow steps 1 through 4 of the CMS instructions above.

2. In the Custom Setup screen, choose the reports you want to install. You must choose the same combination of reports that you installed on your CMS.

3. In the iConsole Application Server Details screen, if no iConsole application server is installed locally you will need to specify which iConsole application server to connect to.

Server

Specify the name or IP address of the machine hosting the application server. Type localhost to specify the local machine.

Port

Specify the TCP port used for communication between front-end Web server and the application server. This defaults to port 80. If you specify a non-default port, you must also update the application server to use the same port.

Use SSL

If you intend to use SSL to communicate over a secure port (for example, 443), select this check box. This ensures that the port used for communication between the front-end Web server and the application server will use SSL.

4. The installation wizard now has all the information it needs. Click Install to start the file transfer.

More infromation:

Advanced Dashboard Configuration (see page 199)

Defining New iConsole Searches

To create new event searches and make them available to all iConsole users, you must first write a stored procedure (SP) in the CMS database. This ‘search SP’ contains SQL statements that define a specific search for captured or imported events. CA DataMinder supports search SPs for both Microsoft SQL Server and Oracle databases.

You must then create and install an XML search definition file onto the CMS. The search definition file defines the search parameters, including customizable parameters, and the layout of the search results screen in the iConsole. It also references the name of a stored procedure (SP) in the CMS database.

After installing the search definition file, you can test a new search in the iConsole. Specifically, you need to confirm that the search correctly references the search SP and returns a valid set of search results.

Finally, after testing a new search, you must publish it. When you publish a new search, it becomes available to all other iConsole users and is listed in the Predefined Searches list in the iConsole Search screen.

This process is summarized below:

1. **Define a Search SP**. Each new event search requires an associated search SP in the CMS database.
2. **Create an XML search definition.** The search definition file specifies the search parameters and screen layouts.
3. **Manage the searches.** You perform the following tasks in the Manage Stored Searches screen in the iConsole.
    a. **Install the XML search definition.** This file gets installed onto the CMS.
    b. **Test the search.** Confirm that the search correctly references the search SP and returns a valid set of search results.
    c. **Publish the new search.** This makes it available to all other iConsole users.

Note: For further information, see the *iConsole Search Definition Guide*. This guide is available to download from CA Support at http://ca.com/support.

Back Up Search Files

All stored procedures (SPs) and search definitions are stored in the CMS database, so these will be covered by your general database backup procedures. Backing up and restoring the CMS database is described in the Database guide; search the index for ‘backups’.

However, we strongly recommend that you also back up any .sql and .xml files created for your custom searches so that you can reapply these if you need to rebuild your CMS database.

Note: The .sql and .xml files for the default searches are available on the CA DataMinder distribution media if you need to rebuild your CMS database.

Customizing the iConsole Stylesheet

In previous CA DataMinder releases, it was possible to customize the look-and-feel of the iConsole by editing the branded.css stylesheet. For example, you were able to customize the colors in the Incident Rate By Policy report.

In the current CA DataMinder release, you can still edit this stylesheet to customize the iConsole. However, such customizations are not officially supported and we cannot guarantee that they will continue to work in future CA DataMinder releases.

Standard Policies

The iConsole standard policies compromise a predefined set of policies drawn from the CA Foundation Policy Pack (FPP). You can customize these standard policies in the iConsole to quickly roll out CA DataMinder across your organization.

The FPP organizes policies into classes, such as 'Corporate and Regulatory Compliance' and 'Personally Identifiable Information (PII)'. Each policy class contains several individual policies. For example, the PII policies include 'Account Number' and 'Credit Card Information' policies.

Individual policies are based on triggers in the user policy, plus other key settings such as document classifications. However, you must edit these standard policies in the iConsole.

Note: For details about editing these standard policies, see the iConsole online help, the *Policy Guide*, or the *FastStart Implementation Guide*.

More information:

Available Policies (see page 179)
Who Do the Standard Policies Apply To? (see page 196)
FPP User Groups Created Automatically on the CMS (see page 196)
Install the Standard Policies (see page 198)

Available Policies

The iConsole standard policies compromise a predefined set of policies. You can customize these standard policies in the iConsole to quickly roll out CA DataMinder across your organization.

Standard policies are organized into classes, such as 'Corporate and Regulatory Compliance' and 'Personally Identifiable Information (PII)'. Each policy class contains several individual policies. For example, the PII policies include 'Account Number' and 'Credit Card Information' policies.

Individual policies are based on triggers in the user policy, plus other key settings such as document classifications. However, you must edit these standard policies in the iConsole.

The following sections provide summary descriptions of the CA DataMinder standard policies.

Note: The following sections do not necessarily list the complete set of policies. For example, the available policies may vary according to your CA DataMinder license. Use the iConsole to view the complete set of policies available to your organization.

More information:

Corporate and Regulatory Compliance Policies (see page 180)
Customer / Supplier Treatment Policies (see page 182)
Employee Behavior Policies (see page 183)
Intellectual Property (IP) Policies (see page 184)
Legal Policies (see page 185)
Non-Public Information (NPI) Policies (see page 185)
Personal Health Information (PHI) Policies (see page 187)
Personally Identifiable Information (PII) Policies (see page 188)
Security General / Corporate Policies (see page 192)
User Defined Policies (see page 195)

Corporate and Regulatory Compliance Policies

Anti-Money Laundering - OFAC

This policy detects suspicious financial transactions such as tax evasion or false accounting, especially with entities that appear on the U.S. OFAC list.

Bid Rigging Detection: Insurance

This policy identifies 'B' bids, and other electronic communications indicative of bid rigging, as it relates to the insurance industry.

Bid Rigging Detection: Municipal Bond Issuance

This policy detects language that indicates possible bid rigging related to Municipal Bond issuance.

Blast E-Mail

This policy monitors for blast e-mail which is sent to more than a specified number of external recipients at one time.

Bribes/Kickbacks/Quid Pro Quos/Blackmail

This policy detects involvement in bribery or blackmail schemes.

Broker Error

This policy detects indications that a broker has made or is attempting to correct an error with respect to trading.

Communication with Regulatory, Legal, and Governmental Authorities

Protect and control communications between an employee and regulatory, legal, and governmental authorities.

Fair and Balanced Advice

This policy detects unbalanced communication by recognizing claims and statements that focus solely on positive or negative aspects of a product, advice, or decision.

Information Destruction Alert

Electronic information can be eliminated as easily as it is created, making the uncontrolled destruction of retained information an unacceptable risk. This policy detects text indicative of a suggestion to eliminate e-mail messages, computer files, or documents. It also detects general references to retention rules.

Investment Advice Prohibition

This policy detects messages that appear to contain investment advice or recommendations.

Securities Parking

This policy is designed to look for evidence of two parties engaged in a possible "trade parking", or "wash trade", arrangement.

Solicitations: Charitable

This policy detects solicitations or requests for contributions to charities, student fundraisers, or other non-commercial and non-political organizations.

Solicitations: General

This policy detects language containing general references to contributions or solicitations for contributions.

Solicitations: Political

This policy detects solicitations or requests for contributions to political causes or campaigns.

Solicitations: Private Investments

This policy detects language containing references to contributions or solicitations for contributions to private investment activities.

Solicitations: Religious

This policy detects solicitations or requests for contributions to religious organizations.

Tax Advice Prohibition

In general, a Representative must be both qualified, and allowed by the firm, in order to offer 'advice' to a customer. This policy is designed to identify messages where a non-tax Professional offers tax advice to a public customer.

Trading in an Outside Account: Order Confirmations

This policy detects order confirmations so as to identify trading activity, for one's personal account, outside of firm-approved processes and/or procedures.

Trading in an Outside Account: Order Placements

This policy is designed to identify trade order placements, for one's personal account, outside of Firm-approved processes and/or procedures.

Whistleblower

This policy detects possible whistle blower situations and allows an organization to take appropriate steps in response.

Customer / Supplier Treatment Policies

Customer Complaints: Response Prohibition

Most companies do not allow their representatives to directly respond to a customer complaint. This policy analyzes outbound external e-mail for indications that a representative has directly responded to a customer complaint, which may or may not have been received initially by e-mail.

Customer Complaints: Unprofessional Responses

This policy analyzes outbound external e-mail for indications that a company representative has directly responded to a customer

complaint, which may or may not have been received initially by e-mail, in an unprofessional and/or un-empathetic manner.

Customer Conditioning

This policy detects communications to a customer that include pressuring language. This may include attempts to force the customer to accept products or services they do not want or need.

Customer Threats

This policy detects language that indicates pressure being used against a customer in order to limit business with competitors. This is an example of anticompetitive behavior and can be a violation of anti-trust regulations.

Exclusivity

This policy detects language that suggests an attempt to establish full control over sales to a third party. This is an example of anticompetitive behavior and can be a violation of anti-trust regulations.

Gifts and Entertainment

Gifts and entertainment form a common part of many business relationships, yet have the potential to create conflicts. This policy identifies when a business expense violates policy or law and becomes a gift.

Guarantees and Assurances

Guarantees, though often considered a part of "fair and balanced" communication, carry with them legal, regulatory, and financial risks, as well as risks to a firm's reputation. This policy detects prohibited guarantees or assurances and can be used to prevent them from reaching customers.

Unqualified Rebates or Benefits

This policy is designed to detect an offer of a rebate when the terms and conditions have not been met. This can be used as a method to offer money to a customer for excluding competitors or accepting otherwise unwanted products.

Employee Behavior Policies

Coercive Behavior and Intimidation

Coercive behavior and intimidation in the workplace can have significant negative impact on employee morale and productivity. This policy detects such behavior so that enforcement is confidential and immediate.

Communication with Competitors

This policy detects electronic communication between an employee and competitor companies.

Communication with the Press/News Organizations

This policy detects electronic communication between an employee and the press or media organizations.

Corporate Criticism

This policy detects criticisms and negative comments about the company, its products, or the management team.

Deceptive Language

This policy detects communications that may include false or misleading information. In addition, it will detect references that indicate inappropriate offline communications.

Discrimination and Racism

This policy detects inappropriate discriminatory language and/or actions based on race, gender, disability, sexual orientation, religion, age, and other legally protected classes. Sexual harassment related issues are covered by the Harassment policy.

Discrimination: Age

This policy attempts to identify communications containing words and phrases that indicate a likelihood that age discrimination is taking place or being referenced.

Fantasy Leagues

This policy identifies events and activities associated with participation in or running a fantasy sports league.

Gambling Prohibition

This policy detects gambling and betting among employees which is subject to various jurisdictional regulations. Fantasy leagues are covered by a separate policy.

Harassment

This policy detects harassment such as quid pro quo requests for sexual contact, or behavior that is designed to alarm or annoy others.

Inappropriate, Offensive and Sexual Language

This policy identifies communications indicative of offensive and sexual language.

Intent to Resign

This policy detects language indicative of an employee who is dissatisfied with their position or workplace and is actively engaged in seeking employment.

Jokes

This policy detects electronic communication of a wide range of joke formats and subjects. It does not address communication that originated outside the firm, but will capture such events if the recipient within the firm attempts to forward them.

Office Relationships: Romantic

This policy detects events of a romantic nature, or language indicating that such a personal relationship exists.

Outside Business Activity/Directorships/Employment

This policy identifies communications that suggests an employee is engaged in external business activities unrelated to the company; serving or considering serving on another company's board of directors; or is participating in other activities that might affect the employee's performance at the company.

Termination/Layoff Discussions

Protect communications concerning potential and pending terminations and layoffs.

UK Resumes/CVs

This policy is designed to detect UK resumes in standard format.

US Resumes/CVs

This policy is designed to detect US resumes in standard format.

Intellectual Property (IP) Policies

Confidential Trade Data

This policy detects confidential information such as trade secrets, proprietary processes and technical competitive differentiators.

Patent Applications

This policy detects non public patent applications.

Product and Design Specifications

This policy detects functional or marketing specifications of material, products, or services.

Proprietary Software Code

This policy detects software code, programs, and executables.

Technical Specifications or Designs

This policy detects technical designs and specification documents related to products or services.

Legal Policies

Attorney Client Privilege

When an uncontrolled privileged communication or document leaves an organization, any privilege associated with it may be waived. This policy prohibits such communication from being sent externally.

Discussion of Legal Proceedings

This policy detects events related to legal proceedings such as pending civil lawsuits, criminal proceedings, and/or administrative hearings or trials. Threats of contemplated litigation against the organization are not intended to be covered by this policy.

Potential Ethical Issues

This policy identifies potential ethical misconduct or claims of ethical misconduct and alerts the proper internal legal representative.

Potential Legal Issues

Often, questions are circulated internally about the legality of a particular action or business practice without informing a legal representative until the problem has been made public or resulted in some harm. This policy identifies such discussions and alerts the appropriate legal representative.

Threats of Litigation

This policy detects discussions indicating an outside party or an internal employee suggesting or overtly threatening to file a lawsuit against the company.

Non-Public Information (NPI) Policies

Board Minutes and Discussions

This policy is designed to detect events occurring between or concerning board members of an organization.

Corporate Contracts

This policy detects the language that is typically used in corporate contracts.

Customer Lists

This policy detects multiple occurrences of various types of customer contact information.

Draft Documentation

This policy can be used to prevent draft documentation, and discussions surrounding it, being sent outside an organization.

Financial Information - Balance Sheet

This policy detects content found on financial balance sheets.

Financial Information - Income Statement

This policy detects content found on financial income statements.

Financial Information - Projections

This policy detects the disclosure of financial projections.

Information Security Label Control

This policy detects sensitive material classified in various ways such as "confidential", "top secret", and "not for distribution".

Inside Information: Front Running/Trading Ahead

This policy detects messages exhibiting evidence that a market participant is attempting to profit financially by placing transactions before (in front of) another market player, or customer, by leveraging the information a "tipper" possesses about what that market player/customer intends to do.

Inside Information: Non-Public Company Information Loss

Protect and control non-public company insider information, such as management discussions.

Inside Information: Non-Public Financial Information Loss

This policy detects unauthorized disclosure of non-public company financial and stock information.

Inside Information: Rumors and Secrets

This policy detects unsubstantiated information or rumors about any organization or client for legal purposes.

Inside Information: Trading Ahead of Research

Disseminating and acting on non-public, inside information is illegal. The content of a research report may influence the price of the security being discussed. Parties may profit from this non-public information by placing trades ahead of the issuance of the research report. This policy is intended to detect language indicative of two or more parties disseminating non-public information regarding advance knowledge of pending research.

Internal Investigations

This policy detects the existence, purpose, and/or results of company specific investigative matters.

Internal IT Support Documents

This policy identifies internal IT system and support documentation.

Licensing Agreements

This policy is designed to detect information containing software license agreements.

Mergers and Acquisitions

This policy identifies discussions and documents pertaining to pending or proposed merger and acquisition transactions in which the organization is or will be participating. Transactions such as IPOs, private placements, and other prospectus offerings are not expressly included in this policy.

Pricing List

Pricing List

This policy is designed to detect nonpublic pricing information.

Project Information

This policy identifies various types of project information such as project plans, timelines, project codes, task lists, and issue lists related to project planning and deployments.

Restricted List

This policy detects items and content on restricted lists in e-mails and files. Restricted/Watch/Grey Lists are associated with services, products, companies, customers, or other defined business elements that have restrictions.

Sales Information

This policy detects company sales information, sales collateral such as tools, models, contracts, fee structures, and deal information, and other elements supporting the sales organization.

Personal Health Information (PHI) Policies

Benefits Enrollment Information

This policy detects benefit applications and other forms that include personal health information.

Diagnosis Information

This policy detects medical diagnosis information including mental, physical and addiction-related ailments.

Individually Identifiable Health Information (IIHI)

This policy detects individually identifiable information in conjunction with medical information related to patients, employees, or customers.

Medical Billings and Claims

This policy detects medical billing information and claims data including submissions to insurance companies, approvals and denials of payment, and continuing correspondence.

Medical History

This policy detects medical history information including diagnosis and prescription details.

Medical Record Numbers

This policy detects medical record numbers used in the identification and treatment of patients.

Medical Record Numbers - Threshold

This policy detects a specified amount (or threshold) of medical record numbers used in the identification and treatment of patients.

Personally Identifiable Information (PII) Policies

Account Number

This policy detects specific account numbers and/or account numbers that fall within a particular range. Numbers may be entered exactly or matched with a template.

Account Number - Threshold

This policy protects and controls a specified amount (or threshold) of specific account numbers and/or account numbers that fall within a particular range.

Account Number and Routing Information

This policy detects both an organization's account number(s) and the associated routing number(s).

Account Number with Additional PII

This policy detects specific account numbers and/or account numbers that fall within a particular range when accompanied by at least one or more pieces of identifying information such as name, address or DOB that could be used for identity theft.

Australian Medicare Card Number

This policy detects one or more Australian Medicare Card Numbers in various formats.

Australian State Drivers License

This policy detects one or more Australian State Drivers License Numbers in various formats.

Australian Tax File Number

This policy detects one or more Australian Tax File Numbers in standard format.

Background Checks

This policy detects background information checks, including private and often sensitive data that might be communicated inappropriately.

Canadian Social Insurance Number

This policy detects one or more Canadian Social Insurance Numbers in various formats.

Canadian Socal Insurance Number - Threshold

This policy detects a specified amount (or threshold) of Canadian Social Insurance Numbers in various formats.

Canadian Social Insurance Number with Additional PII

This policy detects one or more Canadian Social Insurance Numbers when accompanied by at least two pieces of identifying information such as name, address or DOB which could be used for identity theft.

Chinese Identity Card Number

This policy detects one or more Chinese Identity Card Numbers in standard format.

Credit Card Information

This policy detects credit card numbers in various ranges and formats.

Credit Card Information - Threshold

This policy detects a specified amount (or threshold) of credit card numbers in various ranges and formats.

Credit Report

This policy detects inappropriate distribution of credit reports or credit related data issued by consumer reporting agencies (CRAs).

Employee Evaluation Information

This policy is designed to identify employee evaluations, often regarded as private between an employee and an organization.

German Social Insurance Number

This policy detects one or more German National Pension Numbers in standard format.

Hong Kong Identity Card Number

This policy detects one or more Hong Kong Identity Card Numbers in standard format.

Indian Permanent Account Number

This policy detects one or more Indian Permanent Account Numbers in standard format.

Indonesian Identity Card Number (Nomor Induk Kependudukan)

This policy detects one or more Indonesian Identity Card Numbers in various formats.

Irish Personal Public Service Number

This policy detects one or more Irish Personal Public Service Numbers in standard format.

Italian National Identification Number

This policy detects one or more Italian National Identification Number in standard format.

Macau Non-Permanent Resident Identity Card (BIRNP)

This policy detects one or more Macau Non-Permanent Resident ID Numbers in standard format.

Macau Permanent Resident Identity Card (BIRP)

This policy detects one or more Macau Permanent Resident ID Numbers in standard format.

Malaysian National Registration Identification Card Number

This policy detects one or more Malaysian National Registration Numbers in standard format

Pakistan National Identity Card Number

This policy detects one or more Pakistan National Identity Card Numbers in standard format

Singapore National Registration Identity Card

This policy detects one or more Singapore National Registration Identity Card Numbers in standard format

Social Security Number

This policy detects one or more US Social Security Numbers in various formats.

Social Security Number - Threshold

This policy detects a specified amount (or threshold) of US Social Security Numbers in various formats.

Social Security Number with Additional PII

This policy detects one or more US Social Security Numbers when accompanied by at least one or more pieces of identifying information such as name, address or DOB that could be used for identity theft.

Taiwan Identity Card Number

This policy detects one or more Taiwan Identity Card Numbers in standard format.

Thailand Population Identification Code

This policy detects one or more Thailand Population Identification Codes in standard format.

UK Drivers License

This policy detects one or more UK Driving License Numbers in various formats.

UK Drivers License - Threshold

This policy detects a specified amount (or threshold) of UK Driving License Numbers.

UK Employee Compensation Information

Protect and control information related to the compensation of their UK employees to identity outside the organization, to a particular group (such as HR), or to a select circle of individuals that are allowed to receive and send such compensation information.

UK National Insurance Number

This policy detects one or more UK National Insurance numbers (the U.K. equivalent of the U.S. SSN), in various formats.

UK National Insurance Number - Threshold

This policy detects a specified amount (or threshold) of UK National Insurance Numbers in various formats.

UK National Insurance Number with Additional PII

This policy detects one or more UK National Insurance Numbers when accompanied by at least two pieces of additional identity information such as name, address or DOB that could be used for identity theft.

UK Tax Identification Number

This policy detects one or more UK Tax Identification Numbers in various formats.

UK Tax Identification Number - Threshold

This policy detects a specified amount (or threshold) of UK Tax Identification Numbers in various formats.

Unencrypted Wire Transfer Information

This policy assists organizations that want to be alerted to or prevent unencrypted disclosure of wire transfer information.

US Drivers License

This policy detects one or more US Drivers License Numbers in various formats.

US Drivers License - Threshold

This policy detects a specified amount (or threshold) of US Driver License Numbers in various formats.

US Employee Compensation Information

This policy detects information related to compensation for US employees being disclosed to parties outside the organization.

US Passport Number

This policy detects US Passport Numbers in various formats.

US Passport Number - Threshold

This policy detects specified amount (or threshold) of US Passport Numbers in various formats.

US Taxpayer Identification Number (TIN)

This policy detects US Taxpayer Identification Numbers in various formats.

US Taxpayer Identification Number (TIN) - Threshold

This policy detects a specified amount (or threshold) of US Taxpayer Identification Numbers in various formats.

Vietnam ID Card Number

This policy detects one or more Vietnam ID Card Numbers in standard format.

Security General / Corporate Policies

Audio Files

Sensitive information may be recorded and sent out of the organization. Protect and control the transmittal of audio media files.

E-mail to Personal Addresses

This policy identifies electronic communication with attachment(s) being sent to non-commercial domains (Hotmail, Yahoo, Gmail, and domains ending in .gov, .edu, .info, and so on), which immediately raises concerns as to whom the information is being distributed.

Forwarding Senior Management E-mail or Documents

This policy detects the forwarding of content originally sent by senior management.

Graphic and Image Files

This policy identifies graphic and image files in various formats.

Large Message or File Size

This policy identifies users sending messages over a certain size or files over a certain size.

Large Print Job Warning

This policy detects print jobs that exceed a specified number of pages and warns the user.

Network Security Threats

This policy identifies common hacking utilities and terms such as spoofing, buffer overflow tools, log wiping tools and password database cracking tools.

Password Protection/Encryption: Prohibition

This policy detects content that has been protected with a password or has been encrypted.

Random Sample

Regulators suggest that adding a targeting a reasonable percentage of messages for random review, in addition to normal lexicon-based reviews, is a prudent practice since such random reviews may discover issues not normally detected by ordinary means. This policy will randomly select messages, based on a percentage that is defined by the firm, to be automatically included in a reviewer’s queue.

Sharing of Usernames and Passwords

This policy detects the disclosure and sharing of passwords both inside and outside the organization.

Suspicious E-mail Behavior

This policy identifies electronic communication with blank subjects whose context suggests that the sender is attempting to avoid detection.

Transfer of Attachments - Threshold

This policy identifies electronic communication with a specified number (or threshold) of attachments, which could suggest a drive dump or other inappropriate bulk transfer of files.

Transfer of Personal E-mail File Folders

This policy identifies inappropriate bulk transfer of e-mail file folders which includes .PST and .NSF files.

Video Files

This policy identifies video media files in various formats.

User Defined Policies

By default, these policies are empty. They allow you to define your own policy criteria. You can use them to test your CA DataMinder setup. They also enable you to define your own custom policies if you have a particular requirement that is not fulfilled by the policy pack.

Common Content

For each user-defined policy, you can define the key text that you want CA DataMinder to detect.

Immediate Disqualifiers

Immediate Disqualifiers are words or phrases which immediately result in a non-match if CA DataMinder detects them in an email, file, or web page. That is, the email, file, or web page definitely does not match the policy criteria. CA DataMinder does not apply policy to these items even if they contain other sensitive words or phrases.

Note: A single word or phrase is sufficient to prevent CA DataMinder from applying policy.

Positive Indicators

Positive Indicators are preferred words or phrases. If CA DataMinder detects these words or phrases in an email, file, or web page, it increases the probability that the item matches the policy criteria.

A single Positive Indicator word or phrase is sufficient to trigger policy if no other excluding criteria are detected (such as excluded file names or URLs).

User Defined 1

This policy detects the key words and phrases that you specify in the Common Content settings.

You can define the severity, the policy action, and the resulting message that is seen by users. For the CA DataMinder Enterprise Edition, you can also assign smart tags to classify any captured items.

- For email policies, you can specify excluded search text (CA DataMinder ignores emails containing these words or phrases). You can also specify **included** sender addresses (CA DataMinder only applies this policy to emails from these senders).
- For Files In Motion and Data At Rest policies, you can specify excluded file names (CA DataMinder ignores these files) and a minimum file size (CA DataMinder ignores any smaller files).
- For Web policies, you can specify excluded URLs (CA DataMinder does not apply policy to these web pages).

User Defined 2

This policy detects a second set of key words and phrases. It provides the same settings as the *User Defined 1* policy, with one difference for email policies.

For email policies, you can specify **excluded** sender addresses (CA DataMinder ignore emails from these senders).

User Defined 3

This policy detects a third set of key words and phrases. It provides the same settings as the *User Defined 1* policy, with one difference for email policies.

For email policies, you can specify **excluded** sender addresses (CA DataMinder ignore emails from these senders).

Who Do the Standard Policies Apply To?

The iConsole standard policies only apply to members of the FPP Custom group or its subgroups.

FPP User Groups Created Automatically on the CMS

When you install the iConsole standard policies, the following user groups are created automatically below the top-level Users group:

FPP Base Group

This user group functions as a receptacle for the default FPP policies. It does not contain any user accounts.

FPP Custom Group

The policy for this group gets updated when you edit the policies in the iConsole. This group contains new user accounts created automatically when you install CA DataMinder endpoint agents. This group also contains various accounts used by CA DataMinder to apply policy to file events and unrecognized email senders.

User Accounts in FPP Custom Group

The following user accounts are added, or moved to, the \FPP Custom Group folder:

New users

After installing the iConsole standard policies, any new CA DataMinder user accounts created when you deploy CA DataMinder endpoint agents are added to the FPP Custom Group.

Note: Any CA DataMinder user accounts that already existed before you installed the iConsole standard policies are **not** moved to the FPP Custom Group. These user accounts stay in their existing users groups and are not governed by the iConsole standard policies. The standard policies only apply to users in the FP Custom Group and its subgroups.

UnknownInternalSender

This account has no management or administrative privileges; its sole purpose is as a conduit to enable policy engines to apply policy to internal emails from unrecognized senders.

When you install a CMS, the Unknown Internal Sender setting in the machine policy defaults to this UnknownInternalSender user account

DefaultFileUser

This account has no management or administrative privileges; its sole purpose is as a conduit to enable policy engines to apply policy to scanned, captured or imported files if no other means are available to determine the policy participant.

When you install a CMS, the Default Policy for Files setting in the machine policy defaults to this DefaultFileUser user account.

DefaultClientFileUser

This account is similar to the DefaultFileUser account. The account is used solely by the Client File System Agent (CFSA) when scanning local workstations. The CFSA uses this account to apply the same policy to scanned files across all workstations.

When you install a CMS, the Default Policy for Data At Rest setting in the machine policy defaults to this DefaultClientFileUser user account.

Install the Standard Policies

To roll out the iConsole standard policies, you must repeat the installation on your CMS, your iConsole application servers and your iConsole front-end Web servers.

To install on the CMS

Before you start, make sure that your database engine is properly configured.

1. Launch the CA DataMinder installation wizard. To do this, run setup.exe in your CA DataMinder distribution image.
2. In the Installation Type screen, choose Advanced Installation.
3. In the Advanced Install Options screen, choose iConsole Standard Policies and then click Install.

   This launches the Standard Policies installation wizard in a separate window.
4. In the Standard Policies installation wizard, navigate to the final screen and click Install to start the file transfer.

To install on iConsole application servers

Follow steps 1 through 4 of the CMS instructions above.

To install on iConsole front-end Web servers

1. Follow steps 1 through 3 of the CMS instructions above.
2. In the iConsole Application Server Details screen, if no iConsole application server is installed locally you will need to specify which iConsole application server to connect to.

   Server

   Specify the name or IP address of the machine hosting the application server. Type localhost to specify the local machine.

   Port

   Specify the TCP port used for communication between front-end Web server and the application server. This defaults to port 80. If you specify a non-default port, you must also update the application server to use the same port.

   Use SSL

   If you intend to use SSL to communicate over a secure port (for example, 443), select this check box. This ensures that the port used for communication between the front-end Web server and the application server will use SSL.
3. Navigate to the final screen and click Install to start the file transfer.

Advanced Dashboard Configuration

You can make various advanced configuration changes to the dashboard. For example, you can:

- Change the dashboard snapshot periods (snapshots refer to total incident counts for the whole snapshot period).
- Determine which events are included in the Reviewed Events and Escalated Events metrics by specifying the required audit values.
- Limit the policies or policy classes represented in dashboard charts.

More information:

Configure the Dashboard on SQL Server CMSs (see page 199)
Configure the Dashboard on Oracle CMSs (see page 200)
Dashboard Configuration Parameters (see page 200)
Dashboard Event Totals Seem Wrong After Drilling into Report or Chart (see page 205)

Configure the Dashboard on SQL Server CMSs

To configure the dashboard on SQL Server CMSs, run a database statement to change parameters to new values.

Syntax

In SQL Server Enterprise Manager Studio, run the following database statements.

```
EXEC rut_dlp_set_aggregation_values @<parameter>[,@<parameter>]
```

Where @<parameter> defines the new parameter value.

Each query can include multiple parameter definitions, separated by commas. If a parameter is not explicitly specified it retains its current value. Available parameters are described in the following sections.

Example

The following example specifies a 3 month dashboard snapshot period. Snapshots refer to total incident counts for the whole snapshot period. Snapshots are calculated for the following metrics: Reviewed Events, New Events, and Escalated Events:

```
EXEC rut_dlp_set_aggregation_values @history=3
```

More information:

Dashboard Configuration Parameters (see page 200)

Configure the Dashboard on Oracle CMSs

To configure the dashboard on Oracle CMSs, run commands using a utility such as Oracle SQL Plus to change parameters to new values.

Syntax

Run this command:

```
begin dlp_agg.
  rut_dlp_set_aggregation_values (<parameter>[,<parameter>]);
  commit; end;
```

Where <parameter> defines the new value for an dashboard parameter.

Each query can include multiple parameter definitions, separated by commas. If a parameter is not explicitly specified, it retains its current value. Available parameters are described in the following sections.

Example

The following example specifies a 3 month dashboard snapshot period. Snapshots refer to total incident counts for the whole snapshot period. Snapshots are calculated for the following metrics: Reviewed Events, New Events, and Escalated Events:

```
begin dlp_agg.
  rut_dlp_set_aggregation_values
  (history=3);
  commit; end;
```

More information:

Dashboard Configuration Parameters (see page 200)

Dashboard Configuration Parameters

You can customize the following dashboard parameters.

More information:

History (see page 201)
AF1_Reviewed (see page 202)
AF1_Escalated (see page 203)
Root Policy Node (see page 204)

History

Use this parameter to specify the dashboard snapshot periods (in months). Snapshots refer to total incident counts for the whole snapshot period. Snapshots are calculated for the following metrics: Reviewed Events, New Events, and Escalated Events.

Syntax

SQL Server

```
@history=<n>
```

Oracle

```
history=><n>
```

Examples

If <n> is set to 3, a data warehousing job that runs on 20 June will calculate snapshot totals based on all qualifying incidents with a capture timestamp of 20 March or later:

```
@history=3
history=>3
```

Notes

Reviewers viewing the iConsole dashboard must be aware that snapshot data relates to a specific point in time. In particular, snapshots are calculated when the last data warehousing job ran. Snapshots are not calculated in real time when the reviewer views or refreshes the dashboard.

AF1_Reviewed

There are three related dashboard parameters:

```
AF1_reviewed
AF2_reviewed
AF3_reviewed
```

Use these parameters to determine which events are included in the Reviewed Events metric.

By default, these parameters are set to NULL. This means event gets added to the Reviewed Events metric if any audit value is set in Field 1, Field 2 or Field 3. To narrow the Reviewed Events metric definition, set these parameters to any combination of supported audit values.

Syntax

AF1, AF2 and AF3 refer respectively to audit fields Field 1, Field 2 and Field 3.

SQL Server

```
@AF1_reviewed=NULL or <Audit value>
@AF2_reviewed=NULL or <Audit value>
@AF3_reviewed=NULL or <Audit value>
```

Oracle

```
AF1_reviewed=>NULL or <Audit value>
AF2_reviewed=>NULL or <Audit value>
AF3_reviewed=>NULL or <Audit value>
```

Enclose the <Audit value> in single quotes.

Values are not case-sensitive.

The data warehousing job uses a LIKE operator when querying the values in an event's audit fields.

Examples

The following examples match events with the audit status of 'Escalate'.

```
@AF1_reviewed='esc'
AF1_reviewed=>'esc'
```

Notes

Field 1, Field 2, and Field 3 are configurable in the Administration console. Click Tools, Audit Options, General tab to edit these fields. Field 1 is usually set to 'Audit Status'. For details, see the CA DataMinder Administration console online help; search for 'event auditing: customizing the audit features'.

Reviewers use these fields to update audit details for captured events in the iConsole.

AF1_Escalated

There are three related dashboard parameters:

```
AF1_escalated
AF2_escalated
AF3_escalated
```

Use these parameters to determine which events are included in the Escalated Events metric.

By default, AF1 is set to 'escalate' while AF2 and AF3 are set to NULL. This means event gets added to the Escalated Events metric only if they have an any audit status of 'Escalate'. To broaden the Escalated Events metric definition, set these parameters to any combination of supported audit values.

Syntax

AF1, AF2 and AF3 refer respectively to audit fields Field 1, Field 2 and Field 3.

SQL Server

```
@AF1_escalated=NULL or <Audit value>
@AF2_escalated=NULL or <Audit value>
@AF3_escalated=NULL or <Audit value>
```

Oracle

```
AF1_escalated=>NULL or <Audit value>
AF2_escalated=>NULL or <Audit value>
AF3_escalated=>NULL or <Audit value>
```

Enclose the <Audit value> in single quotes.

Values are not case-sensitive.

The data warehousing job uses a LIKE operator when querying the values in an event's audit fields.

Examples

The following examples match events with AF2 values of 'Escalate to Management' or 'Escalate to HR', and so on.

```
@AF2_escalated='esc'
AF2_escalated=>'esc'
```

Notes

Field 1, Field 2, and Field 3 are configurable in the Administration console. Click Tools, Audit Options, General tab to edit these fields. Field 1 is usually set to 'Audit Status'. For details, see the CA DataMinder Administration console online help; search for 'event auditing: customizing the audit features'.

Reviewers use these fields to update audit details for captured events in the iConsole.

Root Policy Node

Use this parameter to limit the policies or policy classes represented in dashboard charts. By default, dashboard charts include all incidents in the data warehouse. However, you can use this parameter to restrict policy charts to include only incidents associated with specific policies or policy classes.

Syntax

This parameter specifies a list of integer ClassUID values (in decimal) for policy classes as listed in the Wgn3ClassificationNode database tables. All policies in the subtree of any specified ClassUID node are included in dashboard charts.

By default, this parameter is set to 3000000 and 4000000. These ClassUID values represent the root nodes for predefined and custom policies respectively. If necessary, you can set include this parameter and explicitly set it to different ClassUID values.

SQL Server

```
@root_policy_node_list='<n>[,<n>]'
```

Oracle

```
root_policy_node_list=>'<n>[,<n>]'
```

Note: If you specify a custom root policy node but later want to revert to the default ClassUID values, set this parameter explicitly to 3000000 and 400000. Simply omitting this parameter from the data warehousing job does *not* reinstate the default root policy nodes.

Examples

To display only incidents captured by Compliance policies (ClassUID 3020000), set this parameter to:

```
@root_policy_node_list='3020000'
root_policy_node_list='3020000'
```

To display only incidents captured by Personally Identifiable Information (ClassUID 3010300) or Personal Health Information (ClassUID 3010400) policies, set this parameter to:

```
@root_policy_node_list='3010400,3010300'
root_policy_node_list='3010400,3010300'
```

About Policy Classes

CA DataMinder contains several *policy classes* (such as Non Public Information). In turn, each policy class is implemented though various *policies* (such as Sales Information). For categorization purposes, you can associate individual triggers with a particular policy class. This information gets stored with an event's metadata when the trigger activates. All triggers support policy classes. You can then use the iConsole to search for events associated with a policy class.

Find details of how CA DataMinder stores policy classes as ClassUID values in the 'Policy Classification Nodes' section under the Wgn3ClassficationNode table description in the *Database Schema and Views Reference Guide*.

Dashboard Event Totals Seem Wrong After Drilling into Report or Chart

In certain conditions, if a reviewer drills down into a report or dashboard chart, the number of results does not match the event total shown in the report or chart. This apparent disparity can occur if further events have been captured or reviewed in the intervening period since the snapshot totals were last calculated.

Snapshot totals are recalculated each time a data warehousing job runs. By default, these jobs run every hour. Consequently, snapshots (such as 'total unreviewed events') reflect the total number of events *at the time when the job ran*.

If the number of events (or incidents) in the CMS database rises or falls before the next data warehousing job runs (for example, because a manager reviews some previously unreviewed events), then the snapshot total shown in the report or dashboard will no longer tally with the actual number of underlying events in the CMS database. If a reviewer were to drill down into the report or dashboard at this point, they would see an apparent disparity in the number of events.

Example

Consider this timeline:

15.00 PM	A data warehousing jobs finds 100 unreviewed events and adds them to the data warehouse.
15.15 PM	A manager refreshes their dashboard. The snapshot total for Unreviewed Events is 100.
15.16 PM	The same manager drills down into the dashboard to see the underlying events. The Search Results screen does indeed find 100 unreviewed events.
15.30 PM	A reviewer audits 25 of the unreviewed events in the iConsole.
15.45 PM	The manager refreshes their dashboard; the snapshot total for Unreviewed Events is still 100. This is because there has been no new data warehousing job since 15.00 PM.
15.46 PM	The manager drills down into the dashboard again. But this time, the Search Results screen only finds 75 unreviewed events!
16.00 PM	The next data warehousing job runs and finds 75 unreviewed events. The snapshot total and number of underlying events are back in sync.

Note: Such potential disparities only affect snapshots of event counts based on audit status. They cannot occur with snapshots based on non-changing event attributes such as events counts by policy.

Chapter 9: Data Warehouse

This section contains the following topics:

About the Data Warehouse

The Data Warehouse is a set of database tables containing CA DataMinder event data that has been transformed into a format suitable for generating reports and iConsole dashboards.

For the current CA DataMinder release, the Data Warehouse tables are installed in the CA DataMinder Central Management Server database.

Data Warehouse Requirements

This section lists the general requirements and considerations for the CA DataMinder data warehouse.

CMS Database Accounts

You can enable the data warehouse when you install the CMS. When you enable the data warehouse, you must supply credentials for two accounts for the CMS database.

Reporting User

External reporting applications (such as BusinessObjects Enterprise) use this database account to connect to the Data Warehouse and CMS database.

This database account inherits the security model of the CA DataMinder user who is running the report. For example, if the user running the report has been assigned to the Management Group security model, then the report results are also subject to RLS restrictions based on the user’s management group. Conversely, if the user has been assigned to the Unrestricted security model, the report results are not subject to any RLS restrictions.

Unrestricted Search User

This database account corresponds to the 'Unrestricted' security model. CA DataMinder consoles and external reporting tools can use this database account when searching the CA DataMinder Data Warehouse and CMS database for events. Unlike normal Search User database accounts, the Unrestricted Search User is *not* subject to row level security (RLS) when searching the database. If a reviewer has 'Unrestricted' security model, the reviewer can see any events when they run a search or report. Search results or reports are not restricted by policy class or the reviewer’s management group.

This account is useful if, for example, an external auditor requires unrestricted (view-only) access to captured events in your CMS database.

Size Considerations

Installing the data warehouse tables can increase the size of the existing CMS database by up to 50%. Therefore, ensure you have sufficient disk space for the volume of data that you want to keep in the data warehouse. To ensure sufficient disk space, you may need to modify the default settings for data warehousing jobs to prevent the data warehouse growing excessively.

Competition for Memory, CPU and Disk Resources on the CMS Host Server

The data warehouse tables are stored in the CMS database. Consequently, any memory, CPU and disk resources consumed by data warehouse queries are not available to the CMS database and can result in performance issues. For example, large sort or join operations and increased buffer cache requirements can exert pressure on memory resources. Likewise, insufficient CPU threads can exert pressure on the CPU and overloaded disks can increase disk response times.

We therefore recommend that you set up low impact monitoring of these resources on the CMS host server to detect any performance issues.

Also, configuration changes on the host server may ease some performance issues. For example, you can add more memory or increase the number of CPUs. You can also reduce CPU parallelism. You can uninstall other applications from the host server. You can spread the disk workload or isolate different types of disk activity.

Whatever configuration changes you make, we recommend that investigate the performance issue first to identify the genuine cause. Sometimes, the symptoms of a problem may obscure the underlying cause. For example, a memory shortage may result in a very small buffer cache, in turn causing an excessive number of physical disks reads.

Do I Need To Collect Event Participant Data?

This section summarizes the data warehouse requirements if you intend to run BusinessObjects reports for CA DataMinder.

Collect Event Participant Data

Important! You *must* collect event participant data if you intend to run BusinessObjects reports!

This data enables the data warehouse to associate captured events with event participants. For example, it can associate emails with the sender and recipients.

You can configure the data warehouse to collect event participant data when you install your CMS. Or you can retrospectively configure the data warehouse to collect this data.

If you already use the iConsole dashboard, your Data Warehouse already contains event and audit data. However, it does not contain event participant data. You must therefore resynchronize, or empty and repopulate, the data warehouse.

More information:

Configure Support for BusinessObjects Reports (see page 217)

Enable the Data Warehouse at Install Time

The Data Warehouse is installed automatically when you install a new CA DataMinder CMS, but you must explicitly enable the data warehouse if you want to use the iConsole dashboard or run BusinessObjects reports for CA DataMinder.

To enable the Data Warehouse

1. Follow the standard instructions for installing a CA DataMinder CMS and navigate the installer screens to the Data Warehouse Configuration screen.

   See the reference below for details about installing a CMS.

2. Fill in the following fields. Then click Next.

   Enable data warehousing for this CMS

   Select this check box enable the data warehouse.

   Collect event participant data

   Select this check box to collect event participant data.

   Important! You must collect event participant data if you intend to run BusinessObjects reports.

3. In the Data Warehouse Database Account screen, define the following database accounts.

   In all cases, click the ... button to specify the account credentials. In the resulting User Credentials dialog, specify the username and password for the database account. If this account is a new account, select the Create User check box.

   Data Warehouse User

   External reporting applications (such as BusinessObjects Enterprise) use this database account to connect to the Data Warehouse and CMS database.

   Unrestricted Search User

   This database account corresponds to the 'Unrestricted' security model. CA DataMinder consoles and external reporting tools can use this database account when searching the CA DataMinder Data Warehouse and CMS database for events. Unlike normal Search User database accounts, the Unrestricted Search User is *not* subject to row level security (RLS) when searching the database. If a reviewer has 'Unrestricted' security model, the reviewer can see any events when they run a search or report. Search results or reports are not restricted by policy class or the reviewer’s management group.

Database Administrator User

If either of the database accounts specified above are new, specify the Database Administrator User that the installation wizard can use to log in to SQL Server or Oracle to create these new accounts.

For Oracle databases, this Database Administrator account *must* have the following system privileges:

```
CREATE SESSION
RESOURCE
DBA
SYSDBA
```

4. Continue to the final wizard screen and click Install.
5. (Only applicable to SQL Server Express CMSs) If your data warehouse is hosted in a SQL Server Express database, you must manually enable and schedule the processing jobs that populate the data warehouse.

More information:

Do I Need To Collect Event Participant Data? (see page 211)

Do I Need To Collect Event Participant Data?

This section summarizes the issues to consider when deciding whether to enable or disable the Event Participant Fact table in the data warehouse.

The Data Warehouse installer includes a 'Collect event participant data' check box. This check box determines whether the Event Participant Fact table is enabled or disabled. This table in the Data Warehouse associates events with participants.

If the Event Participant Fact table is enabled

This table has the potential to grow very large. Our testing indicates that the table can increase the size of the CMS database by 30-40%. The reason is because an individual event can have many participants.

However, enabling the Event Participant Fact table has one important advantage. Namely, users assigned to a Management Group security model can successfully run BusinessObjects reports for CA DataMinder.

If the Event Participant Fact table is disabled

The BusinessObjects reports for CA DataMinder do not support filtering by user or user group. In turn, the reports do not support the Management Group security model. (Under this model, reviewers can only view events where at least one participant was in their management group when the event was captured.) Therefore, if users assigned to a Management Group security model run a BusinessObjects report for CA DataMinder, the report returns zero results.

Users assigned to a Policy security model can run these reports, but they cannot filter the reports by user or user group (even if users and groups are shown as configurable report parameters).

Users who only need the iConsole dashboard (but not the BusinessObjects reports for CA DataMinder) can leave the Event Participant Fact table disabled. The dashboard does not use data in this table.

Note: If you already use the iConsole dashboard but now want to start running BusinessObjects reports for CA DataMinder, you must enable *and populate* the Event Participant Fact. See the reference below for details.

Can I subsequently enable or disable the Event Participant Fact table?

Yes. You can manually enable or disable the Event Participant Fact table by running wgninfra.exe commands. See the following section for details.

Note: If you want to start running BusinessObjects reports for CA DataMinder, you must enable *and populate* the Event Participant Fact. It is not sufficient to simply enable this table. See the reference below for details.

Data Warehouse is Populated Automatically

BusinessObjects reports for CA DataMinder and iConsole dashboards are based on event data in the data warehouse. This event data has been extracted and transformed from source tables in the CMS database.

When the data warehouse is enabled, CA DataMinder automatically schedules jobs to populate the data warehouse with event data. By default, these data warehousing jobs run during a daily off-peak processing window. Depending on the system configuration, typical processing speeds can be up to a 200,000 events in 5 minutes.

- On CMSs with less than 100,000 events, CA DataMinder processes *all events* when the scheduled job first runs.
- On CMSs with over 100,000 events, CA DataMinder prioritizes events captured in the last 3 days. CA DataMinder process these events when the scheduled job first runs. It then gradually processes the remaining older events whenever a scheduled job runs.

By default, the off-peak processing window starts at midnight and lasts for 300 minutes. If required, you can change the default job parameters and job frequency. See the 'Advanced Configuration' section for details.

Advanced Configuration

More information:

Retrospectively Enable the Data Warehouse (see page 213)
Enable the Data Warehouse on SQL Server Express CMSs (see page 218)
Reconfigure the Data Warehouse (see page 220)
Change the Job Frequency (see page 220)
Change the Purge Schedule (see page 223)

Retrospectively Enable the Data Warehouse

(Applies only if you did not enable the data warehouse when you installed the CMS.)

A data warehouse is installed automatically when you install a CMS. But if you chose not to enable the data warehouse when you installed the CMS, the data warehouse remains disabled until you enable it.

How to Retrospectively Enable the Data Warehouse

1. (Optional) Set credentials for the Reporting User (see page 207) database account.

   You only need this database account if you intend to run BusinessObjects reports.

2. (Optional) Set credentials for the Unrestricted Search User (see page 207) database account.

   You only need this database account if you intend to run BusinessObjects reports.

3. Enable the Data Warehouse in the Administration console.
4. (Optional) Configure the Data Warehouse to support BusinessObjects reports.

More information:

Set Credentials for the Reporting User (see page 214)
Set Credentials for the Unrestricted Search User (see page 215)
Enable the Data Warehouse After Installing (see page 216)
Configure Support for BusinessObjects Reports (see page 217)

Set Credentials for the Reporting User

You must specify a Reporting User database account if you enable data warehousing. External reporting applications (such as BusinessObjects Enterprise) use this database account to connect to the Data Warehouse and CMS database.

You can use the Administration console to add or modify credentials for the Reporting User database account. For example, if the password has been changed on the database server (for example, for security reasons), you can supply CA DataMinder with the new password.

To set credentials for the Reporting User

1. Log on to the Administration console using an account that has the 'Admin: Manage security models' privilege.
2. Click Tools, Set Reporting User Credentials.
3. Enter the user name and password in the Set Reporting User Credentials dialog.
4. (Optional) If necessary, provide credentials for an existing Database Administrator account. See below for details.

   For Oracle CMS databases, this Database Administrator account *must* have the following system privileges:

   ```
   CREATE SESSION
   RESOURCE
   DBA
   SYSDBA
   ```

When must I provide Database Administrator details?

Credentials for the Reporting User are securely stored in the CMS database and in the CMS internal file system. The two sets of credentials must be in sync.

You do *not* need to provide Database Administrator details if a DBA has already updated the Reporting User credentials in the CMS database. In this situation, CA DataMinder only needs to update the Reporting User credentials stored in the CMS internal file system.

You *do* need to provide Database Administrator details if the CMS database has not been updated yet. In this situation, CA DataMinder simultaneously adds the Reporting User credentials to the CMS internal file system and the CMS database. CA DataMinder uses the Database Administrator account to log in to SQL Server or Oracle and update the CMS database.

Set Credentials for the Unrestricted Search User

Before you enable the Data Warehouse, you may need to specify the Unrestricted Search User database account.

This database account corresponds to the 'Unrestricted' security model. CA DataMinder consoles and external reporting tools can use this database account when searching the CA DataMinder Data Warehouse and CMS database for events. Unlike normal Search User database accounts, the Unrestricted Search User is *not* subject to row level security (RLS) when searching the database. If a reviewer has 'Unrestricted' security model, the reviewer can see any events when they run a search or report. Search results or reports are not restricted by policy class or the reviewer’s management group.

You can use the Administration console to add or modify credentials for the Unrestricted Search User database account.

To set credentials for the Unrestricted Search User

1. Log on to the Administration console using an account that has the 'Admin: Manage security models' privilege.
2. Click Tools, Manage Security Models.
3. In the Manage Security Models dialog, select the Unrestricted model and click Modify.
4. In the Modify Security Model dialog, click Set Credentials.
5. In the Set Model Credentials dialog, enter the user name and password.
6. (Optional) If necessary, provide credentials for an existing Database Administrator account. See below for details.

   For Oracle CMS databases, this Database Administrator account *must* have the following system privileges:

   ```
   CREATE SESSION
   RESOURCE
   DBA
   SYSDBA
   ```

When must I provide Database Administrator details?

Credentials for the Unrestricted Search User are securely stored in the CMS database and in the CMS internal file system. The two sets of credentials must be in sync.

You do *not* need to provide Database Administrator details if a DBA has already updated the Unrestricted Search User credentials in the CMS database. In this situation, CA DataMinder only needs to update the Unrestricted Search credentials stored in the CMS internal file system.

You *do* need to provide Database Administrator details if the CMS database has not been updated yet. In this situation, CA DataMinder simultaneously adds the Unrestricted Search User credentials to the CMS internal file system and the CMS database. CA DataMinder uses the Database Administrator account to log in to SQL Server or Oracle and update the CMS database.

Enable the Data Warehouse After Installing

The Data Warehouse is installed automatically when you install a new CA DataMinder CMS, but you must explicitly enable the data warehouse if you want to use the iConsole dashboard or run BusinessObjects reports for CA DataMinder.

To enable the Data Warehouse

1. Log on to the Administration console using an account that has the 'Admin: Manage security models' privilege.
2. Click Tools, Configure Data Warehouse.
3. In the General Options section, select the 'Enable Data Warehouse population' check box.
4. (Optional) Configure other Data Warehouse settings (see page 220) as required.

Configure Support for BusinessObjects Reports

(Applies only if you want to run BusinessObjects reports for CA DataMinder.)

BusinessObjects reports for CA DataMinder show results by user and group. These reports therefore require event participant data. Before reviewers can run these reports, you must populate the Data Warehouse with this data.

To configure the Data Warehouse to support BusinessObjects reports

1. Log on to the Administration console using an account that has the 'Admin: Manage security models' privilege.
2. Click Tools, Configure Data Warehouse.

   The Configure Data Warehouse dialog displays.
3. In the General Options section, select the 'Collect event particpant data' check box.
4. (Applies only if you already use the iConsole dashboard). Resynchronize, or empty and repopulate, the Data Warehouse.

   If you already use the iConsole dashboard, your Data Warehouse already contains event and audit data. However, it does not contain event participant data. You must now add this data to the Data Warehouse. Do *one* of the following:

   - Go to the Advanced Options section and select the 'Resynchronize Data Warehouse data on next run' check box.

     Use this method if you have never purged events from the CMS database. Events in the CMS already correspond with events in the data warehouse. This operation is relatively fast.

   - Go to the Advanced Options section and select the 'Purge all Data Warehouse data and repopulate on next run' check box.

     Use this method if you regularly purge events from the CMS database. The Data Warehouse probably contains data for events that no longer exist in the CMS. You must eliminate this discrepancy before you run BusinessObjects reports. Specifically, you must empty and then repopulate the entire Data Warehouse so that it only contains events that currently exist in the CMS database. This operation takes longer than a resync.

Enable the Data Warehouse on SQL Server Express CMSs

If your data warehouse is hosted in a SQL Server Express database, you must manually enable and schedule the processing jobs that populate the data warehouse.

To enable the data warehouse on SQL Server Express CMSs

1. Create a SQLCMD script to run data warehousing jobs.

   Specifically, save the following SQL Server query to a batch file, such as Data_Warehouse_Job.cmd:

   ```
   sqlcmd -E –S "<host>\<instance>" -d "WGN_<server>" -Q "EXEC
   rut_dlp_aggregation_process"
   ```

   Where:

   -E

   > Specifies a Windows-trusted database connection.
   >
   > **Note:** Verify that the scheduled task runs as a Windows user who has a SQL Server login that has been granted the db_owner privilege on the WGN_<server> database.

   -S "<host>\<instance>"

   > Specifies the host server and a named (non-default) instance of SQL Server. For example:
   >
   > ```
   > -S "localhost\SQLEXPRESS"
   > -S "UXW2K8DLP1\SQLEXPRESS"
   > ```
   >
   > If you omit the '-S' parameter, the script uses the default instance of SQL Server on the localhost.
   >
   > To specify the default instance of SQL Server on a host server, you only need to specify "<host>".
   >
   > To specify the host server and the instance of SQL server using a static port number, specify the host server (without an instance name) and port number, separated with a comma:
   >
   > ```
   > "<host>,CA Portal"
   > ```
   >
   > For example:
   >
   > ```
   > -S "UXW2K8DLP1,1629"
   > ```
   >
   > **Note:** You may need to specify a static port if your organization disables the SQL browser service for security reasons.

   -d "WGN_<server>"

   > Defines the database name. This database name incorporates the name of your CMS host server. For example:
   >
   > ```
   > -d "WGN_UXW2K8DLP1"
   > ```

   -Q

Specifies the database SQL command that initiates the data warehouse processing job:

```
"EXEC rut_dlp_aggregation_process"
```

2. Set up a scheduled data warehousing aggregation task using the Windows Task Scheduler:

    a. Create a task to run the data warehousing batch file that you previously created.

    b. Configure the task to run as the Windows user account that corresponds to the appropriate SQL Server login.

    c. Specify a task schedule as required. We recommend that the data warehousing job runs hourly.

More information:

Change the Job Frequency (see page 220)

Reconfigure the Data Warehouse

You can reconfigure the Data Warehouse at any time. For example, you may want to change the settings for off-peak processing jobs or data purges.

To reconfigure the Data Warehouse

1. Log on to the Administration console using an account that has the 'Admin: Manage security models' privilege.
2. Click Tools, Configure Data Warehouse.
3. Configure the following settings:

   General Options

   > These settings enable or disable the Data Warehouse. They also specify which data gets copied into the Data Warehouse. In particular, they specify whether to include event participant data. Other settings enable regular purges of older data from the Data Warehouse.
   >
   > **Important!** You must collect event participant data if you intend to run BusinessObjects reports..

   Additional Population and Maintenance

   > These settings configure off-peak processing jobs for the Data Warehouse.
   >
   > If purging is enabled, the purges are performed by the off-peak processing job.
   >
   > By default, these jobs run at midnight for 300 minutes, but you can reschedule them. Be aware that the data processing associated with these data warehousing jobs can generate a heavy workload on the CMS. We strongly recommend that you run these jobs during offpeak times.

   Advanced Options

   > These settings configure batch sizes for data warehousing jobs. Other settings enable you to resynchronize the Data Warehouse with data in the CMS database or to purge and repopulate the entire Data Warehouse.

   Note: For details about the available settings, see the online help.

Change the Job Frequency

By default, data warehousing jobs run every hour. But if necessary, you can change the job frequency.

Change the Job Frequency on SQL Server CMSs

Note: If you change the job frequency, you must be logged onto the CMS database as the primary user. This is the account that CA DataMinder uses to access the CMS database.

SQL Server CMSs

Edit the following SQL Server job:

```
DLP_Aggregation_<CMS database>
```

Where <CMS database> is the name of CMS database.

SQL Server Express CMSs

Set up a scheduled data warehousing aggregation task using the Windows Task Scheduler:

a. Create a task to run the data warehousing batch file that you previously created.

b. Configure the task to run as the Windows user account that corresponds to the appropriate SQL Server login.

c. Specify a task schedule as required. We recommend that the data warehousing job runs hourly.

Change the Job Frequency on Oracle CMSs

Note: If you change the job frequency, you must be logged onto the CMS database as the primary user. This is the account that CA DataMinder uses to access the CMS database.

You can specify a repeat interval for a specified data warehousing job. First, you need to check the job name and confirm its next run time. Then you set the job frequency.

Syntax

Check the Job Name

Run the following SQL query:

```
SELECT OWNER, JOB_NAME, SUBSTR(REPEAT_INTERVAL,1,40)
AS REPEAT_INTERVAL,
NEXT_RUN_DATE
FROM ALL_SCHEDULER_JOBS;
```

This query returns the JOB_NAME value.

Set the Job Frequency

Run the following PL/SQL command. This example sets the job to run at 2:00 AM every night:

```
BEGIN DBMS_SCHEDULER.SET_ATTRIBUTE
(name => 'DLP_AGGREGATION_<DBNAME>',
attribute => 'repeat_interval',
value => 'FREQ=<Period>; <Frequency>' );
END;
```

Where:

DLP_AGGREGATION_<DBNAME> is the JOB_NAME returned from the first SQL query.

<Period> and <Frequency> are DBMS_SCHEDULER parameters that define the job schedule.

Examples

The following example sets the job to run at 2:00 AM every night:

```
BEGIN DBMS_SCHEDULER.SET_ATTRIBUTE
(name => 'DLP_AGGREGATION_<DBNAME>',
attribute => 'repeat_interval',
value => 'FREQ=DAILY; BYHOUR=2' );
END;
```

The following example sets the job to run every 3 hours:

```
BEGIN DBMS_SCHEDULER.SET_ATTRIBUTE
(name => 'DLP_AGGREGATION_<DBNAME>',
attribute => 'repeat_interval',
value => 'FREQ=HOURLY; interval=3' );
END;
```

See your Oracle documentation for full DBMS_SCHEDULER details.

Change the Purge Schedule

If you purge large amounts of data from the data warehouse, this generates a heavy workload on the CMS. We therefore recommend that you run data at off-peak times.

Purges are performed by off-peak processing jobs. By default, these jobs run from midnight until 5am. To change when purges run, you must reschdeule the off-peak processing jobs. Settings in the 'Additional Population and Maintenance' section determine when these jobs run.

More information:

Reconfigure the Data Warehouse (see page 220)

Chapter 10: Quarantine Manager

This section contains the following topics:

Quarantine Manger Overview

Regulatory requirements require that certain categories of documents sent to multiple external recipients must be approved by an appropriate representative. The CA DataMinder quarantine feature enables your organization to enforce this requirement.

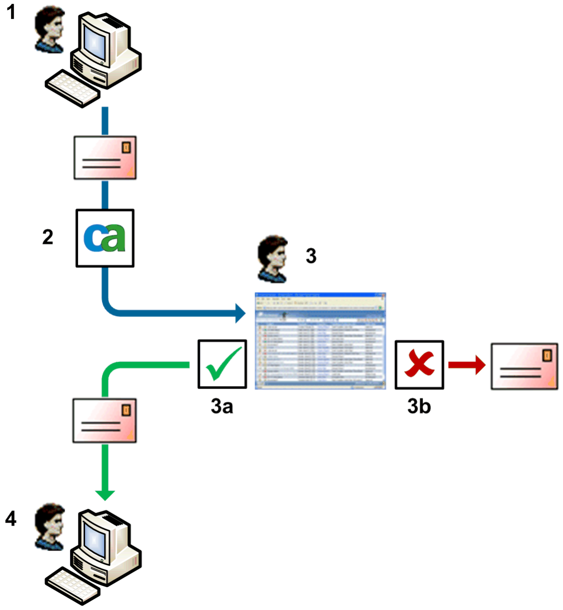

CA DataMinder and quarantined emails

1 A user sends an email containing a potential non-compliance issue.

2 CA DataMinder intercepts and quarantines the email.

3 A reviewer determines whether to release the email from quarantine (**3a**) or to reject it (**3b**).

4 Finally, released emails are forwarded to the intended recipient.

Quarantine Manager Architecture

The role of the Quarantine Manager is to ensure that emails released from quarantine are sent on to their original recipients. To achieve this, the Quarantine Manager regularly queries the CMS for released or timed-out emails and forwards these to their intended recipients. The architecture is summarized below:

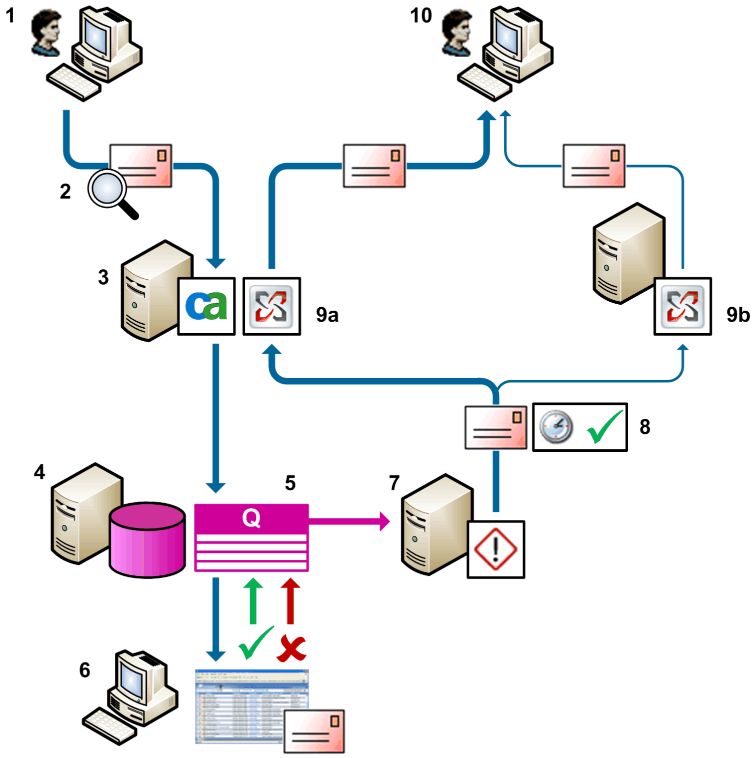

Quarantine procedure: Example based on Exchange server integration

This example shows how the Quarantine feature operates in conjunction with Exchange server integration. However, it can also operate in conjunction with Outlook and Notes endpoint agents.

An email is sent (**1**) and monitored by CA DataMinder (**2**) as it transits through the Exchange server (**3**). A control trigger quarantines the email. The CMS (**4**) maintains a queue of quarantined emails (**5**). A reviewer checks quarantined emails in the iConsole (**6**). The reviewer can either release or reject a quarantined email.

The Quarantine Manager (**7**) regularly checks the quarantine queue on the CMS, checking for emails that have been released or which have timed out (**8**). It then forwards these emails, either via the original Exchange server (**9a**) or, if so configured, though an alternative Exchange server (**9b**) to the intended recipient (**10**).

Note: For simplicity, this diagram omits the policy engine hub and policy engines.

Quarantine Manager Requirements

Before deploying the Quarantine Manager, note the following requirements:

CA DataMinder Server

You can install the Quarantine Manager on the CMS or a utility machine.

Email Client or Access to SMTP Server

To enable Quarantine Manager to release emails captured by:

- Exchange server agent or Outlook endpoint agent, the Quarantine Manager host machine must be running Outlook. Outlook *must* be the default email application on the host machine.
- Domino server agent or Notes endpoint agent, the Quarantine Manager host machine must be running Notes.
- The NBA or Milter MTA agent, the Quarantine Manager host machine must be able to access an SMTP server. You specify the target SMTP server when you configure the Quarantine Manager registry values.

QM Domain User

The Quarantine Manager uses this domain account to access the CMS and to log on to the specified Exchange or Domino server or, for Sendmail and Postfix mails, the specified SMTP server. This account requires various rights and privileges.

More information:

Specify a QM Domain User (see page 231)

Pre-deployment Considerations

Before deploying the Quarantine Manager, be aware of the following issues.

If Using Outlook Endpoint Agents and the Exchange Server Agent

If you want to quarantine emails without notifying the sender, you must use the Exchange server agent. It is not possible to quarantine emails silently using an Outlook endpoint agent. This has implications if your CA DataMinder deployment uses both Outlook endpoint agents and an Exchange server agent.

By default the Exchange server agent does not reprocess emails that have already been processed by an Outlook agent. Therefore, an email cannot normally be quarantined by the Exchange server agent if other policy triggers have already been applied by an Outlook agent.

To fix this, you must explicitly configure the Exchange server agent to reprocess (and apply quarantine actions) to emails already processed by an endpoint agent.

Follow these steps:

1. Go to this registry key on the Exchange server:
   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
     \CurrentVersion\Exchange
   ```
2. Add this registry value to the \Exchange registry key:
   ```
   ReprocessClientEmails
   ```
   The ReprocessClientEmails registry value is described in the 'Exchange, Domino and IIS SMTP' chapter of the *Message Server Integration Guide*.

Emails Detected At Network Boundary May Not Be Quarantined

Under certain conditions, a ‘quarantine' action applied by the NBA or Milter MTA agent *may* only affect external recipients; internal recipients *may* already have received the email. In general, this problem only affects emails sent to both internal and external recipients. It is also depends on how your CA DataMinder server agents and user policies are configured.

Because the NBA and the Milter MTA agent capture emails at the network boundary, an email may already have been received by some internal recipients before it is detected by the NBA or Milter MTA agent. This happens if the email was routed directly to these recipients by Exchange or Domino without passing through the NBA or a Sendmail or Postfix server. As a result, if the NBA or Milter MTA agent then quarantines the email, the email is only withheld from external recipients pending its release. Any internal recipients *may* already have received the email.

Note: The Milter MTA agent is hosted on Sendmail or Postfix servers, which are typically located on the boundary of your corporate network.

Multiple Quarantine Managers

For maximum reliability, you can install multiple Quarantine Managers.

Each Quarantine Managers registers with the CMS. The CMS automatically selects one to be the ‘active’ Quarantine Manager. The CMS keeps in regular contact with the registered Quarantine Managers and if the active one fails for any reason, the CMS selects a standby Quarantine Manager to be the ‘active’ one.

If a Quarantine Manager is installed on the CMS, this always defaults to be the ‘active’ one when the CMS starts up.

To allow a different Quarantine Manager to be active on startup

1. Edit the ActiveOnStartup (see page 238) registry value on that Quarantine Manager’s host server.

2. Restart the CMS.

More information:

Quarantine Manager Registry Values (see page 238)

Deploy the Quarantine Manager

Setting up CA DataMinder to quarantine emails that require an urgent review is a multi-step procedure. Briefly, you must:

1. Specify a domain user.
2. Allow access to designated mailboxes.
3. Install the Quarantine Manager
4. Configure the Quarantine Manager
5. Mark emails for quarantine. To do this, you must setup control triggers and actions in the relevant user policies.

Individual steps are described in the following sections.

More information:

Specify a QM Domain User (see page 231)
Enable Access to the Designated Mailbox (see page 233)
Install the Quarantine Manager (see page 236)
Configure the Quarantine Manager (see page 237)
Mark Emails for Quarantine (see page 237)
Quarantine Manager Log Files (see page 237)

Specify a QM Domain User

The Quarantine Manager must be able to access the CMS directly to check for emails released from quarantine. It also needs to be able to log on to the specified Exchange, Domino or Sendmail or Postfix server. Because the Quarantine Manager runs as a process within the CA DataMinder infrastructure, you need to assign various rights and privileges to the logon account for the infrastructure service.

Therefore, you must either create a new domain user in Active Directory, or choose an existing domain user. This is your 'QM domain user'. Throughout this section, the term 'QM domain user' refers to the logon account for the CA DataMinder infrastructure service on the Quarantine Manager host machine.

The QM domain user must also be a local administrator on the Quarantine Manager host server and have the 'Log on as a service' security privilege. You supply the QM domain user credentials when you install the Quarantine Manager.

You must also ensure the QM domain user can access the designated Exchange mailbox or the designated Notes or SMTP server in order to release quarantined emails.

Finally, you must create or designate a corresponding CA DataMinder user account. See the following sections for details.

More information:

Assign the 'Log on as a service' Privilege (see page 231)
Add to Local Administrators Group (see page 232)
Create a 'QM' CA DataMinder User Account (see page 232)
Install the Quarantine Manager (see page 236)

Assign the 'Log on as a service' Privilege

You must manually assign the 'Log on as a service' security privilege to the QM domain user. To do this:

1. Ensure that you are logged on to the target server with local administrator rights.
2. On the target server, open the Local Security Policy applet or, if this server is a domain controller, open the Domain Controller Security Policy applet. Both applets are available in Administrative Tools.
3. Expand the Local Policies branch and select User Rights Assignment. This security area determines which users have logon privileges on the local computer.
4. Assign the Log on as a service privilege to the QM domain user account.

Add to Local Administrators Group

On the Quarantine Manager host server, add the QM domain user to the Administrators group.

Create a 'QM' CA DataMinder User Account

After specifying your QM domain user, you must create a matching CA DataMinder user account. The CA DataMinder user account must have the same account name as the QM domain user. The Quarantine Manager uses this CA DataMinder account to log on to the CMS when checking for emails released from quarantine.

To create a QM user account

1. In the CA DataMinder Administration console, create a new user.

   See the Administration console online help for details; search for ‘new users’.

2. When you specify the user name, include the domain prefix to ensure compatibility with the account name for the QM domain user (for example, UNIPRAXIS\QMUser).

3. Specify the management group for the Quarantine Manager. We recommend that you set the management group to the top-level Users group.

4. Set the User Role to 'System Process'.

   This role confers the following administrative privileges on the QM user:

   Admin: Use single sign-on:

   > This privilege enables the Quarantine Manager to log on to the CMS automatically (without needing to provide authentication), even if the CMS machine policy setting ‘Allow single sign-on?’ is set to False.

   Admin: Disable security model filtering:

   > This privilege enables the Quarantine Manager to bypass built-in security measures and search for events without management group restrictions. In effect, assigning this privilege guarantees each reviewer can retrieve all the quarantined events associated with users in their respective management groups.

   Events: Control quarantined events:

   > This privilege permits the Quarantine Manager to access emails released from quarantine.

Enable Access to the Designated Mailbox

When an email is released from quarantine, CA DataMinder uses a designated mailbox to forward the email to its intended recipients. Therefore, before you install the Quarantine Manager, you must ensure that the QM domain user is a mailbox-enabled user. That is, the QM domain user must be able to access the following mailboxes:

Exchange Servers

The mailbox specified by the MapiEmailServerName registry value.

Notes Servers

The mailbox specified by the NotesEmailServerName registry value.

SMTP Servers

The mailbox specified by the SmtpEmailServerName registry value.

Note also the mailbox access requirements in the following sections.

Exchange Servers: Grant the Send As Permission

The QM domain user must have the ‘Send As’ Exchange permission when releasing emails from quarantine. This permission enables the QM domain user to forward released emails so that they appear (to the recipients) to have come directly from the original sender.

Important! If the QM domain user does not have the ‘Send As’ permission for a particular user, the Quarantine Manager cannot forward that user’s e-mails when they are released from quarantine!

Example: Granting the Send As Permission

Exchange Server 2007 and 2010 support various methods for assigning the Send As permission. The following example describes how to use Active Directory Users and Computers to grant this permission. This method allows the QM domain user to send emails as any existing user plus any new user accounts added in the future.

1. Open Active Directory, Users and Computers.
2. Click View, Advanced Features.

   Note: When you enable Advanced Features, the Security tab is displayed in the Domain Properties dialog (step 4).
3. Right-click the appropriate domain and click Properties.

   The Domain Properties dialog appears.
4. Click the Advanced button in the Security tab.

   The Advanced Security Settings dialog appears. This dialog lists the domain users.
5. Double-click the account name for the QM domain user.

   Note: If the QM domain user is not listed, click the Add button and add the account to the list.

   The Permission Entry dialog appears.
6. Do one of the following:
   - Click 'User Object' in the Apply To list.
   - (Windows Server 2008 and 2012) Click 'Descendant User Objects' in the Applies To list.
7. Select the Send As checkbox in the Permissions list and click OK.

   The Permission Entry dialog closes and returns you to the Advanced Security Settings dialog.
8. Click Apply. Then click OK.

   The Properties window closes. The Send As permission has been successfully granted to your QM domain user.

More information:

Quarantine Manager Registry Values (see page 238)

Notes Servers: Set Credentials

When configuring the credentials for the Notes server and mailbox using the NotesEmailServerName and NotesEmailID registry values, you must also set the password for the Notes 'Current User' on the Quarantine Manager host server. This ensures that the Quarantine Manager has access to the Notes mailbox specified in the NotesEmailID registry value.

So that this password is not stored in a registry value (which would represent a security loophole), you must configure wgncred.exe to securely cache the password itself using the following component ID:

```
Quarantine Manager for Lotus Notes: QMGRNOTES
```

More information:

Set Account Credentials with WgnCred.exe (see page 550)
Quarantine Manager Registry Values (see page 238)

SMTP Servers: Set Credentials

When integrating with Sendmail or Postfix, the Quarantine Manager uses the mailbox on the SMTP server associated with the QM domain user account to send e-mails released from quarantine on to their intended recipients.

Access to this mailbox is specified by the SmtpEmailID and SmtpEmailServerName registry values. The password SmtpEmailID user account is passed to the Sendmail or Postfix server with the account name. Use wgncred.exe to securely cache this password. The required component ID is:

```
Quarantine Manager for SMTP: QMGRSMTP
```

More information:

Set Account Credentials with WgnCred.exe (see page 550)
Quarantine Manager Registry Values (see page 238)

Install the Quarantine Manager

You install the Quarantine Manager using the CA DataMinder server installation wizard.

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.
3. In the Advanced Install Options screen, choose CA DataMinder Platform and then click Install.

   This launches the CA DataMinder server installation wizard in a separate window.

4. In the server installation wizard, navigate to the Custom Setup screen.
5. In the Custom Setup screen, choose Quarantine Manager.
6. In the Server Type screen, we recommend that you choose a utility machine. The Quarantine Manager can also run on a CMS or gateway, but utility machines are explicitly designed to host CA DataMinder add-ins such as the Quarantine Manager.
7. In subsequent screens, specify the parent server (typically the CMS), the location of the \Data folder, plus details about the local CA DataMinder database.
8. In the Service Accounts screen, specify the QM domain user as the logon account used by the CA DataMinder infrastructure service.
9. In the final wizard screen, click Install to start the file transfer.

   When the installation is complete, you must manually configure the Quarantine Manager by editing the registry.

More information:

Server Installation Features (see page 38)
Specify a QM Domain User (see page 231)
Configure the Quarantine Manager (see page 237)

Configure the Quarantine Manager

After installation, you must configure how the Quarantine Manager handles quarantined emails. To do this, you must edit the registry on the Quarantine Manager host machine.

- **Specify the Quarantine mailboxes:** To specify the mailbox used by the Quarantine Manager when forwarding released emails, you must edit the MapiEmailID or NotesEmailID registry value.
- **Specify the email Server:** To specify which Exchange, Domino or Sendmail/Postfix server the Quarantine Manager connects to when forwarding released emails, you must edit these registry values:

  ```
  MapiEmailServerName
  NotesEmailServerName
  SmtpEmailServerName
  ```

  Important! Because these registry values are not specified during installation, when the Quarantine Manager first starts a warning message is added to the Activity log and the Quarantine Manager shuts down. After supplying the required registry values, you must restart the local infrastructure to enable the Quarantine Manager to start.

More information:

Quarantine Manager Registry Values (see page 238)

Mark Emails for Quarantine

Now you need to set up the Quarantine feature to identify emails that need urgent reviewing. To do this, edit control triggers and actions in the relevant user or group policy. For details, see the Administration console online help; search the index for ‘quarantined emails’.

Quarantine Manager Log Files

The Quarantine Manager writes messages to the Activity log and also to its own log saved on the QM host machine. The QM log is configured in the registry. Activity log messages generally record the outcome of quarantine operations, while the QM’s own log provides more diagnostic details.

More Information:

Quarantine Manager Registry Values (see page 238)

Quarantine Manager Registry Values

To configure how the Quarantine Manager processes e-mails released from quarantine and how it interacts with the CMS database, you must edit values in the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\Quarantine Manager
```

Within this registry key, the registry values you need to edit are:

MapiEmailServerName

Type: REG_SZ

Data: Set this value to the name of the Exchange Server hosting the mailbox of the MapiEmailID user account (see below).

MapiEmailID

Type: REG_SZ

Data: Set this value to the Exchange mailbox you want to use. You need only specify the mailbox; do not include the domain prefix. For example, set this value to 'QMuser' not 'Unipraxis\QMuser'.

The Quarantine Manager will use this mailbox to send emails released from quarantine on to their intended recipients.

NotesEmailServerName

Type: REG_SZ

Data: Set this value to the name of the Domino Server hosting the mailbox of the NotesEmailID user account (see below).

NotesEmailID

Type: REG_SZ

Data: Set this value to the Domino mailbox (that is, the NSF file) you want to use. The Quarantine Manager will use this user account to send emails released from quarantine on to their intended recipients.

You need to specify the path to the NSF file relative to the Data folder on the Domino server. Paths are not case-sensitive and the .nsf suffix is optional. For example:

mail\unipraxis\qmuser.nsf

You must also cache the password for the Notes 'Current User' on the Quarantine Manager host server. Use wgncred.exe to do this.

SmtpEmailServerName

Only required for Sendmail, Postfix, and IIS SMTP integrations.

Type: REG_SZ

Data: Specifies the name of an SMTP server that you want to use for sending emails released from quarantine. For example, this may be a Sendmail or Postfix server.

This registry value can also specify the TCP port used for communication between the Quarantine Manager and the SMTP server. If omitted, the port number defaults to 25. To specify a non-default port number, append the number to the server name, separated by a colon. For example:

```
unipraxis.com:25777
```

SmtpEmailID

Only required for Sendmail, Postfix, and IIS SMTP integrations.

Type: REG_SZ

Data: Set this value to the user account that the Quarantine Manager will use to log on to the Sendmail or Postfix server. This registry value is only required if the Quarantine Manager authentication method is not ‘None’ (as specified by SmtpAuthType; see below).

The Quarantine Manager will use this user account to send emails released from quarantine on to their intended recipients.

If you do need to specify an SmtpEmailID user account, you must cache this user’s password using wgncred.exe.

If the SmtpAuthType is set to NTLM and the Windows Mail server is set to use 'Integrated Windows Authentication' then you need only supply the username rather than qualify it with a domain. That is, 'qmgr' not 'unipraxis\qmgr'.

SmtpAuthType

Only required for Sendmail, Postfix, and IIS SMTP integrations.

Type: REG_SZ

Data: Defaults to None. This specifies which standard SMTP authentication type the Quarantine Manager uses to connect to the Sendmail or Postfix server. The following values are supported:

None

Plain

Login

NTLM

CRAM-MD5

We recommend you choose None for unauthenticated connections. However, your Sendmail or Postfix server must be configured to accept connections from the Quarantine Manager host machine.

We do not normally recommend Plain or Login authentication because under these protocols the logon password is sent as unencrypted plain text across the network.

NTLM and CRAM-MD5 authentication can be used to connect to Sendmail or Postfix servers on Windows and Unix machines respectively. However, although these protocols do not send unencrypted logon credentials, you must still ensure that these credentials are protected.

If you use Plain, Login, NTLM or CRAM-MD5 authentication, you must set up the SmtpEmailID registry value to pass the logon account details to the SMTP server. You must also use Wgncred.exe to cache the password.

ActiveOnStartup

Type: REG_SZ

Data: Defaults to 1. In a multiple Quarantine Manager environment, this value determines which Quarantine Manager is active on startup.

To specify a standby Quarantine Manager, set this value to zero. To specify the active Quarantine Manager, set this value to 1.

MaxRetryAttempts

Type: REG_DWORD

Data: Defaults to 5 attempts. This is the number of times the Quarantine Manager tries to resend a released email.

ReleaseTime_Mins

Type: REG_DWORD

Data: Defaults to zero minutes. This specifies how long an e-mail is held in quarantine before it is automatically released. The zero minute default means that emails are kept in quarantine indefinitely.

SleepTime_Secs

Type: REG_DWORD

Data: Defaults to 30 seconds. This is how long the Quarantine Manager waits between querying the CMS for e-mails released from quarantine. Increasing this interval reduces the impact that the Quarantine Manager has on database performance at the cost of a longer delay before a message can be sent on to its intended recipients.

SaveMessages

Type: REG_DWORD

Data: Defaults to 0. After the Quarantine Manager sends a released email on to its intended recipients, it is deleted from the mailbox. Set this to 1 to retain the email after it has been sent.

LogLevel

Type: REG_DWORD

Data: Defaults to 2. This setting determines the level of logging for the Quarantine Manager’s own log file (general operational details are recorded in the CA DataMinder Activity log). For example, you can configure the Quarantine Manager to only log errors or warning system messages.

Log entries are written to the wgnqmgr_<date>.log file, where <date> is the date and time when the log file was created; the file is located in CA's \data\log subfolder of the Windows All Users profile. The supported logging levels are:

1 Errors only

2 Errors and warnings

3 Errors and warnings, plus informational and status messages

Note: Setting LogLevel=3 will cause the log file to grow extremely rapidly. This level of logging is provided for testing and diagnostic purposes, for example, it shows storage and retrieval on every resource item.

LogMaxNumFiles

Type: REG_DWORD

Data: Defaults to 10. This specifies the maximum number of log files. When the maximum number of log files exists and the maximum size of the latest is reached (see below), the oldest log file is deleted to enable a new one to be created. Log entries are written to a wgnqmgr_<date>.log file, where <date> is the date and time when the log file was created

LogMaxSizeBytes

Type: REG_SZ

Data: Defaults to 1,000,000. This specifies the maximum size for each log file. When the current log file reaches its maximum size, the Quarantine Manager creates a new log file. Log entries are written to a wgnqmgr_<date>.log file, where <date> is the date and time when the log file was created.

LogFilePath

Type: REG_SZ

Data: Defaults to empty. This specifies the folder you want to write log files to. If the path is not defined, the log file is saved to CA's \data\log subfolder of the Windows All Users profile on the machine hosting the Quarantine Manager. The QM domain user must have write access to the specified folder.

More information:

Notes Servers: Set Credentials (see page 235)
SMTP Servers: Set Credentials (see page 235)
Specify a QM Domain User (see page 231)

Email Release Procedure

After the Quarantine Manager has received a released e-mail event from the CMS:

1. It connects to the Exchange, Domino or Sendmail/Postfix server specified by:
   `MapiEmailServerName, or`
   `NotesEmailServerName, or`
   `SmtpEmailServerName`
2. Then it connects to the mailbox associated with the Quarantine mailbox user, in turn specified by one of these registry values:
   `MapiEmailID`
   `NotesEmailID`
   `SmtpEmailID`
3. Finally, it sends the email on to the intended recipient(s).

More information:

Automatic Release of Timed-out Emails (see page 243)
Failure to Forward Released Emails (see page 243)
Encrypted Emails (see page 243)
Quarantine Manager Registry Values (see page 238)

Automatic Release of Timed-out Emails

To ensure that business critical messages are not delayed unnecessarily, the Quarantine Manager can automatically release an email from quarantine after a specified period, even if it has not been reviewed. The interval between queries from the Quarantine Manager to the CMS database is configurable by a registry value. By default, the Quarantine Manager will never automatically release an email event.

Failure to Forward Released Emails

If the Quarantine Manager fails to send a released email to its intended recipients, it adds an entry to the Activity Log. It will then try to resend the email. The MaxRetryAttempts registry value determines the number of retries.

If the Quarantine Manager cannot access the specified email server (for example, because it is temporarily shut down for maintenance), it waits before attempting to send the email again, gradually increasing the interval between attempts up to a maximum of 5 minutes.

If the number of failed attempts matches the limit specified by MaxRetryAttempts, the Quarantine Manager quits trying to resend and the email is blocked.

More information:

Quarantine Manager Registry Values (see page 238)

Encrypted Emails

See also the known issues for encrypted emails and the Quarantine Manager (see page 609).

Chapter 11: Account Import

Note: CA DataMinder integrates directly with CA Identity Manager. This integration allows you to use CA Identity Manager to maintain your CA DataMinder user accounts. For more information, see the Technical Information (see page 582) section of this guide.

This section contains the following topics:

Overview

To simplify mass deployments, Account Import enables administrators to import user details into CA DataMinder from an external Lightweight Directory Access Protocol (LDAP) directory or a source file. Account Import can import new users and groups into the existing CA DataMinder user hierarchy, or it can reorganize existing users to synchronize them with an external hierarchy. It can also import user attributes such as email addresses and employee IDs.

For machines, Account Import can bulk create new accounts and pre-assign them to parent servers. It can also re-parent existing client machines and gateways in advance of the CA DataMinder rollout.

This section describes how to run import operations from a command line or by using the Account Import wizard. It also provides full details about parameters files and command files.

Import Methods and Sources

Account Import can import user details from several sources and supports two import methods: command line operations and the Account Import wizard. In all cases, Account Import converts all source data to a *command file* before importing.

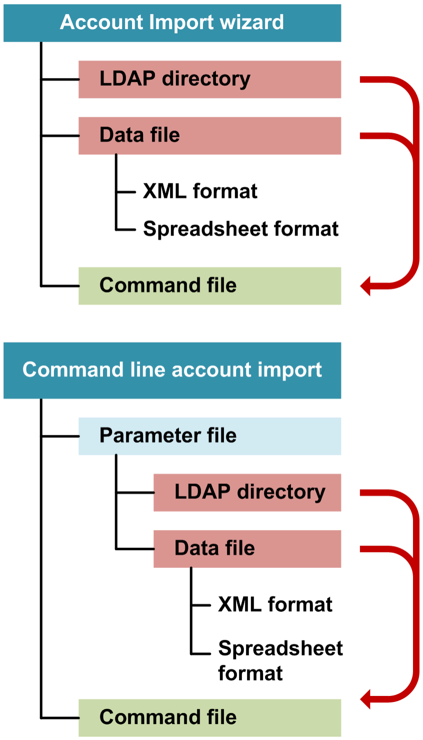

Import Methods

You can import user details by running:

Account Import wizard

This is the simplest method. You run the wizard from the Administration console. In a multiple-CMS environment, the wizard creates new users on the currently selected CMS.

Command line import operations

Account Import supports command line operations for importing user details from an LDAP directory or other external source. You can schedule regular command line operations to keep your LDAP directory and CA DataMinder user hierarchy synchronized.

More information:

Command Line Import Operations (see page 263)

Import Sources

Account import can import user information directly from an LDAP directory, parameter file, data file, or command file:

LDAP directory

The Lightweight Directory Access Protocol (LDAP) enables directory services to manage directory objects. Objects and attributes in an LDAP directory are exposed to any other application that uses the LDAP protocol. CA DataMinder can import user details from the following LDAP directories:

- Domino Server 7.x or 8
- Microsoft Active Directory
- Novell eDirectory (NDS)
- Netscape/Sun ONE Directory Server

Note: You can import from an LDAP directory using SSL.

Parameter files

(Only supported for command line import operations) These files contain import parameters similar to the configuration options provided by the Account Import wizard.

Data files

These are structured files of user data, in XML or spreadsheet-compatible format. Data files contain encoded versions of all or part of an external user hierarchy. These files contain the user details necessary for CA DataMinder to create, or re-create, this external hierarchy on the CMS.

For details about the XML schema, see the *Account Import XML Schema Guide*.

Important! Spreadsheet-based data files cannot be used to import email addresses!

Command files

These are import configuration files containing CA DataMinder user and machine import commands (for example, ‘create new user’ or ‘set user attribute’). Typically, you import command files to make specific changes or additions to your *existing* CA DataMinder user hierarchy.

More information:

Set Up Secure Sockets Layer (SSL) (see page 265)

Synchronizing Users

When synchronizing your CA DataMinder user hierarchy with your principal user directory, we recommend that you run a single import operation to synchronize all of your users in one go.

Do not run several smaller operations, each synchronizing a specific set of users. With this approach, there is a risk that a partial synchronization may inadvertently move unknown users (that is, users not present in the specified source LDAP directory) to an ‘exceptions’ group.

If a single operation is not practical, you can run a separate operation for each network domain. For example, if your users are spread across multiple domains, you may prefer to run a separate Account Import operation for each domain. The ‘Synchronize users from this domain’ option ensures that only users in the specified domain can be reorganized within the CA DataMinder user hierarchy.

Account Import Log Files

The results of each import operation are written to an individual Account Import log file. This lists all changes or additions to CA DataMinder user or machine hierarchies and includes any errors that may have occurred. For example, if Account Import failed to find a specified parent group or parent server, this is recorded in the log file. If Account Import fails to recognize an import operation in a parameter file or command file, this is also recorded in the log file.

Account Import log files are saved in CA's \data\log subfolder of the Windows All Users profile.

Log files take the format: ldap_200201200945.log.

To view the log file in the Administration console, choose Manage, Log Files or click .

Account Import Privileges

Your administrators must have the necessary administrative privileges when running Account Import operations:

To do this	You need this privilege
Assign roles to users (Command files only)	Admin: Edit User Roles
Create new machines; reorganize the machine hierarchy (Command files only)	Machines: Edit the user Hierarchy
View log files of user and machine import operations	Machines: View Log files
Create new user; reorganize the user hierarchy	Users: Edit the user Hierarchy
Assign temporary passwords to user accounts (Command files only)	Users: Reset Users Password

Synchronizing Email Addresses

Important! User email addresses on the CMS *must* remain synchronized with your principal user directory (typically an LDAP directory such as Active Directory).

One of the most important reasons for using Account Import is to synchronize email addresses in the CMS database with addresses in the LDAP directory. Such synchronization is essential for CA DataMinder features that rely on email address mapping.

During synchronization, any addresses in the CMS database that are not present in the LDAP directory are deleted from the CMS. For example, if an email address has been deleted from the LDAP directory since the last synchronization, it is deleted from the CMS database during the next synchronization.

Account Import Wizard

Start the Account Import wizard in the Adminsitration console. Click Tools, Account Import Wizard. The wizard steps you through each stage of the import process.

Note: Some wizard screens may not appear, depending on which import options you choose.

1. **Select source of account data screen**

   In the first wizard screen, specify the source for the imported user details.

   Synchronize to Data Source

   > Choose the LDAP Database or Data File check boxes to synchronize your existing CA DataMinder user hierarchy with these external sources.
   >
   > If you choose both the LDAP and Data File check boxes (that is, you want to simultaneously import from an LDAP directory and a data file), you can specify how Account Import handles duplicate records (any user listed in both sources). By default, the user record in the XML directory is imported while the record in the data file is ignored, but you can override this default.

   Input from Command File

   > Specify the command file that contains the changes or additions to your existing CA DataMinder user hierarchy.
   >
   > **Note:** By default, when you export any branch of the CA DataMinder user hierarchy to a command file, the target file name has an .acc extension. You can then edit this file before re-importing it.

2. **LDAP Logon screen**

   (Only applicable if importing from an LDAP directory. See step 1.)

   You must supply logon details for the source LDAP Directory:

   LDAP Server

   > Identify the server hosting the source LDAP directory. Enter its name or an IP address.

   Port number

   > Enter the TCP/IP port number used to connect to the LDAP server. CA DataMinder uses this port to communicate with the LDAP server.

   Base DN/Domain

   > Identify the LDAP server's base DN or domain. For example, to specify an Active Directory domain, enter one of these formats:
   >
   > `company.com`
   > `dc=company,dc=com`
   >
   > **Note:** If Account Import can detect the default DN, it is shown automatically. Also, some configurations, for example Domino Server, may require you to leave this field empty.

   User

   > Enter your user name on the LDAP Server. The format for this name depends on the type of LDAP database. For example, if you import users from a Microsoft Exchange server, this name will be the same as your domain user name, with your domain and name separated with a backslash, such as:
   >
   > `unipraxis\frankschaeffer`
   >
   > On other LDAP databases, this name may be a fully qualified LDAP distinguished name, for example:
   >
   > `cn=frankschaeffer,o=unipraxis`

   Password

   > Enter the password for your LDAP user.

   Note: If the LDAP server permits anonymous access, leave both the User and Password fields blank.

3. **LDAP Search Filters screen**

   (Only applicable if importing from an LDAP directory. See step 1.)

   Where possible, the wizard automatically detects the type of LDAP directory (for example, Microsoft Active Directory) and key details about the LDAP directory structure. The wizard provides 'best guess' default search filters, but you can override these if necessary. Specifically, you must ensure that the following fields contain correct values:

   User Name Attribute

   Specify the LDAP attribute that holds the user names.

   User Search Filter

   Specify the LDAP search filter needed by the wizard to extract users from the LDAP database.

   Group Search Filter

   Specify the LDAP search filter needed by the wizard to extract the LDAP containers that correspond to CA DataMinder user groups.

   Note: If you override the default search filters and specify different object classes and categories, ensure that the new filter conforms to RFC 2254.

4. **LDAP Source Directory screen**

   (Only applicable if importing from an LDAP directory. See step 1.)

   Specify the root directory for user data extracted from the LDAP directory. All users and groups at and below this root directory will be copied into CA DataMinder. Click Browse to select the root-level LDAP tree level.

   For example, select ‘ou=Unipraxis/ou=Sales’ to import all users from this level downwards:

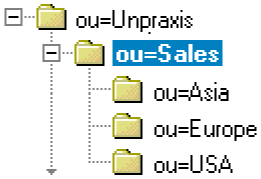

5. **Users Tree Root screen**

   Specify the target parent group in the CA DataMinder user hierarchy. You can only choose one of your management groups as the parent group. All users and groups imported from LDAP and or a data file are added to this parent group.

   Note: If you choose to reorganize existing CA DataMinder users to match the directory structure in LDAP or the structure specified the data file (you choose this in step 6), the reorganization only affects CA DataMinder users within the target parent group.

6. **Synchronization Scope screen**

   (Applicable if importing from an LDAP directory or a data file. See step 1.)

   Now define the synchronization scope. The Account Import wizard enables you to synchronize your CA DataMinder user hierarchy with an external source. You can select any combination of the following synchronization options.

   Create new users

   This option creates new CA DataMinder accounts for unknown users. That is, it creates a new account for each imported user who has no corresponding account in CA DataMinder.

   Note: If a user is created with a user name matching a user account that was previously deleted, CA DataMinder can automatically recreate the deleted user.

   Re-organize existing users

   This option rearranges the existing hierarchy of CA DataMinder users and groups to synchronize it with the Group Structure. You define the Group Structure in the next wizard screen.

   If you do not select this option, all existing CA DataMinder users stay in their current group.

   Copy user attributes

   This option updates existing user accounts with e-mail addresses and attributes imported from corresponding users in the LDAP directory or data file. You specify these attributes in later wizard screens. See steps 9 and 10 for details.

   Note: The full name associated with each CA DataMinder user account is imported automatically from the LDAP directory.

   - **Email addresses can be deleted:** This option specifies whether emaildelete commands are carried out by the synchronization.

     Important! We do not recommend that you use this parameter, as existing email events may no longer be associated with the correct user—see the /ed parameter for details.

Synchronize users from this domain

This option prefixes names for new user accounts with the specified domain (such as unipraxis\srimmel). If importing users from an LDAP directory, specify the domain for the LDAP host server.

Note: This option is essential if single sign-on is enabled on your CMS.

Set policy exemption state

This option exempts specified CA DataMinder user accounts from policy.

If you are importing user details from a data file, verify that the data file uses the correct format to identify policy-exempt users.

LDAP policy exemption attribute

(Available only if importing users directly from an LDAP directory) Click the LDAP attribute that you want to use to identify exempt users. CA DataMinder automatically exempts any imported users from policy if they have this attribute.

Value

If required, you can use specific attribute values to filter the users that you want to exempt from policy. For example, you select the Office attribute and set the value to London, you can exempt all users in your London office from policy.

7. **Import Options screen**

These options determine how Account Import handles anomalous users and groups, whether you must confirm the changes, and how new user names are composed.

Group Structure

These options determine how imported users are organized into parent groups in the CA DataMinder user hierarchy.

- **Use LDAP hierarchy to group users:** This option creates a new set of user groups that match the hierarchical structure of the source LDAP directory or data file. The new group structure is rooted at the source LDAP directory specified in step 4. It is created below the CA DataMinder parent group specified by the User Tree Root in step 5.
- **Place all users in User Tree Root:** This option imports all users into a flat, non-hierarchical group structure. That is, all imported users are added to the parent group specified by the User Tree Root in Step 4.
- **Use LDAP attributes to group users:** This option derives a hierarchy of parent groups based on a concatenation of specified LDAP attributes, or attributes specified in a data file. For details, see step **8**.

Create empty groups

Available only if you selected ‘Create new users’ in step **6**.

The LDAP directory structure may contain empty containers. These are containers that hold subcontainers or other items, but no users. When importing users from the LDAP directory, you can set up the import wizard to ignore these empty containers or to create corresponding empty user groups in CA DataMinder.

If you select this option, the wizard creates empty user groups for each empty LDAP or data file container.

If you clear this option, the wizard ignores empty containers. For example, an LDAP directory may include the following branch:

`LDAP:` `ou=Unipraxis/ou=London/ou=Sales`

If the ‘Sales’ container is empty of users but the ‘London’ container is not empty, the wizard creates the following hierarchy in the Administration console:

`CA DataMinder:` `Unipraxis/London`

Move unknown users to...

(Available only if you selected ‘Re-organize existing users’ in step 6.)

If your existing CA DataMinder user hierarchy contains users or groups not present in the LDAP directory, you can move them to an 'exceptions' group. The exceptions group can be any existing group in the user hierarchy. If you do *not* select this option, any non-LDAP users and groups are preserved in the CA DataMinder user hierarchy.

This option only affects CA DataMinder users within the specified target parent group.

Users prepended with a domain name other than the one set on the Synchronization Scope screen are not moved (see 'Synchronize users from this domain' in step 6).

Exempt unknown users from policy

Select this check box to exempt unknown users from policy.

If your existing CA DataMinder user hierarchy contains users not present in the LDAP directory or XML data file, you can exempt them from policy.

Exempt users are users who have a CA DataMinder account on the CMS but who are exempt from policy. That is, CA DataMinder does not monitor email, web or file activity for policy-exempt users.

Manual confirmation

If you select this option, you must confirm the all of the resulting changes to the user hierarchy. If you do not select this option, synchronization is automatic and you cannot confirm or reject individual changes.

8. **Create Group from LDAP attributes screen**

   (Only available only if you selected ‘Use LDAP attributes to group users’ in step 7.)

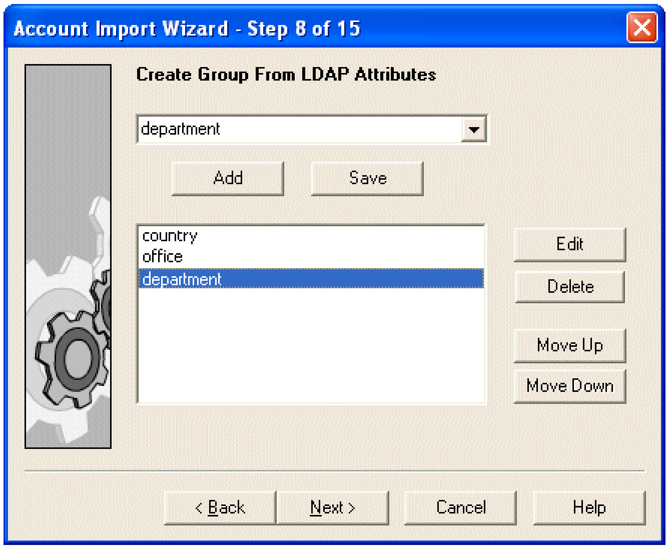

Account Import Wizard - Step 8 of 15

Create Group From LDAP Attributes

department

country
office
department

If required, Account Import can derive a hierarchy of parent groups based on a concatenation of specified LDAP attributes.

Choose which LDAP attributes to use, and specify the order in which they are used to derive a group hierarchy. For example, when these LDAP attributes are arranged in the following order:

```
country
office
department
```

They produce this group hierarchy in CA DataMinder:

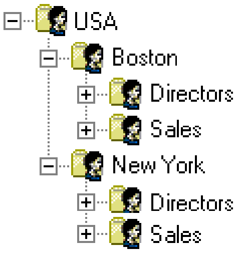

Adding custom attributes

Account Import only displays the most commonly used LDAP attributes in this screen. If you need to add an attribute not listed here (for example, an employee attribute custom created for your organization), use the Edit and Save buttons to add this attribute to the group-defining list.

Modifying attribute values

If you need to modify the values of an LDAP attribute before using these values to derive a group hierarchy in CA DataMinder, you can append a conversion expression, enclosed in square brackets, to the attribute name. Use the Edit and Save buttons to add the attribute-plus-expression to the attribute list.

For example, an LDAP directory for Unipraxis includes an 'office' attribute. Values for the office attribute have a 'UX-' prefix, such as UX-Boston and UX-New York. Use the following conversion expression to strip out the prefix in the resulting CA DataMinder user groups (that is, Boston and New York):

```
office[if"UX-"["???{%untilEnd%}"] else["{?%untilEnd%}"]]
```

9. **Email attributes screen**

   (Available only if importing from an LDAP directory *and* you selected ‘Copy user attributes’ in step 6.)

   Account Import can sychronize e-mail addresses in the CMS database with addresses in an external source, typically an LDAP directory. *Such synchronization is essential for CA DataMinder features that rely on e-mail address mapping!*

   Add the LDAP attributes that contain e-mail addresses. Each imported address is associated with a CA DataMinder user. The association is based on an *anchor attribute* that ties each LDAP user to a specific CA DataMinder user (see Step 11).

   Note: If you use the ICAP agent to integrate with BlueCoat ProxySG servers, you must import the distinguishedName attribute.

10. **User Attributes screen**

(Only available only if importing from an LDAP directory *and* you selected 'Copy user attributes' in step 6.)

Use this screen to import LDAP attributes and map them to attributes for CA DataMinder user accounts. CA DataMinder attributes are listed on the left. LDAP attributes are listed on the right.

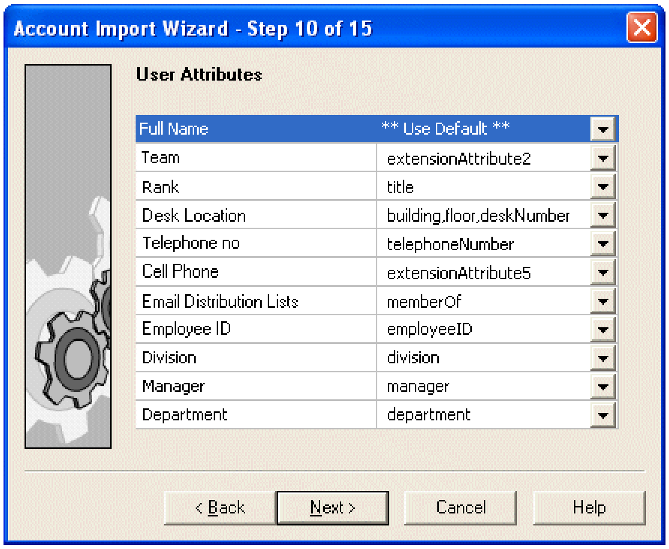

Full Name	** Use Default **
Team	extensionAttribute2
Rank	title
Desk Location	building,floor,deskNumber
Telephone no	telephoneNumber
Cell Phone	extensionAttribute5
Email Distribution Lists	memberOf
Employee ID	employeeID
Division	division
Manager	manager
Department	department

CA DataMinder lets you define custom attributes for user accounts. For example, you can create an Employee ID attribute and assign a unique ID to each user in your organization. When the import operation runs, the Account Import updates the attributes for each CA DataMinder user with the corresponding attribute values in the LDAP directory.

To map an LDAP attribute to CA DataMinder attribute, select CA DataMinder attribute then choose an LDAP user attribute from the drop-down list.

Combining LDAP attributes

> To combine multiple LDAP attributes and write them as a single value to a CA DataMinder attribute, double-click the LDAP attribute, then manually type a comma separated list of the LDAP attributes you want to combine.For example:
>
> Desk Location = building,floor,deskNumber

Renaming attributes

> You can rename any CA DataMinder or LDAP attribute. Double-click the attribute and type its new name.

Modifying attributes

> You can modify the imported value for any LDAP attribute before writing it to an attribute of a CA DataMinder user account. Double-click the LDAP attribute and then append a conversion expression to the attribute name. Enclosed the expression in square brackets.

11. **Anchor Attribute screen**

    You can use a CA DataMinder account attribute to synchronize CA DataMinder users with LDAP users (or users in a data file). This CA DataMinder attribute is the *anchor attribute*. It can be the user name, the user full name, or any of the defined user attributes. Account Import uses the anchor attribute to establish a link between the target user account in CA DataMinder and the source user account in the LDAP directory (or data file). Account Import then updates the account details in CA DataMinder with the imported information (the user's parent group, e-mail addresses and other attributes).

    User name

    You specified the LDAP attribute mapped to CA DataMinder user names in the User Name Attribute field in step 3.

    Full name

    You specified the LDAP attribute mapped to CA DataMinder user full names was specified in step 10.

    Attribute index

    You specified the LDAP attributes mapped to CA DataMinder account attributes in step 10. Enter a value in the Attribute Index field, where index 1 refers UserAttribute1, index 2 to UserAttribute2, and so on.

User renames allowed

This check box prevents CA DataMinder user names being inadvertently overwritten if the CA DataMinder name is different to the value of the LDAP attribute (or XML user tag) that identifies the user name. For example, this can happen if a user recently got married.

To stop the user name in the CA DataMinder database being overwritten during a synchronization process, clear this check box.

Note: This check box is automatically selected and disabled if you choose to anchor the user synchronization on the user name. This is because the synchronization does not match against a CA DataMinder user unless the user names are the same.

12. **Import Assessment screen**

    Wait while Account Import identifies all the changes and additions that must be made to the CA DataMinder user hierarchy.

13. **Confirm Changes screen**

    (Only available if you selected the ‘Manual confirmation’ option in step 7.)

    Confirm or reject all of the changes to the existing user or machine hierarchies.

    Display Changes

    Click this button to view the proposed changes to the CA DataMinder user hierarchy. The changes can take several minutes to appear if the import operation involves substantial additions or changes to the user hierarchy. When the list of changes appears, you can accept or reject all of the changes.

    Email addresses can be deleted

    (Applicable only if importing from a command file. See step 1)

    If you select this check box, any *emaildelete* commands in the command file are executed during the import operation. (An emaildelete command deletes an email address associated with a CA DataMinder user.)

    Important! Use this parameter with caution! After running an emaildelete (see page 290) command, existing mail events may no longer be associated with the correct user.

14. **Importing screen**

    The wizard now has all the information it needs. Wait while it imports the user data and updates the CA DataMinder user hierarchy.

15. **Import Complete screen**

    Details about the import operation are recorded in a log file.

More information:

Delete an Email Address (see page 290)
About Single Sign-On (see page 164)
Modify LDAP Values with Conversion Expressions (see page 298)
Features That Use Email Address Mapping (see page 495)
Combining Multiple LDAP Attributes in Single CA DataMinder Attributes (see page 280)

Example Import Operation

The example below is based on the examples in Account Import wizard, steps **4** and **5**. It assumes that the Create New Users option is selected when you run the Account Import wizard.

1. The source LDAP directory is ‘ou=Sales’.

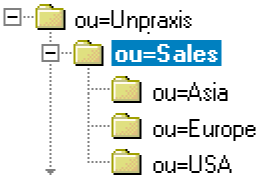

Example LDAP directory structure

2. The target CA DataMinder parent group is ‘Users’.

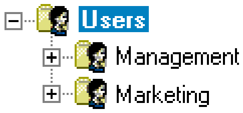

CA DataMinder user hierarchy: Before importing

3. The following changes are imported to the CA DataMinder user hierarchy:

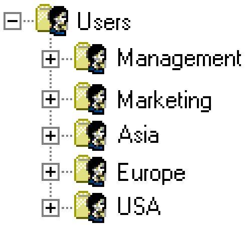

CA DataMinder user hierarchy: After importing

The original LDAP directory structure is preserved in the CA DataMinder user hierarchy.

Handling for Unknown Users

When you import users from an LDAP directory, you can specify how CA DataMinder handles 'unknown' users or groups. These are users and groups in your existing CA DataMinder user hierarchy that are not present in the LDAP directory or XML data file.

Specifically, you can move unknown users to an 'exceptions' group and you can exempt them from policy.

Exempt from policy

You can exempt unknown users from policy.

Exempt users are users who have a CA DataMinder account on the CMS but who are exempt from policy. That is, CA DataMinder does not monitor email, web or file activity for policy-exempt users.

'Exceptions' group

The exceptiopns This can be any existing group that falls within your management group. However, be aware that this reorganization only affects CA DataMinder users within the specified parent group. Unknown users outside of this branch of the user hierarchy are not reorganized.

The following example specifies the Management group as the destination for imported users. If required you can move any unknown users within this group or its subgroups to an exceptions group. However, any unknown users outside of this group or its subgroups are not reorganized. For example, unknown users in the Sales group are not reorganized into the exceptions group.

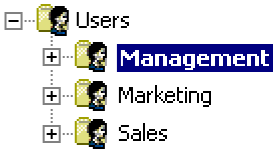

Example CA DataMinder user hierarchy

The Management group is the target parent group for imported users. Unknown users in the Marketing and Sales groups are not reorganized.

Command Line Import Operations

Account Import supports command line operations for importing user details from an LDAP directory or other external source. You can schedule regular command line operations to keep your LDAP directory and CA DataMinder user hierarchy synchronized.

You can import user details from an LDAP directory via a command line using Secure Sockets Layer (SSL).

Note: Command line import operations have been tested against Microsoft Active Directory and iPlanet 5.2 LDAP. Command line operations may import user details successfully from other LDAP compatible directories, but these have not been tested.

More information:

Set Up Secure Sockets Layer (SSL) (see page 265)

Command Syntax

The command line syntax for importing user details from an LDAP directory, or XML file is:

```
wgninfra -run wigan/infrastruct/accounts/AccountImport [<logfile>] <parameters>
```

Where:

wigan/infrastruct/accounts/AccountImport

Defines the required Java class for an LDAP import operation.

Note: The LDAPTrans command has been deprecated, but is still available for use. However, we recommend that you use the syntax shown above.

[<logfile>]

Optionally defines a non-default log file. See the following sections for details.

<parameters>

Defines one or more import parameters. Typically, a command only includes the /op parameter—this specifies an import configuration file, which contains all the other parameters required for the current import operation.

Example Command

The command below specifies that import parameters are defined in the my_parameters.opt parameter file and that import results are written to my_logfile.log.

```
wgninfra -run wigan/infrastruct/accounts/AccountImport my_logfile.log /op
my_parameters.opt
```

Cache the LDAP Logon Credentials

Before you begin ‘real’ command line import operations from an LDAP directory, we recommend that you run a single ‘actionless’ import operation with the /cc parameter, solely to cache your LDAP credentials. For example:

```
wgninfra -run wigan/infrastruct/accounts/AccountImport /cc /un UNIPRAXIS\srimmel /pw
abc123
```

This command securely caches the LDAP logon credentials for the user UNIPRAXIS\srimmel (the password is abc123). All subsequent LDAP import operations can use these cached credentials by invoking the /ec parameter.

Log Files

If the command does not specify a log file, the results of each import operation are written to a CA DataMinder log file. Alternatively, you can specify a custom log file. If:

- The log file already exists, log entries for the new import operation are appended to it.
- You do not specify a path, the log file is saved in CA's \data\log subfolder of the Windows All Users profile.
- You do not specify a file name extension for the log file, CA DataMinder appends a timestamp and .log extension to the file name, for example, mylog_200201200945.log

More information:

Account Import Log Files (see page 248)

Set Up Secure Sockets Layer (SSL)

CA DataMinder supports SSL when importing user details from an LDAP directory via a command line.

To set up SSL

1. Obtain a server authentication certificate that identifies your LDAP server and install it on your CMS.

   This certificate enables the CMS to recognize and trust the LDAP server.

   a. Browse to the \system\jre142_12\bin subfolder in the CA DataMinder installation folder on the CMS

   b. From a command prompt, run:

   ```
   keytool -import -keystore ..\lib\security\jssecacerts
   -alias <alias> -file <cert_file>
   ```

   Where <alias> is a relevant and descriptive name for the certificate and <cert_file> is the certificate file name.

2. Restart the CMS.
3. Edit the Account Import parameter file to include SSL support.

   Specifically, add the /us option (see page 269).

Parameter Files

A range of parameters are available for configuring the import operation. These are listed on the following sections.

For ease of maintenance, we strongly recommend that you use a parameter file to specify your import parameters—see the /op data source parameter for details.

- Comment markers (see page 267)

  `#, REM, rem, Rem`

- Source and target file parameters (see page 268)

  ```
  /ou <file name>
  /in <file name>
  ```

- Data source parameters (see page 268)

  ```
  /op
  /df
  /lp
  ```

- LDAP logon parameters (see page 269)

  ```
  /sv <server>
  /un <ldap user>
  /pw <ldap password>
  /cc
  /ec
  ```

- SSL support parameter (see page 269)

  ```
  /us
  ```

- LDAP filter attributes (see page 270)

  ```
  /uf <LDAP filter>
  /nf <LDAP filter>
  ```

- Source container and target group parameters

  ```
  /dn <Root DN>
  /lr <context>
  /wr [set the product group or family]
  /ga <LDAP attribute list>
  /ft
  /ce
  /me
  /eg <exceptions group>
  /pd <domain>
  /ml <LDAP attribute list>
  /ed
  /al <LDAP attribute list>
  ```

- Operation parameters (see page 274)

  ```
  /ca
  /re
  /at
  ```

- Error handling parameter (see page 274)

  ```
  /iw
  ```

- Exempt users parameters

  ```
  /ee
  /pe
  /pl
  ```

- User mapping and identification parameters (see page 276)

  ```
  /ua <LDAP attribute>
  /fn <LDAP full name attribute>
  /an <CA DataMinder account attribute>
  /nu
  ```

More information:

Parameter Rules (see page 267)
Example Parameter File (see page 278)

Parameter Rules

These rules apply to all relevant Account Import parameters:

Single quotes

If a parameter is set to a comma separated list of values (this typically applies to the /al and /ga parameters), you can enclose each list item in ‘single quotes’. This helps when the list item itself contains a comma—this typically happens when you need to specify a separator for multiple-value LDAP attributes. For example:

```
/al ‘proxyAddresses’,’Mail’
```

Here, all values in the multiple-value LDAP attribute proxyAddressses are assigned to UserAttribute1, where they are separated by commas, while the Mail LDAP attribute is assigned to UserAttribute2.

Prefix special characters with a backslash

Certain separator characters have a reserved meaning. If a parameter value includes quotes, commas and backslashes, you must prefix the characters with a backslash to ensure they are interpreted correctly. For example:

```
/lr "o=\"Unipraxis PLC\""
```

Here, the LDAP root container is o="Unipraxis PLC". Backslash prefixes are needed to ensure the double quotes are handled correctly.

Double quotes

You must enclose the entire parameter value in "double quotes" if the value contains a space. In the example below, the target group for imported users is ‘LDAP users’:

```
/wr "LDAP Users"
```

More information:

Command Syntax (see page 263)
Conversion Expression Syntax (see page 299)

Comment Markers

#, REM, rem, Rem

These formats identify comments in a parameter file:

```
# Your comment goes here
REM Your comment goes here
rem Your comment goes here
Rem Your comment goes here
```

Source and Target File Parameters

/ou <file name>

This parameter diverts user details extracted from an LDAP directory or data file to a specified text file. For example, you may want to keep a record of all changes to the CA DataMinder user hierarchy.

This parameter is typically used in association with the /in parameter (see below) to break import operations into two stages: 'Extraction from LDAP or data file' and 'Import into CA DataMinder'.

/in <file name>

This parameter specifies a Command File to use. This file can be created manually or by exporting in command file format—see the Administration console online help; search for 'users, hierarchy exporting'.

Note: If this parameter is specified, all other parameters are ignored.

More information:

Command Files to Import Users (see page 280)

Data Source Parameters

/op

Specifies a configuration file containing other parameters.

/df

Specifies a data file. These are structured files of user data, in XML or spreadsheet-compatible format, containing all or part of an external user hierarchy.

/lp

This parameter takes no value. If an import operation imports simultaneously from both an LDAP directory *and* a data file, this parameter specifies that the LDAP source takes precedence if duplicate users are detected; if this parameter is absent, the data file takes precedence.

More information:

Example Parameter File (see page 278)

LDAP Logon Parameters

/sv <server>

Identifies the server hosting the LDAP directory. You can specify the server name or IP address.

If the LDAP port number is not 389, you can prefix it with a colon after the server name, for example:

```
/sv UNI-HARDY-XP:319
```

/un <ldap user>

Specifies the user account used by CA DataMinder to access the LDAP database.

/pw <ldap password>

Specifies the password for the LDAP user specified by the /un parameter.

/cc

Caches in a secure location the credentials used to access the LDAP database. These are the credentials specified by /un and /pw.

The purpose of this parameter is to allow all subsequent import operations to log on to the LDAP database using these cached credentials. This eliminates the need to include sensitive account details in your import configuration file. See the /ec parameter below for details.

Before you begin ‘real’ LDAP import operations, we recommend that you run a single ‘actionless’ import operation with the /cc parameter, solely to cache your LDAP credentials.

/ec

Enables CA DataMinder to log on to the LDAP database using the cached credentials previously specified by the /cc parameter. This is the logon method we recommend.

More information:

Cache the LDAP Logon Credentials (see page 264)

SSL Support Parameter

/us

Adds SSL support when importing user details from an LDAP directory via a command line. To use SSL, you must first install a server authentication certificate.

More information:

Set Up Secure Sockets Layer (SSL) (see page 265)

LDAP Filter Attributes

If necessary, you can filter your import operations to only extract the users you want from the LDAP database. If you do specify specific object classes and categories, ensure that the new filter conforms to RFC 2254.

/uf <LDAP filter>

Specifies which RFC 2254 filter to use when extracting users from the LDAP database.

/nf <LDAP filter>

Specifies which RFC 2254 filter to use when extracting the LDAP containers (or nodes) that correspond to CA DataMinder user groups.

Source Container and Target Group Parameters

/dn <Root DN>

Specifies the LDAP server's base DN or domain. For example, to specify an Active Directory domain, enter one of these formats:

```
unipraxis.com
```

or

```
dc=unipraxis,dc=com
```

/lr <context>

Specifies the base context for import operations in the LDAP directory. All users and groups at and below this level will be copied into CA DataMinder. For example:

```
ou=dept,ou=room
```

/wr <group>

Specifies the target parent group in the CA DataMinder user hierarchy. All imported users and groups will be added to this parent group. You must specify the path to this group, relative to the root-level 'Users' group.

/ga <LDAP attribute list>

Derives a user's group from the LDAP attributes in this comma separated list. You can specify a single /ga parameter, set to a comma-separated list of LDAP attributes, or you can specify multiple instances of the /ga parameter, each set to a single LDAP attribute; the instances are processed in the order in which they occur in the command or configuration file. For example:

```
/ga division,department,team
```

Or

```
/ga division
/ga department
/ga team
```

/ft

Specifies that users imported into CA DataMinder will have a flat hierarchy. That is, new accounts for all imported users will be created in a single group. The target group is the group specified by the /wr parameter—see above.

/ce

The LDAP directory structure, or the structure specified in a data file, may contain empty containers. These may hold subcontainers or other items, but no users. This parameter creates corresponding empty user groups in CA DataMinder.

Note: To use the /ce parameter, the /ca parameter must also be set.

/me

If your existing CA DataMinder hierarchy contains users not present in the LDAP directory or data file, this parameter moves them to an 'exceptions' group, defined by the /eg parameter—see below.

If required, you can use this parameter in conjunction with /ee (see page 275) parameter to move unknown users *and* exempt them from policy.

Note: Users prepended with a domain name other than the one set in the /pd <domain> parameter (see below) are not moved.

/eg <exceptions group>

Used in association with /me. This parameter specifies the target 'exceptions' group. This can be any group in the CA DataMinder user hierarchy. You must specify the full path to the group, relative to the root-level 'Users' group. For example, this specifies the Users/Non-LDAP users subgroup:

```
/eg "Non-LDAP users"
```

If this parameter is omitted and /me is set, Account Import creates a default 'Exceptions' group, immediately below the root-level 'Users' group.

/pd <domain>

Prefixes new CA DataMinder user names with the specified domain name. You do not need to add a backslash. If the user names in the LDAP directory or data file do not have a domain prefix (that is, the user name does not contain a backslash), this setting will automatically add one.

/ml <LDAP attribute list>

Specifies which LDAP attributes are written to the email address table in CA DataMinder.

Important! You must also include the /at parameter, otherwise any /ml attributes you specify will not be written to CA DataMinder user accounts—see the /at parameter.

You can specify a single /ml parameter, set to a comma-separated list of LDAP attributes, or you can specify multiple instances of the /ml parameter, each set to a single LDAP attribute. For example:

```
/ml mail,proxyAddresses,legacyExchangeDN
```

Or

```
/ml mail
/ml proxyAddresses
/ml legacyExchangeDN
```

Important! For ease of maintenance, we strongly recommend you use multiple instances of /ml.

The /ml parameter also enables you to modify email addresses in the LDAP directory before writing them to the CMS database. To do this, you specify a conversion expression.

Note: If you use the ICAP agent to integrate with BlueCoat ProxySG servers, you must use the /ml parameter to import the distinguishedName attribute.

/ed

Specifies that *emaildelete* commands are carried out during the import or synchronization process. If you specify this parameter:

- For a synchronization process, any email address present in the CA DataMinder database but not in the data source is deleted from its associated CA DataMinder user.
- When importing from a command file, any *emaildelete* commands in the command file are executed.

If you do not specify this parameter (this is the default), *emaildelete* commands are ignored.

Use with caution

Important! Use this parameter with caution! If an emaildelete (see page 290) command removes an email address from a user's address list, any events associated with the deleted email address are no longer associated with that user.

If you use the /ed parameter to clean up a misconfigured import operation, be aware that valid email addresses may also be removed. Instead, you may prefer to remove problematic email addresses using the Administration console or a manually produced command file. Use individual *emaildelete* commands to specify the user and associated email address that you want to remove.

/al <LDAP attribute list>

Specifies which LDAP attributes are written to account attributes of CA DataMinder users.

Important! You must also include the /at parameter, otherwise any /al attributes you specify will not be written to CA DataMinder user accounts—see /at.

You can specify a single /al parameter, set to a comma-separated list of LDAP attributes, or you can specify multiple instances of the /al parameter, each set to a single LDAP attribute; the instances are processed in the order in which they occur in the parameter file. For example:

```
/al division,employeeID,rank
```

Or

```
/al division
/al employeeID
/al rank
```

Important! For ease of maintenance, we strongly recommend you use multiple instances of /al.

LDAP attributes are assigned to CA DataMinder account attributes in the order in which they occur. That is, the first LDAP attribute is assigned to UserAttribute1, the second to UserAttribute2, and so on. In both examples above, the LDAP attribute Rank is assigned to UserAttribute3.

The /al parameter also enables you to:

- **Combine multiple attribute values** and write them to a single CA DataMinder user attribute. For example, the LDAP directory may contain three attributes, Building, Floor and Desk number. To combine these attributes into a single value and write it to a single CA DataMinder user attribute, the syntax is shown below. You can choose which separator to use; the syntax is:

  ```
  /al <attribute1><SV separator><attribute2><SV separator><attribute3>
  ```

  For example:

  ```
  /al Building,Floor,DeskNumber
  ```

- **Modify LDAP attribute values** before writing them to an attribute of a CA DataMinder user account. To do this, you specify a conversion expression.

More information:

Conversion Expression Syntax (see page 299)
Delete an Email Address (see page 290)
Combining Multiple LDAP Attributes in Single CA DataMinder Attributes (see page 280)

Operation Parameters

Important! Each command must specify at least one of the following parameters:

/ca

Specifies that CA DataMinder creates new user accounts for unknown users. That is, it creates a new account for any user in the LDAP directory or data file who has no corresponding account in CA DataMinder. Where necessary, it also creates new groups to hold these new users, based on the equivalent containers in the LDAP directory.

/re

Reorganizes the existing hierarchy of CA DataMinder users to synchronize it with the hierarchical structure of the LDAP directory or data file.

/at

Updates existing CA DataMinder user accounts with attributes imported from corresponding users in the LDAP directory or data file.

Important! You must include the /at parameter if you want to specify attributes using the /al and /ml parameters.

Error Handling Parameter

/iw

When you use a command line to synchronize users' e-mail addresses in the CMS with users in an XML file, Account Import will stop the import operation if a user's management group in the XML file does not exist in the CA DataMinder hierarchy.

This parameter enables Account Import to continue with the import operation under these circumstances by ignoring any warnings. If this parameter is not set, Account Import will stop the import operation under these circumstances.

Note: This parameter is ignored if either the /in or /ou parameter is specified.

More information:

Source and Target File Parameters (see page 268)

Policy Exemption Parameters

You can configure your Account Import operations to exempt users from policy based on LDAP or XML attributes.

/ee

Exempts unknown users from policy. If your existing CA DataMinder user hierarchy contains users or groups not present in the LDAP directory or XML data file, you can exempt them from policy.

Exempt users are users who have a CA DataMinder account on the CMS but who are exempt from policy. That is, CA DataMinder does not monitor email, web or file activity for policy-exempt users.

Note: If required, you can use this parameter in conjunction with /me (see page 270) parameter to exempt unknown users from policy *and* move them to an 'exceptions' group.

/pe

If you are importing user details:

- Directly from an LDAP directory, this parameter specifies that CA DataMinder user accounts are exempt from policy if the corresponding LDAP user account has a specific attribute or attribute value. You must use this parameter in conjunction with the /pl parameter.
- From an XML data file, this parameter specifies that CA DataMinder user accounts are exempt from policy if the *policyexempt* attribute in the <user> tag is 'True'.

/pl <attribute name>=<attribute value>

(Only required when importing user details from an LDAP directory)

Specifies the LDAP attribute or attribute value which, if present, exempts CA DataMinder user accounts from policy. You must use this parameter in conjunction with the /pe parameter.

The *=<attribute value>* is optional. Do not add spaces around the '=' separator. If no value is specified, the attribute only needs to be populated in order to exempt the CA DataMinder user from policy. For example:

```
/pl extensionAttribute1
/pl Office=London
```

The first example exempts all users from policy who have extensionAttribute1 set. For example, your organization may use this LDAP attribute to identify a specific team of users who are not subject to CA DataMinder control.

The second example exempts all users based in the London office. For example, your CA DataMinder license may only cover employees in your US offices.

More information:

Conversion Expression Syntax (see page 299)
Delete an Email Address (see page 290)
Combining Multiple LDAP Attributes in Single CA DataMinder Attributes (see page 280)

User Mapping and Identification Parameters

/ua <LDAP attribute>

Specifies the user name attribute. Use this parameter if you need to specify a custom or non-standard LDAP attribute.

If this parameter is omitted, CA DataMinder automatically detects the type of LDAP directory (for example, Microsoft Active Directory) and key details about the LDAP directory structure. It provides a 'best guess' when selecting the LDAP user name attribute, for example, sAMAccountName.

/fn <LDAP full name attribute>

Specifies the LDAP attribute that contains each user's full name.

If this parameter is omitted, CA DataMinder automatically detects the type of LDAP directory (for example, Microsoft Active Directory) and provides a 'best guess' when selecting the LDAP full name attribute, for example, DisplayName.

/an <CA DataMinder account attribute>

Specifies which CA DataMinder account attribute to use as the anchor for mapping LDAP (or data file) users to CA DataMinder users. This can be the user name, the user display name, or any of the ten user attributes. Use one of these keywords:

/an username

Uses the LDAP attribute specified by the /ua parameter.

/an fullname

Uses the LDAP attribute specified by the /fn parameter.

/an attribute1

Uses the first LDAP attribute specified by the /al parameter.

/an attribute2

Uses the second LDAP attribute specified by the /al parameter.

And so on.

CA DataMinder uses the specified user attribute to locate the corresponding user in the LDAP directory.

Anchor requirements

The requirements for the CA DataMinder attribute used to anchor user import operations are as follows:

- Each user in your CA DataMinder enterprise must have a unique attribute value.
- The attribute values must not have been modified (using a conversion expression). That is, it must match exactly the corresponding attribute value in the LDAP database.
- The attribute cannot be a multiple value attribute.

/nu

When carrying out a synchronization process, it is possible that the user name in the CA DataMinder database is different to the value of the XML <user> tag or LDAP attribute used for the user name. For example, if a user has recently married.

To stop the user name in the CA DataMinder database being overwritten, add this parameter to the command line.

More information:

Importing a Single LDAP Attribute with Multiple Values (see page 279)

Example Parameter File

The following is an example parameter file for a command line import operation from an LDAP directory. In this example, users in the LDAP directory are all stored in a single container, ou=employees; Account Import derives their position in the LDAP hierarchy by concatenating these LDAP attributes: Division, Department and Team. This example also imports e-mail addresses and various LDAP attributes. Account Import writes these details to each CA DataMinder user account.

Example parameters	Resulting action

`/ca`
`/re`
`/at`
	These define the import actions: create new CA DataMinder user accounts; reorganize the existing CA DataMinder user hierarchy; and update the account attributes for CA DataMinder users.

`/sv UNI-HARDY-XP`
`/ec`
	These specify the LDAP host server and use previously cached credentials to log on to LDAP.

`/dn unipraxis.com`
`/lr ou=employees`
	These specify the LDAP root (or source container) for the import operation.

`/me`
`/eg "Non-LDAP users"`
	Existing CA DataMinder users not present in the LDAP directory will be moved to an exceptions group called 'Non-LDAP users'. Note the use of double quotes.

`/pd UNIPRAXIS`	User names for new CA DataMinder accounts will be prefixed with 'UNIPRAXIS\'. The backslash is added automatically.
`/ua attribute6`	Defines the 'anchor' LDAP attribute. This example specifies CA DataMinder user attribute 6.
`/ga Division,Department,Team`	Defines the target group for imported users, based on hierarchy from this list of LDAP attributes.

`/ml mail`
`/ml proxyAddresses`
`/ml legacyExchangeDN`
	Updates CA DataMinder email address attributes.

`/al EmployeeID`
`/al Department`
`/al Team,,Rank`
	Updates CA DataMinder user attributes. These /al instances write LDAP attributes to UserAttribute1 through UserAttribute3 and to UserAttribute5.
	Note: You cannot include 'empty' /al parameters; to specify non-consecutive user attributes, you must use comma delimiters.

More information:

Importing a Single LDAP Attribute with Multiple Values (see page 279)

Multiple Attribute Values

Individual attributes in both CA DataMinder and LDAP directories can contain multiple values. This has implications for user import operations from an LDAP directory.

More information:

Importing a Single LDAP Attribute with Multiple Values (see page 279)
Combining Multiple LDAP Attributes in Single CA DataMinder Attributes (see page 280)

Importing a Single LDAP Attribute with Multiple Values

Individual attributes in LDAP directories or data files can contain multiple values. For example, a MemberOf attribute can contain all the mail groups or email distribution lists that a user belongs to.

Account Import automatically writes multiple LDAP values to multiple values of a specific CA DataMinder user attribute. For example, you can import all the MemberOf distribution lists that a CA DataMinder user belongs to as separate values for an attribute renamed to Email Distribution Lists (example **2** in the screenshot below).

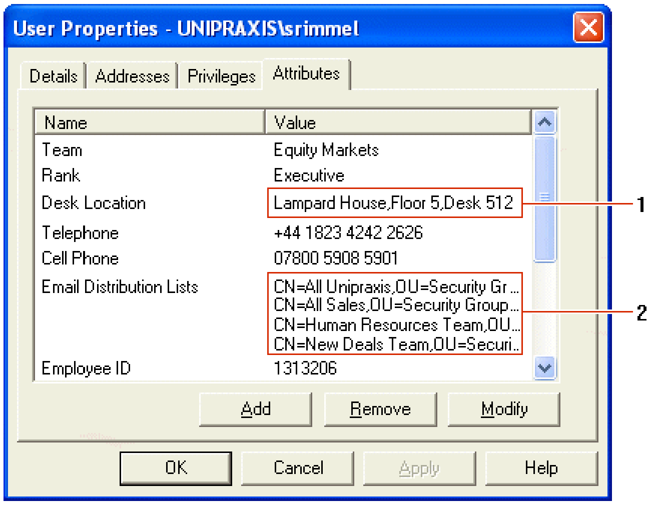

User Properties - UNIPRAXIS\srimmel

Name	Value
Team	Equity Markets
Rank	Executive
Desk Location	Lampard House,Floor 5,Desk 512
Telephone	+44 1823 4242 2626
Cell Phone	07800 5908 5901

Email Distribution Lists	CN=All Unipraxis,OU=Security Gr...
	CN=All Sales,OU=Security Group...
	CN=Human Resources Team,OU...
	CN=New Deals Team,OU=Securi...

Employee ID	1313206

User Properties dialog, Attributes tab example

1 Multiple LDAP attributes combined into a single value for a single CA DataMinder attribute.

2 Multiple values for a single LDAP attribute, written to separate values for a single CA DataMinder attribute.

Combining Multiple LDAP Attributes in Single CA DataMinder Attributes

If necessary, Account Import can write multiple LDAP attributes to a single attribute of a CA DataMinder user account. For example, the LDAP directory may contain three attributes, Building, Floor and Desk number. Using Account Import, you can combine these attributes into a single value and write this value to a single CA DataMinder user attribute renamed to Desk Location. See example **1** in the screenshot below.

To do this using the Account Import wizard, see step **10** in Account Import wizard for an example.

To do this in a command line operation, use the /al <LDAP attribute list> parameter.

Note: In all cases, when assigning lists of LDAP attributes to a single import parameter, remember the parameter rules.

Command Files to Import Users

Account Import also lets you import user details from a command file (a tailored CSV file). These are import configuration files containing CA DataMinder user import commands (for example, 'create new user' or 'set user attribute'). Typically, you import command files to make specific changes or additions to your existing CA DataMinder user hierarchy.

Note: If required, you can even combine user import and machine import operations within a single command file.

More information:

Import from a Command File (see page 295)

Command File Format - Users

Before you can import a command file containing user details, you must ensure that each record in the file conforms to the format required by Account Import. The correct formats for individual records are shown in the tables below and on the following page.

Each user record begins on a new line with a variable that defines the type of operation. Typical operations include adding new users or groups, moving a user or group to a new parent group, specifying a user's role, password or management group, and updating a user’s attributes. You can also include comments within the file, for example, to organize the file into sections.

The order of operations is important. Within the command file, records to create new groups must precede records to add new users to these groups. Likewise, a record to move a group must precede any records to move users into the relocated group.

If you plan to export your user details (say, from your mail server) to a command file, you could write a script or macro to prefix each user record with the appropriate operation variable.

More information:

Example Command file 1: New users and groups (see page 292)
Example Command File 2: User Properties (see page 293)
Format Notes (see page 281)

Format Notes

When creating records in a command file, note the following:

<address>

A user’s email address. This can be any format, for example:

SMTP: srimmel@unipraxis.com

EX: /o=unipraxis/ou=uk/cn=spencer/cn=rimmel

Domino: cn=spencer rimmel/o=unipraxis

A command file can include multiple instances of this command for each user.

<attribute value>

This is the actual number or text associated with the current attribute. For example, the value for an employee ID attribute might be rimspe01. Enclose the text in double quotes if it includes spaces.

attribute*n*

These correspond to the CA DataMinder attributes listed in the User Properties dialog. So attribute1 updates the value for the first attribute listed; attribute2 updates the value for the second listed attribute, and so on.

If you omit any numbered attributes from the CSV file, the corresponding attribute values remain unchanged for the CA DataMinder user.

You can also reference an attribute using the following syntax:

```
attribute:<attribute name>
attribute <attribute name>
```

Note: A single space character is used to separate the 'attribute' keyword from the attribute name.

<current group>

This is the path of the group you want to move. The path is delimited by forward slashes and is relative to the current user’s management group.

<full user name>

The display name that appears in the Summary tab. Enclose the name in double quotes if it includes spaces, for example, "Lynda Steel".

<group name>

This is the name of the new group. No path is needed, for example, Marketing.

<management group>

This a management group assigned to the specified user. Users can have multiple management groups. The path is delimited by forward slashes and is relative to the management group of the administrator running Account Import.

Note: The group must already exist. You cannot assign management groups that fall outside of your own management group.

<parent group>

This is the path of the group into which you want to add or move a user or another group. The path is delimited by forward slashes and is relative to the current user’s management group. See the example group paths and group name requirements in the following sections.

Important! You need to be aware that moving groups can cause security issues and unintended changes to policy—see the Administration console online help; search for ‘groups, move users between groups’.

<password>

The new password assigned to the specified user.

<role>

This is the role assigned to the CA DataMinder user. You can specify the role name or number. The default roles are:

Administrator or 1

Manager or 2

User or 3

"Policy Administrator" or 4

Reviewer or 5

UserRole1 or 6

UserRole2 or 7

Role names are *not* case-sensitive, but you do need to enclose the role in double quotes if the role name contains spaces.

Also, the administrative privileges granted to each imported user comprise only those privileges common to both the specified role and the user running the wizard. In effect, the user running the wizard can only grant privileges which they hold themselves:

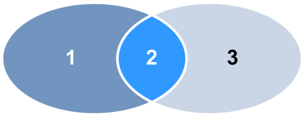

1 The default administrative privileges assigned to a role.

2 Privileges granted to the imported user.

3 Privileges currently granted to the user running the wizard.

<user name>

The name used by the CA DataMinder user account.

If your CMS uses Microsoft Windows user authentication to automatically generate new user accounts, this user name must be prefixed with the domain and a backslash separator, for example, unipraxis\srimmel.

More information:

Group and User Name Requirements (see page 291)

Set the Default Group

Specify any existing group to act as the default group. The group name is not case-sensitive.

```
defaultgroup,<group name>
```

Create a New Group

Specify the new group and any existing group to act as its parent. The group name is not case-sensitive.

```
newgroup,<group name>,<parent group>
```

Delete a Group

This command removes a specified group from the database. The format is:

```
deletegroup,<group name>,<parent group>
```

Import a New User

Specify the user name and a parent group. The parent group must already exist. It is not case-sensitive.

```
newuser,<user name>,<parent group>
```

Move a Group to a New Parent Group

Specify the current group and any existing group to act as its new parent. Note the groups are not case-sensitive. This operation will also move any subgroups.

```
movegroup,<current group>,<parent group>
```

Important! You need to be aware that moving groups can cause security issues and unintended changes to policy—see the Administration console online help; search the index for 'groups, moving users'.

Move a User to a New Group

Specify the user name and their new parent group. The parent group must already exist. It is not case-sensitive.

```
moveuser,<user name>,<parent group>
```

Specify a User's Full Name

Specify the user name (that is, the name used by the user account) and the user's full name. The format is:

```
fullname,<user name>,<full user name>
```

Rename a User

This command allows you to reset the name of the user account, that is, the user name. The format is:

```
renameuser,<user name>,<new user name>
```

Exempt a User

This command exempts the user account from policy. The format is:

```
policyexempt,<user name>,true
```

This command ensures that policy *is* applied to a user account:

```
policyexempt,<user name>,false
```

Delete a User

This command removes a specified user. The format is:

```
deleteuser,<user name>
```

Set a User Password

All users must a supply a password the first time they use the CA DataMinder consoles or certain CA DataMinder utilities. Account Import therefore lets you assign a temporary password to a new user. For security reasons, the new user is then forced to change this password to something more private when they first log on to CA DataMinder.

To specify a temporary password for a new user, the format is:

```
password,<user name>,<password>
```

Reset a User Password

This command allows you to reset a user’s password but does not force that user to change their password when they next log on to CA DataMinder.

To reset a user’s password, the format is:

```
password_noexpiry,<user name>,<password>
```

Specify a User’s Role

This command specifies a user’s role, which in turn defines their default administrative privileges.

```
setrole,<user name>,<role>
```

Specify a New Management Group

This command assigns a single management group to a user. It **replaces** any existing management groups with the new one. The format is:

```
managegroup,<user name>,<management group>
```

Note: If the management group name is left blank, null, or not specified, then the user is assigned the group selected as the User Tree Root. For example:

```
managegroup,"Unipraxis\srimmel",
```

```
managegroup,"Unipraxis\srimmel"
```

```
managegroup,"Unipraxis\srimmel",""
```

```
managegroup,"Unipraxis\srimmel","/"
```

Specify a User's Security Model

This command specifies the security model of a user.

```
securitymodel,<user name>,<security model>
```

Every user must have a security model, you cannot delete it. You can reset the securitymodel of a user to the default using the following command:

```
securitymodel,<user>,MD
```

The command uses the following parameter:

security model

Defines the short name form of the specified model. The Short Names for the available Full Model Name are:

MS

Specifies Management Group (Sender)

MSX

Specifies Management Group (Sender, Self-Exclude)

MD

Specifies Management Group (Standard)

MDX

Specifies Management Group (Standard, Self-Exclude)

PD

Specifies Policy (Standard)

PDX

Specifies Policy (Standard Self-Exclude)

AD

Specifies Unrestricted

Note: You can set a hybrid model by supplying two model names in double quotes, separated by commas.

Example

Use the following command to set a hybrid model of 'Policy (Standard)' and 'Management Group (Sender)':

```
securitymodel,<user>,"PD,MS"
```

Specify a New Policy Role

This command sets a user's policy role. It overwrites any policy roles that are currently specified.

```
policyrole,<user>,<policy role>
```

Delete a Policy Role

This command deletes one of a user's existing policy roles.

```
delpolicyrole,<user>,<policy role>
```

Specify a User's Reporting Name

This command specifies the name used by the BusinessObjects integration to refer to the currently logged on user. The iConsole uses this user name to connect automatically to BusinessObjects, so the iConsole user does no longer need to to enter their BusinessObjects account details manually.

```
reportuser,<user>,<report name>
```

Add a Management Group

This command assigns a management group to a user. It does not replace any existing management groups already assigned. The format is:

```
addmgmtgroup,<user name>,<management group>
```

Note: If the management group name is left blank, null, or not specified (see managegroup for examples), then the user is assigned the group selected as the User Tree Root.

Delete a Management Group

This command removes a specified management group from a user. The format is:

```
delmgmtgroup,<user name>,<management group>
```

Note: If the management group name is left blank, null, or not specified (see managegroup for examples), then the user is assigned the group selected as the User Tree Root.

Add a User Attribute Value

This command adds a value for a specified attribute. Specify the attribute number, user name, and the attribute value you want to add.

```
Attribute1,<user name>,<attribute value>
Attribute2,<user name>,<attribute value>
Attribute3,<user name>,<attribute value>
```

And so on.

Note: You use this command to **add** multiple values to user attributes. The command does **not** overwrite existing attribute values. If you want to modify existing values, use setattribute (see page 289).

If you have renamed a user attribute, you can specify the attribute by its new name (in one of three ways), or its original attribute number when using this command. For example, if Attribute2 is set to 'Gender', use any of the following methods:

```
Attribute2,<srimmel>,<Male>
Attribute:Gender,<srimmel>,<Male>
Attribute Gender,<srimmel>,<Male>
Attribute "Gender",<srimmel>,<Male>
```

Set a User Attribute to a Single Value

This command removes the previous value of the attribute and replaces it with a new value. If no value is currently set, then this command will set one. Specify the attribute number, user name and the attribute value you want to set.

```
setattribute1,<user name>,<attribute value>
setattribute2,<user name>,<attribute value>
setattribute3,<user name>,<attribute value>
```

And so on.

This command overwrites **all** existing values for an attribute.

If you have renamed a user attribute, you can specify the attribute by its new name (in one of three ways), or its original attribute number when using this command. For example, use any of the following methods:

```
setattribute2,<lsteel>,<Female>
setattribute:gender,<lsteel>,<Female>
setattribute gender,<lsteel>,<Female>
setattribute "gender",<lsteel>,<Female>
```

Delete a User Attribute Value

This command deletes a specific value from a specified attribute. Specify the attribute number, user name and the attribute value you want to delete.

```
delattr1,<user name>,<attribute value>
delattr2,<user name>,<attribute value>
delattr3,<user name>,<attribute value>
```

And so on

If you have renamed a user attribute, you can specify the attribute by its new name (in one of three ways), or its original attribute number when using this command.
For example, use any of the following methods:

```
delattr3,<srimmel>,<New York>
delattr:location,<srimmel>,<New York>
delattr location,<srimmel>,<New York>
delattr "location",<srimmel>,<New York>
```

Add an Email Address

This command adds an email address for a user:

```
emailaddress,<user name>,<address>
```

Delete an Email Address

This command deletes an email address associated with a user:

```
emaildelete,<user name>,<address>
```

You typically use this command in conjunction with the /ed parameter in a parameter file.

Important! Use this command with caution! If an *emaildelete* command removes an email address from a user’s address list, any events associated with the deleted email address are no longer associated with that user.

Add Comments

These formats identify comments in a CSV file:

```
# Your comment goes here
// Your comment goes here
REM Your comment goes here
```

REM is not case-sensitive, so REM, rem or Rem are equally acceptable as comment markers.

More information:

Format Notes (see page 281)

Group and User Name Requirements

When specifying group and user names:

- If the user or group name contains spaces, you must enclose the name in "double quotes".

  If you use double quotes,

- If any groups or users have names that contain Far Eastern characters, make sure your CSV file is saved in a Unicode format to ensure that these characters are preserved during import.

Important! If you use the newuser, newgroup, defaultgroup, moveuser or movegroup variables, you must delete any trailing spaces after the group name! If any are present, the import operation will fail.

Example Group Path

The correct path specification for a user group depends on the actual user hierarchy in the Administration Console and the management group for the user running Account Import. For example, if the full path to locate the /North group in the group hierarchy is:

```
Users/Sales/Asia/Hong Kong
```

And your management group is /Sales, then you must use the following paths to add the user lsteel to these groups:

Target	**Use this group path**
Sales	newuser,lsteel[,] No group is needed because the wizard defaults to the management group. The trailing comma is optional.
Asia	newuser,lsteel,Asia
Hong Kong	newuser,lsteel,"Asia/Hong Kong"

Example Command file 1: New users and groups

This command file creates four new groups and two new users. It also reorganizes the group hierarchy. In this example, the administrator running the Account Import wizard has a management group set to the top-level Users group. The wizard interprets all group paths specified in the command file as being relative to this management group.

Example Command File

```
1    newgroup,Directors
2    newgroup,"South Asia"
3    newgroup,Europe,Sales
4    newgroup,Legal,Corporate
5    newuser,unipraxis\lyndasteel,Directors
6    newuser,unipraxis\spencerrimmel,Directors
7    # Move users and groups to a new parent group
8    movegroup,Directors,Corporate
9    movegroup,"South Asia",Sales
10   moveuser,unipraxis\omarabassi,Directors
11   moveuser,unipraxis\frankschaeffer,Directors
```

Resulting Actions:

Lines 1-4: Adds four new groups. Note that because no parent group is specified for the new Directors or South Asia groups, it defaults to the management group of the user running the wizard.

Lines 5-6: Adds two new users to the newly imported Directors group.

Line 7: Adds a comment

Lines 8-9: Moves the Directors group into the Corporate group, and Asia into the Sales group.

Lines 10-11: Moves existing users into the Directors group. You do not specify the users' current group.

Import Results for Example Command File 1

Before importing:	After importing
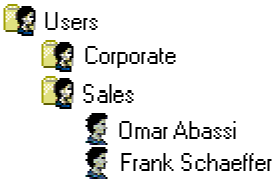 Users Corporate Sales Omar Abassi Frank Schaeffer	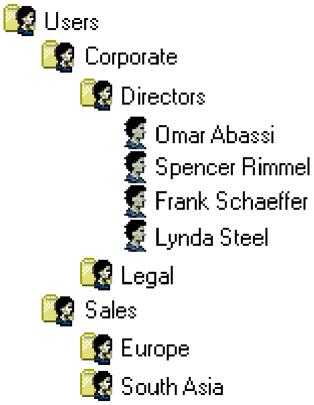 Users Corporate Directors Omar Abassi Spencer Rimmel Frank Schaeffer Lynda Steel Legal Sales Europe South Asia

Note: In this example, the administrator running the wizard has a management group set to the top-level Users group.

Example Command File 2: User Properties

This command file sets various properties for the new users created in Example Command file 1: new users and groups. As before, the administrator running Account Import has a management group set to the top-level Users group.

Note: This command file sets the properties for two new users, but command files can equally be used to set the properties of existing users.

Example Command File

```
fullname,unipraxis\lyndasteel,"Lynda Steel"
fullname,unipraxis\spencerrimmel,"Spencer Rimmel"
setrole,unipraxis\lyndasteel,manager
setrole,unipraxis\spencerrimmel,3
// Set the password and management group
password,unipraxis\spencerrimmel,rimmel
password,unipraxis\lyndasteel,steel
managegroup,unipraxis\lyndasteel,Corporate
addmgmtgroup,unipraxis\lyndasteel,Sales
emailaddress,unipraxis\spencerrimmel,srimmel@unipraxis.com
```

```
11    attribute1,unipraxis\spencerrimmel,131026
12    attribute2,unipraxis\spencerrimmel,0182 3367 0832
13    attribute3,unipraxis\spencerrimmel,Phoebe Rimmel
```

Resulting Actions@

Lines 1-2: Sets the full names for the two new users.

Lines 3-4: Sets the user roles. Note that 3 corresponds to ‘user’.

Line 5: Adds a comment

Lines 6-7: Sets the management groups for the named user.

Lines 8-9: Moves existing users into the Directors group. You do not specify the users’ current group.

Lines 10: Adds an email address.

Lines 11-13: Imports values for user attributes. These correspond to the first three attributes listed in the Options dialog.

Import Results for Example Command File 2

New user:	Lynda Steel	Spencer Rimmel
Role:	Manager	User
Temporary password:	Steel	Rimmel
Management groups	Users\Corporate User\Sales	None
Email address:	Not specified	srimmel@unipraxis.com
User attributes:	Not specified	Attribute 1 (Employee ID): 131026 Attribute 2 (Home telephone): 0182 3367 0832 Attribute 3 (Emergency contact): Phoebe Rimmel

More information:

Example Command file 1: New users and groups (see page 292)

Command Files to Import Machines

To simplify mass deployments, you can bulk create new machine accounts and pre-assign client machines to parent servers in advance of the CA DataMinder rollout. This enables you to deploy multiple client machines using a single source image (which identifies a single parent server) while ensuring that each client machine automatically connects to its 'correct' parent server immediately after installation.

To bulk create new accounts, you import the gateway and client or utility machine details from a command file. You can do this using the Account Import wizard or you can run command line import operations.

More information:

Command Line Import Operations (see page 263)

Import from a Command File

To bulk create new accounts, you import the gateway and client or utility machine details from a command file. You can do this using the Account Import wizard (step 1) or you can run a command line import operation.

More information:

Command Line Import Operations (see page 263)

Command File Format

When you import a command file containing machine details, each record in the file must conform to the format required by the wizard. Each user record must begin on a new line with a variable that defines the type of operation. The supported formats are listed below.

Command Syntax

- **Create a new gateway**

  ```
  newgateway,<gateway>,<parent server>
  ```

  Where <gateway> is the name of the new gateway and <parent server> is its parent server, either the CMS or another gateway.

- **Create a new client machine**

  ```
  newclient,<client name>,<parent server>
  ```

  Where <client name> is the name of the new client machine and <parent server> is its parent server, either the CMS or a gateway.

- **Create a new utility**

  ```
  newutility,<utility name>,<parent server>
  ```

  Where <utility> is the name of the new utility machine and <parent server> is its parent server, either the CMS or a gateway.

- **Delete a machine**

  This command deletes a specified machine. That is, a gateway, client, or utility machine:

  ```
  deletemachine,<machine name>
  ```

- **Move existing client machine or gateway**

  ```
  movemachine,<machine name>,<parent server>
  ```

  Where <machine name> is the name of an existing client or utility machine, or gateway and <parent server> is its new parent server, either the CMS or a gateway.

- **Add comments**

  These formats identify comments in a CSV file:

  ```
  # Your comment goes here
  // Your comment goes here
  REM Your comment goes here
  ```

  REM is not case-sensitive, so REM, rem or Rem are equally acceptable as comment markers.

More information:

Example Command file: Machine import (see page 297)

Example Command file: Machine import

Example CSV Command File

```
1    REM Add new gateways
2    newgateway,GW-MILAN,CMS-HARDY
3    newgateway,GW-NAPOLI,CMS-HARDY
4    REM Add clients to new gateways
5    newclient,UNI-TAYLOR,GW-MILAN
6    newclient,UNI-KEEGAN,GW-MILAN
7    newclient,UNI-VENABLES,GW-NAPOLI
8    newclient,UNI-HODDLE,GW-NAPOLI
9    REM Add new secondary gateway
10   newgateway,GW-ROMA,GW-MILAN
11   newclient,UNI-RAMSEY,GW-ROMA
12   newclient,UNI-REVIE,GW-ROMA
13   movemachine,UNI-ROBSON,GW-NAPOLI
```

Resulting Actions:

Lines 1: Adds a comment.

Lines 2-3: Adds new GW-MILAN and GW-NAPOLI gateways, parented to the CMS.

Line 4: Adds a comment.

Lines 5-8: Adds new clients, parented to the new gateways.

Line 9: Adds a comment.

Line 10: Adds new GW-ROMA gateway, parented to GW-MILAN.

Lines 11-12: Adds new clients, parented to GW-ROMA.

Line 13: Moves existing client machine to a new parent gateway.

Import Results for Example Command File

Before importing:	After importing:
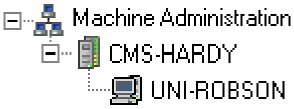	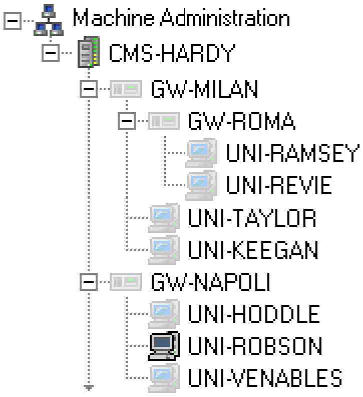 Machine Administration CMS-HARDY GW-MILAN GW-ROMA UNI-RAMSEY UNI-REVIE UNI-TAYLOR UNI-KEEGAN GW-NAPOLI UNI-HODDLE UNI-ROBSON UNI-VENABLES
	The newly imported gateways and client machines are initially represented by 'disconnected' icons. After CA DataMinder has been installed on these machines, they will revert to 'normal' icons.

Modify LDAP Values with Conversion Expressions

A key feature of command line import operations from LDAP is the ability to modify user attribute values in the LDAP directory before writing them to an attribute of a CA DataMinder user account. Conversion expressions can parse, extract and (if necessary) remove or substitute any characters in the attribute value.

To modify imported LDAP values, you configure the /al or /ml parameter to use a conversion expression. For example, if importing email addresses from the proxyAddresses LDAP attribute, email addresses are typically prefixed with a format identifier (such as 'smtp:' or 'x500:'). Account Import uses default conversion expressions to strip out these identifiers so that only the address is imported into CA DataMinder:

```
["smtp:{%untilEnd%}"]
["x400:{%untilEnd%}"]
["x500:{%untilEnd%}"]
```

More information:

Conversion Expression Syntax (see page 299)
Section Syntax (see page 300)
Conversion Expression Variables (see page 300)

Conversion Expression Syntax

Conversion expressions for imported LDAP values can comprise one or more sections. For example:

```
/al <attribute><[section1][section2][section3]>
```

where each [section] represents an element in the final CA DataMinder attribute value.

For example, in the above expression, the sections may resolve like this for a CA DataMinder user, Lynda Steel:

`[section1]`

Resolves to \o=Unipraxis

`[section2]`

Resolves to \ou=Paris

`[section3]`

Resolves to \cn=Lynda.Steel

Example Conversion Expression

In a fictitious LDAP directory for Unipraxis, values for the office attribute may have a ‘UX-’ prefix, such as ‘UX-Boston’ and ‘UX-New York’. The example conversion expression below strips out the prefix in the resulting CA DataMinder user groups (that is, ‘Boston and ‘New York)’:

```
/al office[if"UX-"["???{%untilEnd%}"] else["{?%untilEnd%}"]]
```

Separators between Sections

You cannot add literal separators between sections. Instead, you must add a custom ‘separator expression’ to insert the required character between sections. The example expression below inserts a single semicolon:

```
["zzzz",default=;]
```

Note: This method relies on the specified search text (in this case "zzzz") not being found; the expression then inserts the ‘default’ text (in this case, a semicolon).

To insert a backslash between sections, a modification is needed. This is because backslash characters have a special meaning in these conversion expressions. To insert a literal backslash, use this expression:

```
["zzzz",default=\\\\]
```

More information:

Section Syntax (see page 300)

Section Syntax

Each [section] defines the text that you want to extract from a single, specified component of the Exchange address.

```
["<search_text>{<extracted_text>}",
  repeat=N,
  subsection=A:B,
  substitute=<from>:<to>,
  mandatory,
  default=<term>]
```

where:

- Square brackets [ ] delimit the section. These can also be used to nest expressions, for example, to create multiple search sections:

  ```
  [["srimmel"],["@unipraxis"]subsection=2:,substitute=srimmel@unipraxis:spencer
  ]
  ```

  This example uses square brackets to define two search terms, 'srimmel' and '@unipraxis'.
- "Double quotes" delimit the search expression. For each section, this represents the text that you want to extract from the email address and incorporate into the CA DataMinder user name.
- Curly brackets { } delimit the text actually extracted from the email address (the ‘extracted text’).
- Supported variables <search_text> and <extracted_text> jointly define which part of an email address supplies the extracted text.
- Supported parameters 'repeat', 'subsection', 'substitute', 'mandatory' and 'default' allow you to precisely locate and, if necessary, modify the extracted text.
- Variables are processed in the order they are specified in the expression.

More information:

Conversion Expression Variables (see page 300)

Conversion Expression Variables

<search_text>

Is the actual text you want to detect within the LDAP attribute value. For example: `smtp:`

You use the search text to locate the start of the 'extracted text’.

<extracted_text>

Is the text extracted from the LDAP attribute value and incorporated (after modification, if necessary) into the resulting CA DataMinder attribute value.

The extracted text is the text immediately after the search text in an LDAP attribute value, for example, srimmel@unipraxis.com.

You **must** enclose this text in curly brackets { }.

You can use variables to represent any single component that you want to extract. The supported <extracted_text> variables and wildcards are:

%word%

Matches a whole word (that is, an unbroken sequence of alphanumeric characters).

%digits%

Matches a sequence of digits.

%alpha%

Matches a sequence of letters.

?

Matches any single character.

%a% or %A%

Matches any single letter.

%d% or %D%

Matches any single digit.

%l% or %L%

Matches any single letter or digit, but not others such as / or = characters.

%until(x)%

Matches all extracted text up to (but not including) any character x specified in the brackets. For example:

%until(/)% matches all text until a / is detected.

%until(pq)% matches all extracted text until p or q is detected.

Note: If the character is not detected, the parameter matches all text up to the end of the LDAP attribute value.

(x)%untilEnd%

Matches all extracted text starting from (and including) any character x specified in the brackets, up to the end of the input string. For example:

"smtp:{%untilEnd%}"

Matches the entire input string if it begins with an ‘smtp:’ prefix, but removes this prefix from the extracted text. The expression only extracts text within the curly brackets.

%untilstr(str)%

Matches all extracted text up to (but not including) any string specified in the brackets. For example:

```
%untilstr(/ou,/cn)%
```

Matches all extracted text until /ou or /cn is detected.

Note: If the string is not detected, the parameter matches all text up to the end of the LDAP attribute value.

prepend=<term>

Prepends, or precedes, the specified <term> in front of the output string.

Note: You do not need to enclose the prepended string in quotes. For example:

```
prepend=UK.
```

append=<term>

Appends the specified <term> to the output string.

Note: You do not need to enclose the appended string in quotes.

repeat=<N>

(Optional argument) Takes the Nth matching result. For example:

`repeat=0`

Matches the first matching result.

`repeat=1`

Matches the second matching result.

`repeat=last`

Matches the last matching result.

For example, to derive a CA DataMinder user name from:

```
/O=UNIPRAXIS/OU=ADMIN GROUP/CN=RECIPIENTS/CN=SRIMMEL
```

Use this expression:

```
["/O={%word%}"]\\["/CN={%word%}",repeat=last]
```

This yields the CA DataMinder user 'unipraxis\srimmel'.

subsection=<pos_A>:<pos_B>

(Optional argument) Specifies a subsection of the matching text, from character position <pos_A> up to, but not including, character position <pos_B>.

- If <pos_A> and <pos_B> are positive numbers, the subsection is taken from the start of the matching word. Here, zero refers to the first character, 1 refers to the second character; and so on.
- If <pos_A> and <pos_B> are negative numbers, the subsection is taken from the end of the matching word. Here, zero refers to the last character, -1 refers to the penultimate character; and so on.
- If <pos_A> is blank or zero, the subsection is taken from the start of the matching word.
- If <pos_B> is blank or zero, the subsection is taken to the end of the matching word.

For example:

`subsection=0:1`

Specifies a subsection that comprises only the first character of the matching word.

`subsection=1:3`

Specifies a subsection that comprises the second, third, and fourth characters of the word.

`subsection=0:-2`

Specifies a subsection that comprises the final three characters of the word.

substitute=<term>:<new_term>

(Optional argument) Substitutes text in the 'extracted text' with your own string. <term> is the text that you want to replace. <new_term> is your own text that you want to substitute into the CA DataMinder attribute value.

Note: <term> can only match whole words. For example, if the extracted text is UNIPRAXIS, you cannot substitute "UNI" for "MULTI".

mandatory

Specifies that if no matching 'search text' is found, then Account Import quits trying to modify the imported attribute value. Unless substitute text is specified, a null string is imported into the CA DataMinder attribute. That is, the attribute for that user has no value.

Note: 'mandatory' overrides the default declaration.

default=<term>

Specifies that if no matching 'search text' is found, then a default text string is returned as the 'extracted text'. If this happens, <term> is the default text string For example: unknown user.

Note: You do not need to enclose the default string in quotes.

if"Default"[<then>]else[<else>]

Specifies an "if...then...else" clause. If the Default string is found, the <then> expression is performed. If the Default string is not found, the <else> expression is performed. For example:

```
if"Sales.Unipraxis"[prepend=UK.] else["{@%untilEnd%}"]
```

Changes ‘sales.unipraxis’ to ‘UK.Sales.Unipraxis’, or if ‘Sales.unipraxis’ is not detected, then the parameter matches all text from ‘@’ to the end of the LDAP attribute value. (The prepend variable is described above.)

pp"Default":"<replace>"

Performs a search and replace operation on the strings within the quotes. That is, the first string is replaced with the second. For example:

pp",":"" removes all commas

pp",o":"/o" replaces ,o with /o

The supported variables and wildcards are the same as those supported by the <extracted_text> variable.

You can also specify a section to extract and then use that extracted text in the replace expression. For example:

```
pp",{%a%}=":"/%1%="
```

Converts strings such as ,o= to /o=. This example works with any letter. In each case, the single letter found by the search operation is re-used in the replace expression, where %1% is the extracted string from the search text.

Likewise, this example prepends (or precedes) the search string with a forward slash:

```
pp"{?%untilEnd%}":"/%1%"
```

Chapter 12: Object Storage

This section describes the principal methods used for storing event data in CA DataMinder. It describes how to integrate CA DataMinder with third party object storage solutions, that is, EMC Centera, IBM DB2 Content Manager and NetApp SnapLock. It also provides an overview of the CMS temporary object store.

This section contains the following topics:

CMS Overview

The CMS acts as the central collector of ‘event data’ (captured or imported messages). Each CA DataMinder event comprises metadata, written to the local CA DataMinder database, and a blob (Binary Large Object) file, saved on physical media. The CMS is responsible for building the event database, event object store and temporary object store. These components are summarized below:

Object store

The object store contains the blob files. A blob file contains the email content and any attachments, stored in CA DataMinder format. The blob file is written to disk and saved in the \Data folder or migrated to a third party object store.

Note: If required, the object store can be a remote data folder, for example UNC-specified network file share or a network-attached storage (NAS) device.

Temporary object store

The temporary object store is a subfolder of the object store and it contains cached blob files which are also stored on a third party storage device.

Database

Full details about tables and indexes are in the CA DataMinder Database Schema, available from CA Support at http://ca.com/support.

Concurrent Use of Multiple Object Stores

There is considerable flexibility built into the CA DataMinder support for third party storage solutions. In particular, you can separately configure which object store CA DataMinder uses to store and retrieve data.

For example, you can retain legacy object data in a Centera while storing all new data in a file system or in IBM DB2 Content Manager. In both cases, CA DataMinder enables you to continue retrieving both the legacy data and 'new' data when running event searches.

For this example, you leave the Access Node Addresses policy setting unchanged, so CA DataMinder can continue to retrieve legacy data from the Centera. But you set the Data File Storage Location policy setting to 'File System' or 'IBM DB2 Content Manager', so new data is stored in this location.

More information:

Mandatory Centera Policy Settings (see page 316)
Content Manager Integration Policy Settings (see page 325)

Integrating with EMC Centera

CA DataMinder can integrate with EMC Corporation's Centera content addressed storage (CAS) solution to ensure long-term content integrity and online access for large volumes of fixed data. Integration between these two systems provides your enterprise with an end-to-end solution to your email and Web risk management and storage needs. The diagram below shows how a Centera integrates with CA DataMinder to work as an alternative object store.

Only events captured after integration has been set up will be migrated to EMC Centera. Events captured and replicated to the CMS before Centera integration has been set up will not be migrated to Centera. For this reason, we recommend that you set up Centera integration as soon as possible after deploying CA DataMinder.

For information on migrating existing events to a Centera, see How the Centera Integration Works (see page 311). There is no inbuilt mechanism for migrating to Centera those blobs that were stored on the CMS before Centera integration was implemented. It is possible to do this, however. You will need to write a custom SQL procedure to add a blob queue record for each blob migrated to Centera..

Centera integration is configured through policy settings on the CMS. When a CA DataMinder event is captured or imported, it is replicated to the CMS. The CMS stores an event record in its database and writes a blob file to its \data subfolder. This blob file is subsequently moved to the Centera and the CMS database entry for that event is updated to reflect the new location for the blob.

Note: CA DataMinder uses version 3.1.544 of the Centera API Library to integrate with CentraStar. Integration is recommended with CentraStar 3.1. Integration may succeed with other versions but they may not have been tested. To find whether this version of the Centera API Library permits integration with other versions of CentraStar, please contact EMC Corporation.

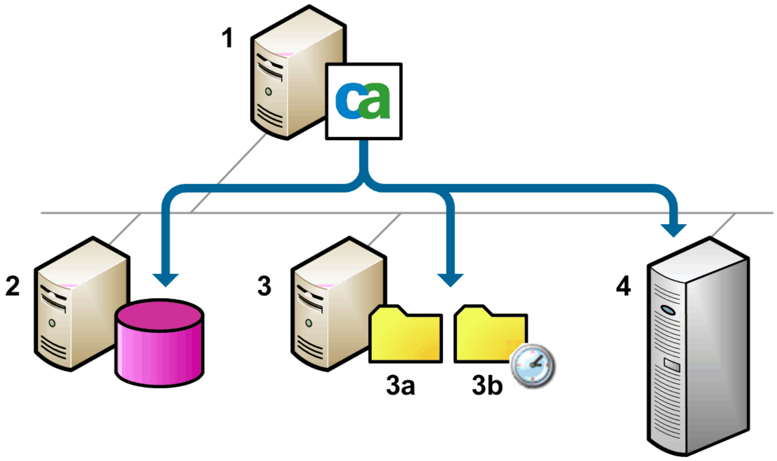

Centera architecture: Data flow

1. **CMS:** Event metadata is written to the database (**2**); event content is saved to the object store or Centera device.
2. **Database server:** The CMS can support a local or remote database. In this example, the database containing metadata runs on a remote server.
3. **Object store:** The object store can be a remote data folder, in this example a UNC-specified network file share, or a network-attached storage (NAS) device. In both cases, it comprises permanent (**3a**) and temporary (**3b**) object stores.
4. **Centera:** As an alternative to a conventional object store, CA DataMinder can integrate with EMC Centera. The blob files for captured or imported events are streamlined from the CMS to the Centera device and stored in clips. Communications with the Centera device use EMC's proprietary TCP/IP-based API.

More information:

How CA DataMinder Events Are Stored (see page 329)

About Centera Integration

This section summarizes the methods used to integrate CA DataMinder with Centera, enabling CA DataMinder events to be migrated to Centera. First, it is necessary to understand how CA DataMinder events are stored on the CMS prior to migration.

How the Centera Integration Works

Centera integration with CA DataMinder is controlled by the CMS. If Centera integration is enabled, when an event is replicated to the CMS from a client machine or gateway, the blob file is first saved to the object store \Data folder. It is then copied to Centera and the source blob file deleted from the CMS.

Note: The blob file is not deleted immediately; it is cached in case it is needed again soon.

The database entry on the CMS is subsequently updated to reflect the new location for the blob, replacing the blob file system location with the Centera content address code.

- **If the Centera is temporarily unavailable:** CA DataMinder maintains a database table containing a list of blobs that need to be stored to Centera. This list can build up if the Centera is offline or if the configuration settings are incorrectly optimized for the installation. We recommend that you monitor the size of this database table as a health check.
- **Migrating existing blob files to a Centera:** There is no inbuilt mechanism for migrating to Centera those blobs that were stored on the CMS before Centera integration was implemented. It is possible to do this, however. You will need to write a custom SQL procedure to add a blob queue record for each blob migrated to Centera.

Note: Any custom SQL procedure to migrate existing blobs must be carefully designed to ensure that the migration queue is properly managed. We strongly recommend that you contact Technical Support for guidance if you need to write a migration procedure; see CA Technical Support (see page 21).

More information:

Managing the Temporary Object Store (see page 331)

Blob Files Are Grouped into Centera Clips

CA DataMinder does not store each blob as a single Centera object. It groups them into clips, significantly reducing write times for event data migrated to Centera. Policy settings on the CMS determine the maximum number of blobs and KB per clip.

By default, blob files are compressed and encrypted when they are stored on the CMS. They are stored to Centera in the same format.

If multiple blobs are stored in Centera clips, CA DataMinder blob files are concatenated into a single Centera blob, referenced from the Centera clip. Blob offset and length details are stored in the clip to optimize data retrieval during event searches.

Note: CA DataMinder does not store blob files with differing trigger-defined Minimum Retention Periods in the same Centera clip.

More information:

Optional Centera Policy Settings (see page 317)
How the Minimum Retention Periods Work (see page 314)

Increasing the Clip Size

The default values for the Maximum Number of blobs per Centera Clip and Maximum Number of KBytes per Centera Clip settings are 1 blob per clip and 1,048,576 KB (1 GB) per Centera clip, respectively. As a general rule, you do not need to change these default values, providing the Centera can keep pace with the CMS when storing data (to check this, you need to monitor the size of the blob queue table in the CMS database). If this is so, there is little advantage in increasing the clip size to achieve faster Centera storage rates.

However, a very fast CMS (for example, a multi CPU, 2 GHz machine) will get better Centera performance by putting more than 10MB in a single Centera clip, so you may want to experiment adjusting the clip size in this situation. Tuning these clip settings for optimum performance also depends on network bandwidth. There is less drop-off in the storage rate for larger Centera clip sizes if you have a Gigabit Ethernet NIC in the CMS. Again, you may need to experiment to find the optimum clip size.

In some circumstances the utilization of all the storage capacity on a Centera Node is limited by the maximum number of objects that can be stored. Generation 3 Centera hardware can store 50 M objects per node. This count includes all types of Centera object, for example CDF, blobs, mirrored copies, and parity fragments. If content is mirrored, then four Centera objects are stored per clip, that is, two CDFs (Clip Definition Files) and two blob files. As a general rule, you should ensure that at least 100 KB is stored in each clip when using mirrored content protection and at least 500 KB when using parity content protection.

- Centera blobs contain multiple CA DataMinder blobs according to the blobs per clip setting. For example, if you store 100 CA DataMinder blobs per clip, 100 events will occupy just four Centera objects.
- If the Centera blob size is less than the embedded data threshold, the Centera blob is stored in the CDF, which means that only two Centera objects are used.

Centera - Blob Retrieval

Blob offset and length details are stored in the Centera clip so that, when retrieving data during an event search, CA DataMinder only retrieves the relevant portion of the Centera blob needed to rebuild the CA DataMinder blob. During an event search, CA DataMinder retrieves the event data from the Centera blob to a regular CA DataMinder blob that can be accessed by the Data Management console or iConsole.

Retrieved blob files are retained in a temporary cache on the CMS. This is the same cache used to retain events retrieved from other third party remote storage locations, such as an email archive. The cache timeout is controlled by the Remote Data Cache Timeout setting in the CMS machine policy; when the timeout expires, the cache is deleted.

Note: When the timeout expires, data is not deleted immediately but is instead moved to a secondary folder within the cache and deleted after another elapsed timeout. This means data may persist in the cache for twice as long as specified.

How the Minimum Retention Periods Work

The Minimum Retention Period in the CMS machine policy determines how long events are retained on the local CA DataMinder machine (though the retention period can be manually overwritten by reviewers in the Data Management console or iConsole, or by a trigger-specific Minimum Retention Period in the user's policy.

CA DataMinder ignores the Centera default retention period. For example, if the CMS's Minimum Retention Period is set to 365 days, and its Event Purge Frequency set to 1 day, then an event purge runs daily on the CMS, deleting events more than one year old from both the CA DataMinder database and the Centera.

Note: The Minimum Retention Period on an Event Import machine (or any CA DataMinder server) does not determine how long blobs are retained on the Centera. Its purpose is solely to determine when events become eligible for purging from the local CA DataMinder machine.

Single-instance Storage

There is little advantage in enabling single instance storage. CA DataMinder blob files contain capture timestamps, making single-instance storage unlikely. Also, storing multiple blobs in a Centera clip or using the Embedded Data Threshold setting will effectively disable single instance storage.

More information:

Optional Centera Policy Settings (see page 317)

Install the Centera Connector

CA DataMinder integration with EMC Centera is provided through the Centera Connector. You install this connector from the CA DataMinder server installation wizard.

To install the Centera Connector

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.
3. In the Advanced Install Options screen, choose CA DataMinder Platform and then click Install.

   This launches the CA DataMinder server installation wizard in a separate window.

4. In the server installation wizard, navigate to the Customer Setup screen.
5. In the Custom Setup screen, expand the CMS Storage Connector feature and select EMC Centera Connector.
6. In the Server Type screen, choose CMS.
7. In the subsequent screens, specify the location of the \Data folder, plus details about the local CA DataMinder database.
8. In the Service Accounts screen, specify the logon accounts used by the local CA DataMinder infrastructure service and other services.
9. In the final wizard screen, click Install to start the file transfer.
10. Now you need to configure Centera integration. Specifically, you need to edit the CMS machine policy; see the following section.

    After editing these policy settings, integration is complete.

More information:

Server Installation Features (see page 38)
Configure Centera Integration (see page 316)
Optional Centera Policy Settings (see page 317)

Configure Centera Integration

To ensure that captured data files are migrated to EMC Centera, you need to configure the integration of your Centera device. You do this by specifying settings in the EMC Centera Integration folder of the CMS machine policy.

As a minimum, you need to set the Data File Storage Location and Access Node Addresses settings (see below), leaving the other settings to use their default values. If you do want to customize the Centera integration by editing the other settings, we recommend that you consult your Centera documentation for guidance. The optimum configuration for your organization will depend on:

- The CMS hardware (memory and processors).
- The Centera configuration (the number and addresses of access nodes, and whether the Centera needs access credentials or multi-cluster failover options).
- The average size of captured or imported events.

The relevant policy settings are described in the following sections.

Mandatory Centera Policy Settings

As a minimum, you need to set these policy settings:

Data File Storage Location

> Find this setting in the \Data Management subfolder of the machine policy. Set this to ‘EMC Centera’.
>
> This setting controls where captured data files (blobs or Binary Large Objects) are stored. By default, the blob file is written to disk and saved in the \Data folder specified when the CMS was installed.

Access Node Addresses

> Find this setting in the \Data Management\EMC Centera Integration subfolder of the machine policy.
>
> Enter the addresses (separated by commas) of the Centera Access nodes here. If this setting is blank, captured data blob files (binary large objects) will not be stored to Centera. This setting may also be used to specify additional Centera connection options—see your Centera documentation for details.

Note: These settings permit the concurrent use of multiple object stores.

More information:

Concurrent Use of Multiple Object Stores (see page 308)

Optional Centera Policy Settings

These settings are optional. Find them in the \Data Management\EMC Centera Integration subfolder of the machine policy.

Maximum Number of Concurrent Storage Operations

Defaults to 10. This setting controls the number of Centera clips that can be stored to Centera simultaneously. Storing more clips at the same time may result in a faster storage rate. You may want to set this number to match the number of storage nodes in the Centera cluster (if there are more than 10 nodes).

Maximum Number of Blobs per Centera Clip

Defaults to 1. This setting defines how many blob files are stored within one Centera clip.

The default value is set to 1 as a precautionary measure. It ensures that data belonging to an event cannot be deleted early if a custom purge script is used to purge events based on criteria other than event expiry date. If you implement such a purge script and need to set this value to more than 1, contact Technical Support for further information. If using the standard purge script, we recommend that this value is set to 100.

Note: Blobs with different retention periods will not be stored in the same Centera clip.

Maximum Number of KBytes per Centera Clip

Defaults to 1,048,576 KB (1 GB). This setting limits the number of blob files stored in a Centera clip by ensuring that the total number of KB stored in the clip does not exceed the value set here.

Clip Buffer Size (KB)

Defaults to 64. This is the size of the Centera internal clip buffer in KB. It specifies the amount of memory to use for temporary storage when sending data to the Centera device.

Content Address Calculation

Defaults to ‘Calculate Content Address During Write’. This setting controls the method used to calculate the Content Address from blob files. Possible values are:

Calculate Content Address Before Write

Content Addresses are calculated on the CMS before data is streamed to the Centera device.

Calculate Content Address During Write

Content Addresses are calculated on the CMS while streaming data to the Centera device.

Calculate Content Address After Write

Content Addresses are calculated on the Centera device itself.

Flush Interval (Seconds)

Defaults to 60. This specifies how long CA DataMinder waits before storing a clip if there are insufficient blob files to fill it. Clip capacity is defined the Maximum Number of blobs per Centera Clip and Maximum Number of KB per Centera Clip settings. If there are not enough blobs to fill the clip according to these settings, CA DataMinder waits for the flush interval to elapse then stores the clip anyway.

Embedded Data Threshold (KB)

Defaults to zero. This is the maximum data size, in KBytes, for data to be embedded in the Centera clip XML instead of being stored separately. The default value of zero KB means that data is never embedded in the clip XML. The maximum value is 100 KB. Embedding data in the clip can improve write performance, but embedded data does not benefit from single-instance storage.

Integrating with IBM DB2 Content Manager

CA DataMinder can integrate with IBM DB2 Content Manager to allow storage management, archiving and retrieval of large volumes of captured data. The diagram below shows how IBM DB2 Content Manager integrates with CA DataMinder to work as an alternative object store.

Only events captured after integration has been set up will be migrated to the Content Manager. Events captured and replicated to the CMS before the Content Manager integration has been set up will not be migrated. For this reason, we recommend that you set up Content Manager integration as soon as possible after deploying CA DataMinder.

IBM DB2 Content Manager integration is configured through policy settings on the CMS. When a CA DataMinder event is captured or imported, it is replicated to the CMS. The CMS stores an event record in its database and writes a blob file to its \data subfolder. This blob file is subsequently moved to the IBM Content Manager and the CMS database entry for that event is updated to reflect the new location for the blob.

Note: CA DataMinder is compatible with Version 8 Revision 3 of IBM DB2 Information Integrator for Content. The integration has been tested with Version 8 Revision 3 of IBM DB2 Content Manager Enterprise Edition, but it may be compatible with other versions. Please contact IBM Corporation for more information.

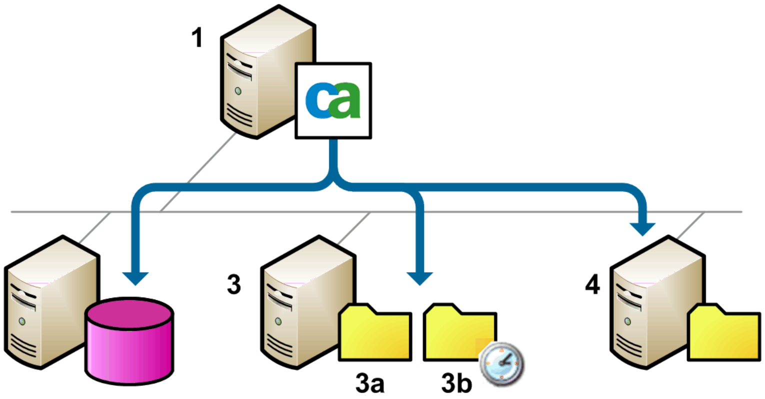

IBM Content Manager architecture: data flow

1. **CMS:** Event metadata is written to the database (**2**); event content is saved to the object store or Content Manager.
2. **Database server:** The CMS can support a local or remote database. In this example, the database containing metadata runs on a remote server.
3. **Object store:** The object store can be a remote data folder, in this example a UNC-specified network file share, or a network-attached storage (NAS) device. In both cases, it comprises permanent (**3a**) and temporary (**3b**) object stores.
4. **Content Manager:** As an alternative to a conventional object store, CA DataMinder can integrate with IBM DB2 Content Manager. The Content Manager can be hosted on more than one machine. The blob files for captured or imported events are streamed from the CMS to the Content Manager and stored in resource items. Communications with the IBM Content Manager use a proprietary TCP/IP-based scheme.

More information:

How CA DataMinder Events Are Stored (see page 329)

About IBM Content Manager Integration

This section summarizes how CA DataMinder integrates with IBM DB2 Content Manager, enabling CA DataMinder events to be migrated to the Content Manager. First, it is necessary to understand how CA DataMinder events are stored on the CMS prior to migration.

How the IBM Content Manager Integration Works

Content Manager integration with CA DataMinder is controlled by the CMS. If the integration is enabled, when an event is replicated to the CMS from a client machine or gateway, the blob file is first saved to the object store \Data folder. It is then copied to the Content Manager and the source blob file deleted from the CMS.

Note: The blob file is not deleted immediately; it is cached in case it is needed again soon.

The database entry on the CMS is subsequently updated to reflect the new location for the blob, replacing the blob file system location with the Content Manager resource identifier.

- **If the Content Manager is temporarily unavailable:** CA DataMinder maintains a database table containing a list of blobs that need to be migrated to the Content Manager. This list can build up if the Content Manager is offline or if the configuration settings are incorrectly optimized for the installation. We recommend that you monitor the size of this database table as a health check.
- **Migrating existing blob files to a Content Manager:** There is no built-in mechanism for migrating to Content Manager those blobs that were stored on the CMS before Content Manager integration was implemented. It is possible to do this, however. You will need to write a custom SQL procedure to add a blob queue record for each blob to be migrated.

Note: Any custom SQL procedure to migrate existing blobs must be carefully designed to ensure that the migration queue is properly managed. We strongly recommend that you contact Technical Support for guidance if you need to write a migration procedure; see CA Technical Support (see page 21).

More information:

Managing the Temporary Object Store (see page 331)

Blob Files Are Grouped into Content Manager Resource Items

CA DataMinder does not store each blob as a single Content Manager resource item. It groups them, significantly reducing write times for event data migrated to Content Manager. Policy settings on the CMS determine the maximum number of blobs and bytes per resource item.

If multiple blobs are stored in Content Manager resource items, CA DataMinder blob files are concatenated into a single item and blob offset and length details are stored with the Content Manager resource ID in the CA DataMinder database to optimize data retrieval during event searches.

By default, blob files are compressed and encrypted when they are stored on the CMS. They are stored to Content Manager in the same format.

Note: CA DataMinder does not store blob files with differing trigger-defined Minimum Retention Periods in the same Content Manager resource item.

More information:

Increasing the Resource Item Size (see page 322)
How the Minimum Retention Periods Work (see page 323)

Increasing the Resource Item Size

The default values for the Maximum Number of blobs per Resource Item and Maximum Number of Bytes per Resource Item settings are 1 and 1 GB respectively. As a general rule, you do not need to change these default values, providing the Content Manager can keep pace with the CMS when storing data (to check this, you need to monitor the size of the blob queue table in the CMS database). If this is so, there is little advantage in increasing the settings to achieve faster storage rates.

Blob Retrieval

Blob offset and length details are stored in the Centera clip so that, when retrieving data during an event search, CA DataMinder only retrieves the relevant portion of the Centera blob needed to rebuild the CA DataMinder blob. During an event search, CA DataMinder retrieves the event data from the Centera blob to a regular CA DataMinder blob that can be accessed by the Data Management console or iConsole.

Retrieved blob files are retained in a temporary cache on the CMS. This is the same cache used to retain events retrieved from other third party remote storage locations, such as an email archive. The cache timeout is controlled by the Remote Data Cache Timeout setting in the CMS machine policy; when the timeout expires, the cache is deleted.

Note: When the timeout expires, data is not deleted immediately but is instead moved to a secondary folder within the cache and deleted after another elapsed timeout. This means data may persist in the cache for twice as long as specified.

How the Minimum Retention Periods Work

The Minimum Retention Period in the CMS machine policy determines how long events are retained on the local CA DataMinder machine (though the retention period can be manually overwritten by reviewers in the Data Management console or iConsole, or by a trigger-specific Minimum Retention Period in the user’s policy.

For example, if the CMS’s Minimum Retention Period is set to 365 days, and its Event Purge Frequency set to 1 day, then an event purge runs daily on the CMS, deleting events more than one year old from both the CA DataMinder database and the Content Manager.

Note: The Minimum Retention Period on an Event Import machine (or any CA DataMinder server) does not determine how long blobs are retained on the Content Manager. Its purpose is solely to determine when events become eligible for purging from the local CA DataMinder machine.

Set Up IBM DB2 Content Manager Integration

CA DataMinder integration with IBM DB2 Content Manager is provided through the Content Manager Connector. You install this connector from the CA DataMinder server installation wizard.

To install the Content Manager Connector

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose CA DataMinder Platform and then click Install.

   This launches the CA DataMinder server installation wizard in a separate window.

4. In the server installation wizard, navigate to the Customer Setup screen.
5. In the Custom Setup screen, expand the CMS Storage Connector feature and select IBM Content Manager Connector.
6. In the Server Type screen, choose CMS.
7. In the subsequent screens, specify the location of the \Data folder, plus details about the local CA DataMinder database.
8. In the Service Accounts screen, specify the logon accounts used by the local CA DataMinder infrastructure service and other services.
9. In the final wizard screen, click Install to start the file transfer.

More information:

Server Installation Features (see page 38)

Configure Content Manager Integration

To ensure that captured data files are migrated to IBM Content Manager, you need to configure the integration. You do this by specifying settings in the CMS machine policy.

Content Manager Integration Policy Settings

Unless stated otherwise, find these settings in the \Data Management\IBM Content Manager Integration subfolder of the machine policy.

Data File Storage Location

This setting is mandatory. Find it in the \Data Management\ subfolder of the machine policy. Set it to 'IBM DB2 Content Manager'.

This setting controls where captured data files (Binary Large Objects) are stored. By default, the blob file is written to disk and saved in the \Data folder specified when the CMS was installed.

Note: You can use this setting to configure the concurrent use of multiple object stores.

Content Manager Database Name

This mandatory setting must be set to the name of the IBM Content Manager database, for example, icmnlsdb.

Content Manager User Name

This mandatory setting must be set to the user name of the IBM Content Manager user permitted to read and write Resource Items.

Content Manager Password

This mandatory setting must be set to the password of the Content Manager user permitted to read and write Resource Items.

Content Manager Connection Options

This setting controls additional Content Manager connection options and should not normally be necessary. Refer to your IBM DB2 Content Manager documentation for more information.

Content Manager Resource Database Name

Defaults to rmdb. This sets the name of the resource manager used to store the CA DataMinder resource items. The default resource manager is automatically set up when Content Manager 8.3 is installed. Resources are stored as items of type S_wgnblob, which is set up by the CA DataMinder Content Manager Integration when it starts for the first time.

Content Manager Interface Service Host Address

This is the TCP connection information used to locate the service provided by WgnIBMCM.exe (that is, the CA DataMinder IBM Content Manager Interface Service). The setting may be left at the default of 127.0.0.1 on port number 56200 if the WgnIBMCM.exe service is installed on the CMS. If the service needs to be installed on another machine, set that machine's name or IP address here.

Note: The IBM DB2 Information Integrator for Content component must be installed on the machine which hosts WgnIBMCM.exe.

Content Manager Interface Service Port Number

This setting along with the Content Manager Interface Service Host Address is the TCP connection information used to locate the CA DataMinder Content Manager Interface Service.

Maximum Number of Concurrent Storage Operations

Defaults to 10. This setting controls the number of resource items that can be stored to the Content Manager simultaneously. Storing more items at the same time may result in a faster storage rate.

Maximum Number of BLOBs per Resource Item

Defaults to 1. This setting defines how many CA DataMinder blob files are stored within one IBM Content Manager blob file.

The default value is set to 1 as a precautionary measure. It ensures that data belonging to an event cannot be deleted early if a custom purge script is used to purge events based on criteria other than event expiry date. If you implement such a purge script and need to set this value to more than 1, contact Technical Support for further information. If using the standard purge script, we recommend that this value is set to 100.

Note: CA DataMinder blobs with different retention periods will not be stored in the same IBM Content Manager blob file.

Maximum Number of KBytes per Resource Item

Defaults to 1,048,576 KB (1 GB). This setting limits the number of CA DataMinder blob files stored in an IBM Content Manager blob by ensuring that the total number of bytes stored in the blob does not exceed the value set here.

Flush Interval (Seconds)

Defaults to 60. This specifies how long CA DataMinder waits before storing a resource item if there are insufficient blob files to fill it. Resource item capacity is defined by the Maximum Number of BLOBs per Resource Item and Maximum Number of KB per Resource Item settings. If there are not enough blobs to fill the resource item according to these settings, CA DataMinder waits for the flush interval to elapse then stores the resource item anyway.

More information:

Concurrent Use of Multiple Object Stores (see page 308)

IBM Content Manager Logging

To configure the level of logging used by the IBM Content Manager, you need to modify a value in the following registry key:

```
HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA DataMinder
  \CurrentVersion\IBM CM Interface
```

Within this registry key, edit the following value:

LogLevel

Type: REG_DWORD

Data: Defaults to 2. This setting determines the level of logging for storage management. For example, you can configure the IBM Content Manager to only log errors or warning system messages.

Log entries are written to the WgnIBMCM_<date>.log file, where <date> is the date and time when the log file was created; the file is located in CA's \data\log subfolder of the Windows All Users profile. The supported logging levels are:

- Errors only
- Errors and warnings
- Errors and warnings, plus informational and status messages

Note: Setting LogLevel=3 will cause the log file to grow extremely rapidly. This level of logging is provided for testing and diagnostic purposes, for example, it shows storage and retrieval on every resource item.

Integrating with NetApp SnapLock

CA DataMinder supports NetApp SnapLock volumes when NetApp is running Data ONTAP version 7.0 or later. Data ONTAP is a unified storage software platform that dramatically simplifies data management, while SnapLock volumes provide high performance and high security data permanence. Integration with CA DataMinder is simple to set up, and enabled through a single machine policy setting.

Only events captured after integration has been enabled will have their retention periods set on the NetApp SnapLock volume. Events captured and replicated to the CMS before SnapLock integration has been set up will not. For this reason, we recommend that you set up SnapLock integration as soon as possible after deploying CA DataMinder.

SnapLock integration is configured through a single policy setting on the CMS. This setting determines the data file storage location. When a CA DataMinder event is captured or imported, it is replicated to the CMS. The CMS stores an event record in its database and writes a blob file to the specified SnapLock volume. In addition, the file’s ‘Last Accessed’ date is set to the event’s expiry date (that is, when the event’s minimum retention period expires and the event becomes eligible for purging).

Set Up SnapLock Integration

To integrate with a NetApp SnapLock volume

1. You must specify the location of the SnapLock volume when you install the CMS.

   But unlike a conventional CMS installation, in the Data Location wizard screen, you specify a UNC path or a mapped drive to the target SnapLock volume.

2. After installing the CMS, CA DataMinder automatically writes blob (Binary Large Object) files to the SnapLock storage location that you specified in step 1**.**

   These blobs contain the email content and any attachments, or the Web page plus any uploaded files, stored in CA DataMinder format.

3. You now need to configure the CMS to write the minimum retention period (also known as the ‘expiry date’) for each event to the SnapLock volume. To do this, you edit the CMS machine policy.

More information:

Mandatory SnapLock Policy Setting (see page 328)

Mandatory SnapLock Policy Setting

Locate the following setting in the CMS machine policy:

Data File Storage Location

Find this setting in the \Data Management policy folder and set it to ‘NetApp SnapLock’.

This setting controls where captured data files (blobs or Binary Large Objects) are stored. For SnapLock integration, blob files are written to the target SnapLock volume specified in step 1 above.

Temporary Object Store

The temporary object store is an event cache, used when CA DataMinder is integrated with third party email archive solutions, for the following tasks:

Imported events

When importing archived emails into the CA DataMinder system, the event metadata is written to the CMS database as normal, and the associated blob is cached in the temporary object store. After the CMS receives confirmation that the event has been successfully imported, the associated blob is moved to the temporary object store. This allows fast event retrieval during event searches during the period immediately following the import operation.

Retrieving events

When running a search to retrieve events archived in third party remote storage locations and display them in the Data Management console, or iConsole, the associated blob files remain in the temporary object store for a configurable length of time.

More information:

Policy Engines Overview (see page 336)

How CA DataMinder Events Are Stored

Each CA DataMinder event comprises metadata, written to the local CA DataMinder database, and a blob file (binary large object), saved on physical media:

- A database entry contains the event metadata. For example, database fields specify an email’s delivery date, 'envelope' details, what policy triggers were applied, and so on.
- A blob file contains the email content and any attachments, stored in CA DataMinder format. The blob file is written to disk and saved in the \Data folder or migrated to EMC Centera, or the IBM DB2 Content Manager.

More information:

Optional Data Location Structure (see page 331)

Configure the Temporary Object Store

Using settings in the CMS policy, you can configure how long events remain in the temporary object store and how long the RDM waits for a response from the CMS when attempting to retrieve events. After the configurable timeout (see below), blobs in the temporary objects store are deleted. Any subsequent searches for these events must use the Remote Data Manager.

To configure the temporary object store, use the following settings in the Infrastructure\Data Management folder in the CMS machine policy:

Remote Data Cache Timeout (Minutes)

Specify how long (in minutes) data retrieved from the Remote Data Manager (RDM) server is retained in a temporary local cache. After this period, the cache is deleted. The minimum timeout is 1 minute; the maximum timeout is 525600 minutes (365 days).

The default timeout is 4320 minutes. This means the cache is deleted every 3 days. You may need to increase the default timeout for newly deployed CA DataMinder installations if you are using the Content Indexer to index large volumes of captured data.

Note: Be aware of the following:

- The timeout represents a minimum retention period. When the timeout expires, data is not deleted immediately but is instead moved to a secondary folder within the cache and deleted after another elapsed timeout. This means data may persist in the cache for twice as long as specified.
- If you intend to schedule regular purges, we recommend that you set the Remote Data Cache Timeout to synchronize with the event purge frequency to avoid blob files remaining on the CMS afterwards. For details, see the Administration console help; search for ‘purging events’.

Remote Data Request Timeout (Seconds)

Specify how long (in seconds) the CMS will wait for a response from the Remote Data Manager (RDM) server. If this timeout expires before the RDM can retrieve an event, CA DataMinder displays a error message. For example, an error message is shown if the Data Management console is unable to display an archived email because this timeout has expired.

Managing the Temporary Object Store

If the disk or network share holding the cache becomes full, the CMS infrastructure is automatically suspended (machine policy settings determine the critical levels of free disk space).

Note: There is no mechanism to automatically delete items from the cache in this situation.

When the infrastructure restarts, the cache contents and cache timeout schedule is persisted securely in the database.properties file. There is no wgninfra -exec command to clear a cache folder, and no log file entries are written when a cache folder is cleared.

The cache is split into three folders: 0, 1 and 2. When the cache timeout expires, each folder cycles through to the next role in this sequence:

```
Current cache to Previous cache to Deleted cache
```

This sequence is demonstrated in the table below:

	Blob added to cache	**1st timeout expires**	**2nd timeout expires**	**3rd timeout expires**
	Starting role:	Role changes to:	Role changes to:	Role changes to:
Folder 0	Current cache	Previous cache	Deleted cache	Current cache
Folder 1	Previous cache	Deleted cache	Current cache	Previous cache
Folder 2	Deleted cache	Current cache	Previous cache	Deleted cache

If the size of the cache is getting larger than anticipated, we recommend that you reduce the cache timeout period. You can also manually delete the cache folders if necessary. They are recreated automatically as required.

Optional Data Location Structure

The CMS stores event blob files in CA's \data folder of the Windows All Users profile. Within this folder are the following subfolders:

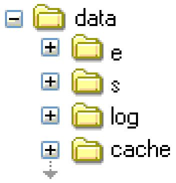

Data

This is the database blob location and contains *.properties files (for example, startup.properties) and *.kdt data blob encryption keys.

e

This stores database event blob files.

s

This stores system blob files. For example, policy files and other system files.

log

This stores log files.

cache

This stores event log files.

You can specify a new location for both the \e and \cache folders, enabling you to separate database event cache files from system blob files. This can be useful if the \data folder is on a remote drive, as logging can potentially fail if the drive becomes inaccessible.

Change the Location of Specific Subfolders in the \data Folder

To change the location of specific subfolders within the \data folder, edit the startup.properties file.

Follow these steps:

1. Stop the 'CA DataMinder infrastructure' service. From a command prompt, run:

   ```
   net stop wgninfra
   ```

2. Edit the startup.properties file. Find this file in the \system subfolder of the CA DataMinder installation folder.
3. Open this file and add the following line(s) to the [Database] section:
   a. To specify a new location for database cache files, add the following parameter:

      ```
      db.CacheLocation=C:\\Dec 07\\cachefiles
      ```

   b. To specify a new location for event blob files, add the following parameter:

      ```
      db.BlobLocation=C:\\Dec 07\\blobfiles
      ```

   Note: You *must* prefix any backslash character in the path with an extra backslash to ensure it is interpreted correctly.

4. While the Infrastructure is still stopped, move the \e and \cache folders to their new locations.
5. Restart the CA DataMinder infrastructure service. From a command prompt, run:

   ```
   net start wgninfra
   ```

Chapter 13: Policy Engines

This section introduces CA DataMinder policy engines. It explains how and why policy engines are used. In particular, it describes how policy engines play a critical role when integrating CA DataMinder with third party e-mail and archiving solutions. This section also provides full instructions for installing, configuring and monitoring policy engines.

This section contains the following topics:

Policy Engines Overview

Policy engines can process events such as emails or files arriving from an external source and apply policy triggers to these events. Typically, events are allocated to individual policy engines by a policy engine hub. Policy engines enable CA DataMinder to integrate directly with:

Email servers

Policy engines can apply policy to emails intercepted by email server agents allow CA DataMinder to monitor and control email activity that would otherwise be missed by Outlook and Notes endpoint integration alone.

CA DataMinder supports email server agents for Exchange, Domino, SMTP servers, Sendmail, and Postfix.

Archive solutions

Policy engines can apply policy to emails sent from a CA DataMinder archive agent and apply smart tags (typically an email category and a retention date) before the emails are archived.

CA DataMinder can integrate with Autonomy ZANTAZ, Symantec Enterprise Vault, and EMC SourceOne.

Scanned items

Policy engines can apply policy to items scanned by the File Scanning Agent (FSA).

Data crossing the network boundary

Policy engines can apply policy to communications intercepted by CA DataMinder Network (also called the 'NBA'). The NBA operates at the boundary between your organization and the Internet, intercepting emails, files and IM conversations,

Imported files, emails, and IM conversations

Policy engines can apply to imported events as part of an Import Policy job.

For example, Import Policy enables you to apply policy to any files stored on your network. For example, you can categorize or apply smart tags to important business documents or reports.

Policy Engine Architecture

Emails are allocated to individual policy engines by the policy engine hub. The hub and policy engines are designed to handle each email with minimal delay. In particular, the hub distributes processing across multiple policy machines in a manner that achieves optimum load-balancing and maximizes throughput. It can also handle hardware failures on remote policy engine machines seamlessly, redistributing events to other policy engines if necessary.

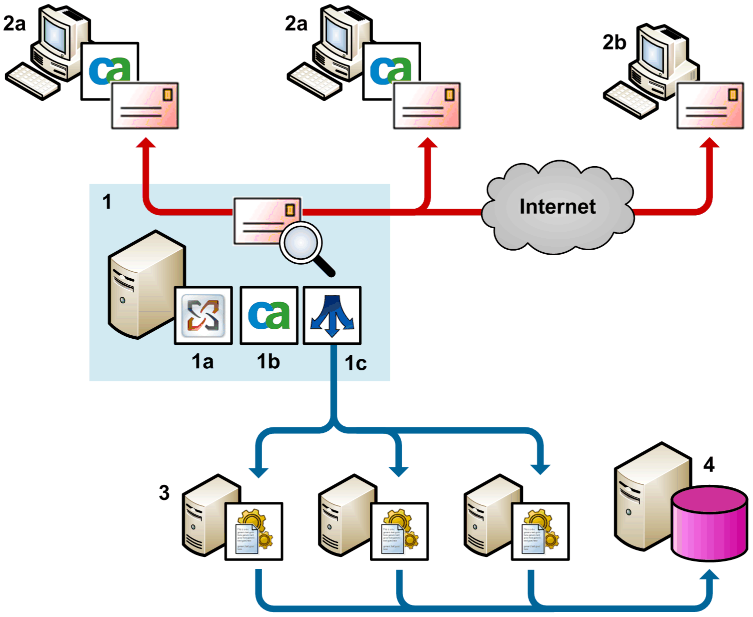

Policy engines and email server integration

This example shows how policy engines can be used to integrate CA DataMinder with an email server.

1. **Email server.** CA DataMinder can integrate with Microsoft Exchange or Lotus Domino (**1a**). This server also hosts the CA DataMinder Exchange server agent or Domino server agent (**1b**) and policy engine hub (**1c**).

   The policy engine hub creates connections between the email server agent and each policy engine host machine and also maintains performance and event processing statistics for each host machine.

2. **Email interception.** Emails transiting through the server, whether sent from internal machines (**2a**) or external machines (**2b**), are detected by the email server agent and passed to the policy engine hub.
3. **Policy engines.** When the policy engine hub receives a new e-mail from the e-mail server agent, it allocates the email to the least heavily loaded policy engine (that is, the policy engine that can process the new email most quickly). The policy engine then analyzes the e-mail and applies policy triggers as necessary.
4. **CMS.** Each policy engine replicates processed emails up to the CMS.

Deploying Policy Engines

Deploying policy engines is the first step in the overall task of integrating CA DataMinder with third party email or archiving solutions, or when setting up the Import Policy feature.

The key deployment tasks are as follows:

1. Specify the required Windows and CA DataMinder user accounts.
2. Install your policy engines.
3. Configure your policy engines.
   - Configure the machine policy on the PE host machine.
   - Edit the registry to enable user lookup operations.
4. Deploy the policy engine hub.
5. Set up the event source. Policy engines process events arriving from an external source.

More information:

Specify PE User Accounts (see page 340)
Configure the Local Machine Policy (see page 343)
Policy Engines Overview (see page 336)
Configure the Policy Engine Registry Values (see page 347)

PE Host Machine Memory Requirement

Important! We strongly recommend that the host machine has enough memory to cache all effective user policies for your organization simultaneously.

When processing an email, a policy engine requires access to the sender’s user policy. The policy engine retrieves the necessary policy from the CMS and, to minimize processing delays, it can retain multiple user policies in memory. Because each policy can take up a significant amount of memory, you must ensure that the host machine has sufficient memory to simultaneously retain all the user policies that is likely to need.

In practice, this means the host machine must be able to simultaneously hold all the ‘effective’ user policies for your organization. For each user, their effective policy is either their own individual policy (if it has been uniquely customized) or the policy inherited from their parent group.

Because most CA DataMinder users (and indeed, most groups) inherit their policy, direct and unchanged, from their parent group, the total number of effective user policies is usually far smaller than the total number of users. For example, if all users and subgroups in your organization inherit, direct and unchanged, the policy for the top level ‘Users’ group, then in effect all users are governed by the same policy. That is, for your entire organization there is only one effective policy.

Nevertheless, to prevent excessive memory allocation, the local machine policy specifies the maximum number of policies that can be held in memory at one time.

More information:

Specify a PE Domain User (see page 340)
Configure the Local Machine Policy (see page 343)

Active and Standby Policy Engines

When you configure the hub, you can designate individual policy engines as either ‘Active’ or ‘Standby’. During normal operations, the hub only allocates emails to active policy engines, but if one or more active policy engines become unavailable it can automatically transfer processing to standby policy engines. This allows you to ensure that CA DataMinder continues monitoring and processing emails even during a contingency. For example, you may want to install active policy engines in your main office and standby policy engines off-site.

To set up ‘Active’ and ‘Standby’ policy engines, you need to configure the policy engine hub. Specifically, you need to edit the ActivePolicyEngines and StandbyPolicyEngines registry values on the hub host machine.

More information:

Policy Engine Hub Registry Values (see page 373)

Email Address Mapping

Before a policy engine can apply policy triggers to an intercepted email, it needs to map the sender's email address to a CA DataMinder user. This mapping identifies the email owner and determines which policy to apply.

Specify PE User Accounts

Before you can deploy your policy engines, you must specify a Windows domain user that allows the policy engines and hub to communicate. You must also create a new CA DataMinder user account.

Specify a PE Domain User

Your policy engines must be able to access the policy engine hub. Specifically, the policy engine service (wgnpesv.exe) must run as a domain user who can access the host machine running the policy engine hub. Likewise, the policy engine hub uses this same domain user to access the remote policy engine machines.

Therefore, you must either create a new domain user in Active Directory, or choose an existing domain user. This is your 'PE domain user'. Throughout This section, the term 'PE domain user' refers to the domain user you are using to pilot your policy engines and the policy engine hub.

The PE domain user must also be a member of the local Administrators group on all policy engines. This requirement is restated where appropriate in the following sections.

Note: When you deploy your policy engines, you will also need to modify the logon properties of this PE domain user on the policy engine host machine.

Create a Corresponding CA DataMinder User

After specifying your PE domain user, you must create a matching CA DataMinder user account. That is, the new CA DataMinder user must have the same account name as the PE domain user. The policy engines will use this CA DataMinder account to log on to the CMS when mapping email addresses onto CA DataMinder users.

1. In the CA DataMinder Administration console, create a new user. See the Administration console online help for details about creating new users; search the index for ‘new accounts’.

   When you specify the user name, you must include the domain prefix to ensure compatibility with the account name for the PE domain user (for example, UNIPRAXIS\PolicyEngineUser).

2. Still in the Administration console, assign the 'Events: Allow bulk session management' administrative privilege to this CA DataMinder user. This permits the new user account to access multiple CA DataMinder user accounts.

   Note: This new user account does not need a management group or any other administrative privileges.

Install Policy Engines

You install policy engines using the CA DataMinder server installation wizard.

To install a policy engine

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose CA DataMinder Platform and then click Install.

   This launches the CA DataMinder server installation wizard in a separate window.

4. In the server installation wizard, navigate to the Customer Setup screen.

5. In the Custom Setup screen, choose Policy Engine and, optionally, the Socket API.

   The Socket API is automatically set up to listen on port number 8538. Using socket connections enables you to call the External Agent API from a remote location, including from a non-Windows system. For example the CA DataMinder Network Boundary Agent uses the Socket API to analyze traffic leaving or entering the corporate network from the internet.

6. In the Server Type screen, choose Gateway.

7. In the subsequent screens, specify the parent server (typically the CMS), the location of the \Data folder, plus details about the local CA DataMinder database.
8. In the Service Accounts screen, specify the logon accounts used by the local CA DataMinder infrastructure service and the policy engine service.

   The policy engine service must run as the **PE domain user**. Click the Browse button for the Policy Engine service, then enter the domain, name and password of the PE domain user .

   Note: The PE domain user must belong to the local Administrators group.
9. In the final wizard screen, click Install to start the file transfer.
10. When the installation is complete, the CA DataMinder Policy Engine service starts automatically.

    Now you must manually configure the local machine policy—see the following section.

More information:

Specify a PE Domain User (see page 340)

Configure Policy Engines

To configure your policy engines, you must:

1. Edit registry values on the policy engine host machine.
2. Edit the local machine policy.
3. (Optional) Set up integration with Voltage SecureMail.

More information:

Configure the Local Machine Policy (see page 343)
Configure the Policy Engine Registry Values (see page 347)
Set Up SecureMail Integration (see page 348)

Configure the Local Machine Policy

Before installing the policy engine and starting the service, you need to configure some policy engine performance parameters and settings to determine how the policy engine applies policy to emails from unrecognized senders and to files when no other means are available to determine the policy participant. To do this, you edit the machine policy for the host machine. You need to modify these settings in the Policy Engine folder of the local machine policy:

Maximum Number of Loaded Policies

Defaults to zero. A zero value means an unlimited number of policies can be retained in memory.

This setting defines the maximum number of user policies that the policy engine can hold in its memory at one time. Because each policy requires a significant amount of memory, this setting can prevent excessive memory usage.

Note that if the policy engine is already holding its maximum number of policies when it needs to load a new policy (in order to process an email from a sender whose policy is not already cached), it discards the least recently used policy before loading the new policy.

However, such policy swaps can significantly slow the processing for an individual email. For this reason, we strongly recommend that your policy engine host machine has sufficient memory so that you do not need to limit number of loaded policies. *In fact, we recommend that the host machine can hold all the effective policies for your organization simultaneously.*

Maximum Number of Concurrent Operations

Defaults to 5. This setting defines the maximum number of emails that can be processed simultaneously by a policy engine. It enables the policy engine to make the most efficient use of system resources. You do not normally need to change the default value. However, you may want to increase the maximum limit if the policy engine is running on a multiprocessor computer.

If this maximum limit is reached, the policy engine delays accepting any further emails from the policy engine hub until the number of emails being processed falls below this maximum limit. This means that when an email completes processing, another is accepted, so maintaining the number of emails at the maximum limit. For example, if five emails finish processing simultaneously, the policy engine immediately accepts five new emails.

Perform LDAP directory lookups?

This setting is provided for diagnostic purposes only. It specifies whether the policy engine can retrieve email address details and distribution list members from an LDAP directory. *We strongly recommend that you do not change this setting!*

Embedded Message Identification

Policy engines need to distinguish between ‘genuine’ emails and ‘embedded content’ emails (that is, EML emails containing embedded IM conversations, Bloomberg messages or other communications such as eFaxes). This is accomplished through the Embedded Message Identification policy setting.

This setting enables policy engines to detect embedded content emails and set the event type as ‘embedded IM’, ‘Bloomberg’, or ‘eFax’. For IM conversations, this setting can also be used to extract or set the IM network.

The default values for this setting enable policy engines to automatically detect:

- Embedded IM conversations in EML emails generated by the Cnv2email.exe utility.
- Embedded Bloomberg messages in EML emails generated by BB2email.exe.

However, if you want policy engines to detect other forms of embedded content (such as eFaxes or IM conversations embedded in EML files that were generated by third party applications), you need to add additional values to this policy setting.

Deadlock Detection Timeout

Defaults to one hour. This setting is designed to maintain processing capacity. It specifies how long a thread must be inactive while processing an event before the policy engine considers the thread to have stalled.

To guard against any problems that might cause a policy engine to take an excessively long time to analyze an event, the policy engine monitors all processing threads. If it detects a deadlock, it creates a new thread for each stalled thread.

Retention Period for Unused Policies

Defaults to 7 (days). This setting defines the frequency of policy time-outs. That is, the amount of time a policy engine retains a policy that has not been used. After this period of time, the policy is unloaded.

Unknown Internal Sender

This setting specifies the name of a CA DataMinder user. It defaults to UnknownInternalSender; this user account is created automatically when you install a new CMS.

Policy engines use this setting to apply policy to emails sent from someone **within** your organization. The policy engine applies the Unknown Internal Sender’s policy if the sender’s address matches an address pattern listed in the Internal Email Address Pattern setting) but no corresponding user exists. For example, this can happen if a new recruit has an account in Active Directory but no CA DataMinder account has been created for them yet.

Important! You can specify a different account if necessary, but this setting must identify a user account, not a group account. Restart the policy engine for the changes to take effect.

External Sender

This setting specifies the name of a CA DataMinder user. It defaults to ExternalSender; this user account is created automatically when you install a new CMS.

Policy engines use this setting to apply policy to external emails. That is, emails sent from someone **outside** your organization. The policy engine applies the External Sender's policy if the sender's address does **not** match an address pattern listed in the Internal Email Address Pattern setting (see below).

Important! You can specify a different account if necessary, but this setting must identify a user account, not a group account. Restart the policy engine for the changes to take effect.

Internal Email Address Pattern

This setting specifies a semicolon separated list of full or partial email addresses.

When the policy engine processes an email, it first checks the sender's email address against these address patterns. If the sender's address does match an internal address pattern, the policy engine attempts to map the sender onto an existing CA DataMinder user account.

The policy engine **only** expands the sender's email address against the LDAP directory if it matches an address pattern in this list. Typically, you use this setting to detect emails sent by users within your organization.

The recipient details of an email are only expanded against the LDAP directory if the recipient's address matches an item in this list. Therefore, if you want to expand recipients' full details (for example, for policy testing), you must ensure that the list is comprehensive enough to match against all addresses you expect to encounter, for example, Exchange and SMTP addresses.

Address Book (MAPI) lookup operations are only performed for recipient email addresses matching an item in this list.

Note: This setting was formerly the policy engine hub registry value UserSpecificAddrPattern.

If an address does **not** match listed patterns (that is, the sender does not match any of the listed address patterns), the policy engine infers the sender is not a CA DataMinder user and:

- **Either** the policy engine applies the External Sender policy. Specifically, it applies any relevant triggers for **outgoing** emails.
- **Or**, if the External Sender setting is not defined, it applies no policy. That is, it applies no outgoing email triggers.

If the sender address **does** match a listed address pattern but no corresponding CA DataMinder user exists:

- **Either** the policy engine applies the Unknown Internal Sender policy. Specifically, it applies any relevant triggers for **outgoing** emails.
- **Or**, if the Unknown Internal Sender setting is not defined, it applies no policy. That is, it applies no outgoing email triggers.

Important! You must restart the policy engine for changes to this setting to take effect.

Special characters

Address patterns can contain special characters, such as wildcards, spaces, quotes and semicolons:

- **Wildcards:** Use * wildcards to match partially against email addresses. For example, ‘unipr*.com’ and ‘unipraxis’ would both match against ‘lsteel@unipraxis.com’.
- **Spaces:** If an address pattern contains spaces, you must enclose it in "double quotes".
- **Double quotes or semicolons:** If an address pattern contains a double quote or semicolon, you must prefix these characters with a forward slash:

  /" or /;

Default Policy for Files

This setting specifies the name of a CA DataMinder user. A policy engine will apply this user's policy to scanned, captured or imported files if no other means are available to determine the policy participant.

For example, if an Import Policy job for FSA scanning job omit to specify the policy participant, or if the specified user account does not exist, the policy engine applies the Default Policy for Files to the imported or scanned files.

More information:

Configure Policy Engines to Detect Embedded Content Emails (see page 453)
Default Policies for Emails From Unrecognized Senders (see page 488)

Configure the Policy Engine Registry Values

After configuring the local machine policy, you may need to add registry values on the policy engine host machine. These values enable the policy engine to apply the correct user policies when processing an email addressed to a distribution list or an alias e-mail address. To perform such lookup operations, the policy engine needs to connect to an organization's LDAP directory service, typically Active Directory. To configure access to an LDAP directory, you modify values in the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\UserProcess
```

Policy Engine Registry Values

Within the specified registry key, the registry values that you may need to add are:

LookupLDAPServers

Type: REG_MULTI_SZ

Data: Specifies a list of servers hosting an LDAP directory. Specifying multiple host servers provides fault-tolerance and load-sharing to ensure that the policy engine processes events as quickly as possible.

If the policy engine is running in a domain, you can leave this registry value unspecified. By default, the policy engine will automatically detect an Active Directory server.

Server names can be 'plain' or include a domain suffix (UNI-EXCH or UNI-EXCH.UNIPRAXIS.COM). If the LDAP port number is not 389, you can add it after the server name; prefix the port number with a colon (UNI-EXCH:319). You can also prefix the server name with the account credentials used to access the LDAP database. The syntax is:

```
<username>:<password>@<server name>
```

Note: When connecting to a non-Microsoft LDAP server (for example Domino), the username must be the distinguished name of a user with read access to the relevant parts of the directory. For example:

```
cn=Spencer Rimmel:MyPassword@unipraxis
```

Where the user name is Spencer Rimmel, the password is MyPassword, and the server name is unipraxis.

Note: Although anonymous access (not supplying a username and/or password) may allow some LDAP attributes to be accessed, other LDAP attributes may be restricted. That is, your policy engines may only able to access some LDAP attributes if you provide user credentials.

LookupSearchFilter

Type: REG_SZ

Data: If necessary, you can specify that lookup operations are filtered against specific LDAP containers or nodes. Policy engines can automatically detect Active Directory; in this case, the value defaults to:

```
(|(objectCategory=group)(objectCategory=person))
```

For other LDAP directories, it defaults to:

(objectClass=*)

If you override the default filter, ensure that the new filter conforms to RFC 2254.

LookupDirectoryType

Type: REG_SZ

Data: Specifies the type of LDAP directory. Policy engines can automatically detect Active Directory; in this case, the value defaults to GC. For other LDAP directories (including Lotus Domino), it defaults to LDAP.

If necessary, you can override these defaults. For example, set this value to NDS for NetWare NDS or NWCOMPAT for NetWare 3.x.

LookupDirectoryBase

Type: REG_SZ

Data: Specifies the LDAP server's base DN or domain. This value defaults to empty, indicating that the policy engine searches the whole directory for objects. However, you can set this value to a specific base DN, for example, to speed up lookup operations or because the account used to access the LDAP directory only has permission to search a specific subset of the directory.

For Domino, set this registry value to the root domain of your organization, for example:

```
"o=unipraxis"
```

Set Up SecureMail Integration

(Optional) You can configure CA DataMinder to detect, analyze, and apply policy to emails encrypted by Voltage SecureMail. As part of the integration setup for SecureMail, you must edit registry values on your policy engines to identify the SecureMail server.

More information:

Integration with Voltage SecureMail (see page 351)

Monitor Policy Engines

There are various sources of diagnostic information when monitoring policy engines. These are performance counters, log files and diagnostic files.

More information:

PE Hub Log Files (see page 383)

Performance Counters

In the Performance applet, you can add a range of useful performance objects. For policy engines, the following performance object is available:

CA DataMinder Policy Engine

Available only as a single instance. This performance object contains counters for the local policy engine. For example, available counters show the number of pending items (events waiting to be processed by the policy engine), items that have been successfully processed, and failed items.

Note: For policy engine hub performance counters, see Performance counters.

CA DataMinder Policy Engine Processing Core

This performance object contains counters for the processes carried out by all policy engines. For example, available counters show the number of pending items (events waiting to be processed by all policy engines), items that have been successfully processed, and failed items. Other counters relate to average processing rates per email, plus the number of policy engine instances and unique policies currently loaded.

Log Files

Policy engines themselves do not generate log entries.

Uninstall Policy Engines

Important! If a policy engine is installed on the same computer as a CA DataMinder server agent (Exchange, Domino, or Enterprise Vault), you **must** uninstall the server agent before uninstalling the policy engine.

1. First, you must unregister the policy engine on the hub. To do this, edit the ActivePolicyEngines and StandbyPolicyEngines registry values on the hub host server.
2. Stop the CA DataMinder Policy Engine service. Either use the Services applet or run the following command from a command line:

   ```
   net stop wgnpesv
   ```
3. Use Add/Remove Programs to manually uninstall a policy engine. This applet is part of the Control Panel. When the wizard starts, go to the Program Maintenance screen and choose Modify to uninstall the policy engine.

   Note: If you choose Remove in the Program Maintenance screen, this removes all CA DataMinder components, not just the policy engine.

Policy Engine Maintenance

If you need to shut down a policy engine (for example, to carry out maintenance work), we recommend you first unregister it on the policy engine hub to avoid logging any unnecessary error messages.

To shut down a policy engine

1. Locate the ActivePolicyEngines registry value on the policy engine hub.
2. Remove the name or IP address of the machine hosting the policy engine you want to shut down.
3. Shut down the policy engine.

Note: Any events currently being processed are failed and logged accordingly.

Chapter 14: Integration with Voltage SecureMail

This section describes how to set up policy engines to integrate with Voltage SecureMail.

This section contains the following topics:

Overview

You can configure CA DataMinder to detect, analyze, and apply policy to emails encrypted by Voltage SecureMail.

How does the integration work? CA DataMinder intercepts Voltage-encrypted emails passing through an email server and passes copies of these emails to a CA DataMinder policy engine. The policy engine establishes a secure connection to the Voltage SecureMail server, which provides the policy engine with an unencrypted version of the email. The policy engine can then apply policy triggers to the email as normal. When policy processing is complete, the policy engine calls back to the email server agent. The callback instructs the email server agent to either block the encrypted email or allow it to continue.

Note: The original encrypted email remains on the email server until policy processing is complete.

The key components are shown below. For simplicity, this diagram shows the Exchange server agent passing encrypted emails to a single policy engine.

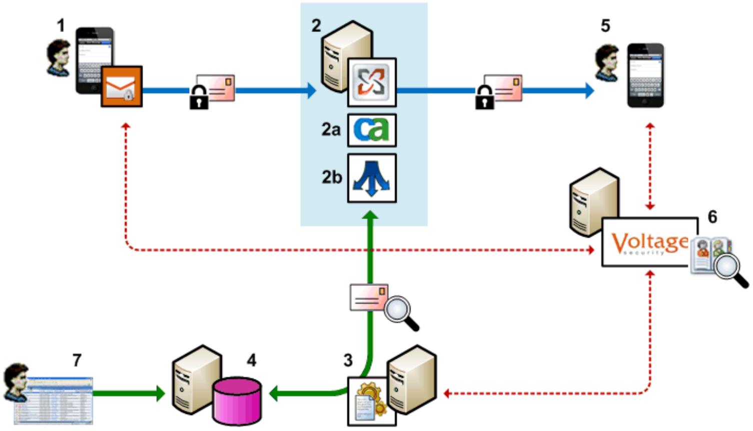

Example deployment architecture: Exchange server agent integration with Voltage SecureMail

A employee (**1**) sends a secure email from their mobile device. The device is running the SecureMail app. The app connects to the Voltage SecureMail server (**6**) to authenticate the sender and encrypt the email.

The encrypted email passes through the Exchange server (**2**). It is intercepted by the Exchange server agent (**2a**) and forwarded to the PE hub (**2b**). The PE hub distributes a copy of the email to a policy engine (**3**).

The policy engine establishes a secure connection to the Voltage SecureMail server (**6**), which sends back an unencrypted version of the email.

The policy engine applies Outgoing Email triggers to the email and calls back to the Exchange server agent with the results of the policy processing (for example, 'block the email'). The resulting email event is replicated to the CMS (**4**).

If policy processing allows the email to continue, the original encrypted email is forwarded to the recipient (**5**). To decrypt the email, the recipient authenticates themself to the Voltage SecureMail server (**6**).

A reviewer (**7**) can search for email events in the iConsole. Unencrypted versions of emails encrypted by SecureMail are available to reviewers and are flagged as 'Secure' in the iConsole.

Requirements

Note the following requirements for CA DataMinder integration with Voltage SecureMail:

Voltage SecureMail

When you enable integration with Voltage SecureMail, CA DataMinder policy engines can access unencrypted versions of emails that were encrypted by Voltage SecureMail clients or the Voltage Zero Download Messenger.

CA DataMinder integrates with the following versions of Voltage SecureMail:

- Voltage SecureMail Mobile Edition
- Voltage SecureMail Cloud Standard Edition
- Voltage SecureMail Cloud Enterprise Edition

Voltage SecureMail Server

You must supply your policy engines with connection details and a shared secret for the Voltage SecureMail web service

Voltage SecureMail Gateway

(Optional) CA DataMinder integration with Voltage SecureMail does not directly require a SecureMail Gateway. However, if you use the CA DataMinder Quarantine Manager to quarantine suspect emails, you must deploy a SecureMail Gateway on your network. The SecureMail Gateway ensures that emails released from quarantine are re-encrypted before they are forwarded to external recipients.

Which agents can detect SecureMail-encrypted emails?

The following CA DataMinder agents can detect and apply policy to emails encrypted by Voltage SecureMail:

- Exchange server agent
- IIS SMTP server agent
- Milter MTA agent (for Sendmail and Postfix email servers)
- CA DataMinder Network (formerly known as the NBA). This agent can detect and apply policy to SMTP emails encrypted by SecureMail.

Policy engines and SSL certificates

CA DataMinder policy engines use the Secure Sockets Layer protocol (SSL) to establish a secure connection to the Voltage SecureMail web service. To ensure that the policy engine and SecureMail web service trust each other, each policy engine must hold a copy of the root certificate that was used to generate the SecureMail certificate.

More information:

Establish an SSL Connection to the SecureMail Web Service (see page 360)

SecureMail Integration and Quarantine Manager

CA DataMinder is able to quarantine emails until they have been approved by an appropriate representative.

If you want to quarantine emails encrypted by Voltage SecureMail, you must deploy the CA DataMinder Quarantine Manager and a Voltage SecureMail gateway.

In the example below, a user sends a secure email to an internal recipient and an external recipient. CA DataMinder detects the email, applies policy, and quarantines the email. When the email is released from quarantine, the version addressed to an external recipient is re-encrypted by the Voltage SecureMail gateway before that version is delivered. The version addressed to an internal recipient is delivered directly and is not re-encrypted.

Note: For simplicity, this diagram omits a policy engine hub and shows the Exchange server agent passing emails to a single CA DataMinder policy engine.

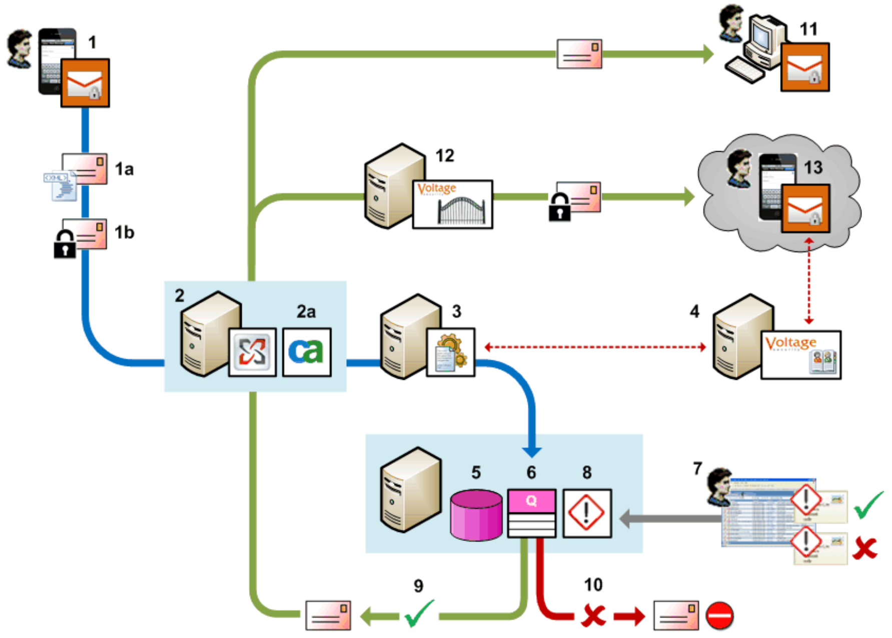

Example: SecureMail integration and quarantined emails on an Exchange server

An employee (**1**) sends a secure email from their mobile device. If they use a Voltage SecureMail clientless solution, the email is not encrypted yet, but contains an x-header to indicate that encryption is required (**1a**). If they use a Voltage SecureMail end-to-end solution, the email is encrypted before it leaves the mobile device.

When the secure email passes through the Exchange server (**2**), the Exchange server agent (**2a**) intercepts the email and passes it to a policy engine (**3**) for analysis.

The policy engine establishes a secure connection to the Voltage SecureMail server (**4**), which sends back an unencrypted version of the email. The policy engine analyzes the email and applies a Quarantine action. A decrypted version of the email is saved in the CMS database (**5**). The email is also saved in a quarantine queue (**6**).

iConsole reviewers can search for quarantined emails and release or reject them (**7**).

The Quarantine Manager (**8**) regularly checks the quarantine queue and forwards released emails to their intended recipients (**9**). At this stage, emails released from quarantine are unencrypted but include an 'encryption request' x-header. Rejected emails are not forwarded (**10**).

When released emails pass through the Exchange server, any emails addressed to internal recipients are sent directly, unencrypted, to the recipient (**11**).

Emails addressed to external recipients are routed through a Voltage SecureMail Gateway (**12**). The Gateway detects the 'encryption request' x-header and re-encrypts the email before sending it to the intended recipient (**13**). To decrypt the email, the recipient must authenticate themself to the Voltage SecureMail server (**4**).

Set Up SecureMail Integration

See the following sections for full details. Briefly, you must:

1. Edit registry values on the policy engine to identify the Voltage SecureMail server and enable integration with SecureMail.
2. Set a shared secret.

   The policy engine uses the shared secret to establish secure connections to the Voltage SecureMail server.
3. Configure the Voltage SecureMail server.

   You must add a new component authentication method and enable the Web Services API.
4. Establish an SSL connection between the policy engine and the SecureMail web service.
5. (Optional) Define a rule on your Voltage SecureMail Gateway to detect emails containing an encryption request x-header.

More information:

Configure the Policy Engine Registry Values (see page 357)
Set a Shared Secret for the Voltage SecureMail Server (see page 358)
Configure the SecureMail Server (see page 359)
Establish an SSL Connection to the SecureMail Web Service (see page 360)
Define Rule for Voltage SecureMail Gateway (see page 362)

Configure the Policy Engine Registry Values

To complete the setup for integration with Voltage SecureMail, you must edit registry values on the policy engine host machine. These values enable the policy engine to connect to the Voltage SecureMail server. To configure access to the Voltage SecureMail server, locate the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\Security Providers\Voltage
```

Within the specified registry key, edit the following registry values:

Enabled

Type: REG_DWORD

Data: Defaults to 0. To enable integration with Voltage SecureMail, set this registry value to 1. This ensures that policy engines attempt to decrypt encrypted emails detected by CA DataMinder email servfer agents. Set this registry value to 0 to disable integration.

ServerURL

Type: REG_SZ

Data: Specifies the protocol, server, and optional port number for the Voltage SecureMail web service that the PE uses to decrypt emails. For example:

```
https://voltage-pp-0000.unipraxis.com
https://voltage-pp-0000.unipraxis.com:425
```

You only need to specify the port number if the web service is *not* using a default port. The defaults are port 80 for HTTP and port 443 for HTTPS.

Set a Shared Secret for the Voltage SecureMail Server

When the policy engine receives a copy of the encrypted email, it sends the email plus a shared secret to the Voltage SecureMail server.

The shared secret (in effect, a password) authenticates the policy engine to the Voltage SecureMail server. It also enables the policy engine to establish a secure connection with the Voltage SecureMail server.

In return, the Voltage SecureMail server provides the policy engine with an unencrypted version of the email.

CA DataMinder provides a command line utility, wgncred.exe, to set account credentials for various components. In this case, you must use it to securely store the shared secret that the policy engine uses to access the Voltage SecureMail server.

Wgncred.exe is installed when you install a policy engine. Find wgncred.exe in the \System subfolder of the CA DataMinder installation folder.

You must run wgncred.exe on each policy engine host server.

To set the SecureMail shared secret

1. From a command prompt, run:

   ```
   wgncred -set
   ```

   A list of components is displayed with their corresponding ID numbers and component identifiers.

2. Choose the Voltage SecureMail Web Service component. The component ID is Voltage.
3. Type the password for this component.

To clear the SecureMail shared secret

1. From a command prompt, run:

   ```
   wgncred -clear
   ```

   A list of components is displayed.

2. Choose the Voltage SecureMail Web Service component. The component ID is Voltage.

More information:

Set Account Credentials with WgnCred.exe (see page 550)

Configure the SecureMail Server

You must configure your Voltage SecureMail server to allow policy engines to authenticate and decrypt emails. In particular, you must:

- Add a new component authentication method.
- Enable the Web Services API

Add a component authentication method to the SecureMail server

A component authentication method enables a trusted component to authenticate itself and download the keys of any user automatically when needed.

In this case, the trusted components are CA DataMinder policy engines. The component authentication method uses the shared secret that you previously stored on your policy engines.

Use the Voltage Administration Management Console to add a component authentication method. You must supply:

- The email patterns that you want to match.
- The IP address of the policy engine that you use to decrypt emails.
- An 'authentication token secret'. This must match the shared secret that you set on the policy engine.

When the new component authentication method is added to the Component Authentication table, the method name defaults to 'Component'. We recommend that you change 'Component' to the name of the policy engine host server.

Note: For full details on configuring a new component authentication method, see the *Voltage SecureMail Management Console Administrator’s Guide*.

Enable the SecureMail Web Services API

After creating a new component authentication method, you must enable the SecureMail Web Services API. This API provides an interface for interacting with SecureMail encryption services over HTTPS.

In this case, the CA DataMinder policy engine that you specified in the component authentication method uses the Web Services API to access the SecureMail encryption services.

Use the Voltage Administration Management Console to enable the Web Services API on each host.

Note: For details, see the 'Enabling Web Services API' section in the *Voltage Web Services API Supplement*.

Establish an SSL Connection to the SecureMail Web Service

CA DataMinder policy engines use the Secure Sockets Layer protocol (SSL) to establish a secure connection to the Voltage SecureMail web service. To ensure that the policy engine and SecureMail web service trust each other, each policy engine must hold a copy of the root certificate that was used to generate the SecureMail certificate.

To install the SecureMail root certificate on your policy engine

Use the Microsoft Management Console to manage root certificates on your policy engine host servers. If required, you can export the root certificate from your SecureMail server and import it onto your policy engine host servers. Exported certificates are saved as files. Copy the file to your policy engine host server and then double-click the file to launch the certificate import wizard.

For further details, search for 'configuring SSL certificates' and 'importing CA and root certificates' in the *Voltage SecureMail Management Console Administrator Guide*.

How does the SSL connection work?

SSL communications between the policy engines and SecureMail web service are encrypted using public/private key encryption.

When you install a Voltage SecureMail server, an SSL certificate is assigned to the SecureMail web service. This certificate was generated from a root certificate. The root certificate is signed by a certificate authority that is trusted by SecureMail (the ‘trusted certificate authority’).

When the policy engine establishes an SSL connection, it obtains a public key from the same root certificate that was used to generate the SecureMail certificate. A copy of this root certificate must be already installed on the policy engine host server.

Next, the policy engine requests the SecureMail certificate. Because the SecureMail certificate is signed by a certificate authority that the policy engine trusts, the policy engine proceeds with the connection and encrypts the communication using the public key.

The SecureMail web service then uses a private key to decrypt the encrypted communication.

Note: Browsers ship with, and regularly update, a set of certificates signed by trusted certificate authorities to ensure that connections can be verified.

More information:

What is SSL? (see page 361)
About Certificates (see page 361)

What is SSL?

The Secure Sockets Layer protocol (SSL) helps ensure that a network transaction (such as a web request) is only serviced by the intended network host (such as a web site). SSL also prevents transmitted data from being intercepted by a third party. The connection does this by encrypting the traffic using public/private key encryption.

- You obtain a public key via a certificate which is validated against a trusted certificate authority.
- Each client holds a well-known public certificate of an organization that it trusts (the certificate authority). The client then requests the certificate of the server that it needs to connect to. If the server's certificate is correctly signed by a trusted certificate authority, the client proceeds with the connection and negotiates the encrypted communications channel.

Typical SSL applications include online purchasing and webmail, and an increasing number of web sites and applications (such as instant messaging). In particular, the widespread use of social networking sites is a major cause for concern regarding data loss. Your ability to analyze data transmitted from your company network to these external networks is increasingly important.

About Certificates

A certificate is a small file containing data about a website or network host. The certificate is signed to prevent falsification and contains a chain of responsibility (the certification path) that allows a browser or network client to verify the certificate even if the browser or client only has local access to the top-level (or root) certificate in the chain.

Web browsers provide the ability to view the certificate of a website and verify that the certificate is valid. Browsers ship with, and regularly update, a set of Certificate Authority certificates to help ensure that verification can be performed.

Define Rule for Voltage SecureMail Gateway

If you use a Voltage SecureMail clientless solution (also known as 'SecureMail FlagSecure'), emails are not encrypted when they are sent but contain an x-header to indicate that encryption is required. Therefore you must define a rule on your SecureMail Gateway to encrypt the email.

Note: If you use a Voltage SecureMail end-to-end solution, emails are already encrypted before they leave the workstation or mobile device.

You *must* use 'x-voltage: encrypt'

To support CA DataMinder integration with SecureMail, you *must* define a SecureMail Gateway rule to detect the following x-header:

```
x-voltage: encrypt
```

This is the default x-header for SecureMail FlagSecure. *Do not specify a different x-header!*

Note: For details about configuring gateway rules, see the *Voltage SecureMail Management Console Administrator Guide*.

Why is this necessary?

CA DataMinder policy engines can only recognize 'x-voltage: encrypt' x-headers in SecureMail emails. They cannot recognize different x-headers.

CA DataMinder uses this x-header to mark SecureMail emails as 'Encrypted' when a user reviews these emails in the iConsole.

Note: If you use a different x-header, CA DataMinder still applies policy to the email. And the email is still encrypted by the SecureMail Gateway. However, CA DataMinder does not mark the email as 'Encrypted' in the iConsole.

Chapter 15: Policy Engine Hub Advanced Configuration

A policy engine hub allocates emails to individual policy engines. Normally, a hub requires very little configuration. Basic instructions for configuring a hub are included in the various *Integration Guides* in the CA DataMinder bookshelf.

This chapter describes the hub internal architecture and provides advanced configuration instructions.

Note: Advanced hub configuratuion is not necessary in a typical CA DataMinder deployment.

This section contains the following topics:

Policy Engine Hubs Overview

The role of a policy engine hub is to allocate emails to individual policy engines. Policy engine hubs can accept e-mails from email and archive server agents (Exchange, Domino and Enterprise Vault), the Network Boundary Agent (NBA), and also from Event Import.

A policy engine hub handles each email with minimal delay. It distributes email processing across multiple remote policy engines to optimize load-balancing and maximize throughput. It can also handle hardware failures on remote policy engines, seamlessly redistributing events to other policy engines if necessary.

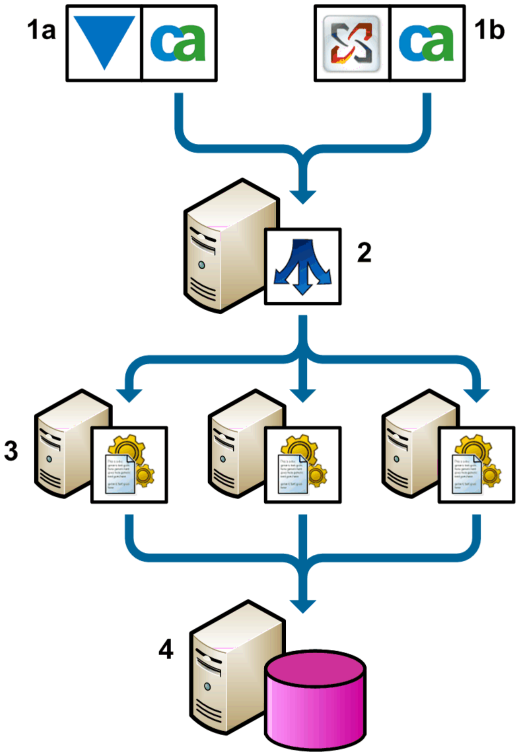

Example policy engine hub deployment

1 Event sources, including CA DataMinder Event Import (**1a**) and e-mail and archive server agents (**1b**) forward emails to the policy engine hub (**2**). The hub distributes e-mails to individual policy engine (**3**) for processing. The resulting events are replicated up to the CMS (**4**).

Policy Engine Hub Architecture

A policy engine hub can have multiple event queues, configurable by email size, with each queue serving multiple policy engines. Distributing emails across multiple queues in this way allows fast-track processing for emails of different sizes. Registry settings on the hub host server define the queue size bands and specify which policy engines are assigned to each queue—see Hub event queues (see page 367) for details. When a policy engine has finished processing an email, it is passed back to the hub before being returned to the source application.

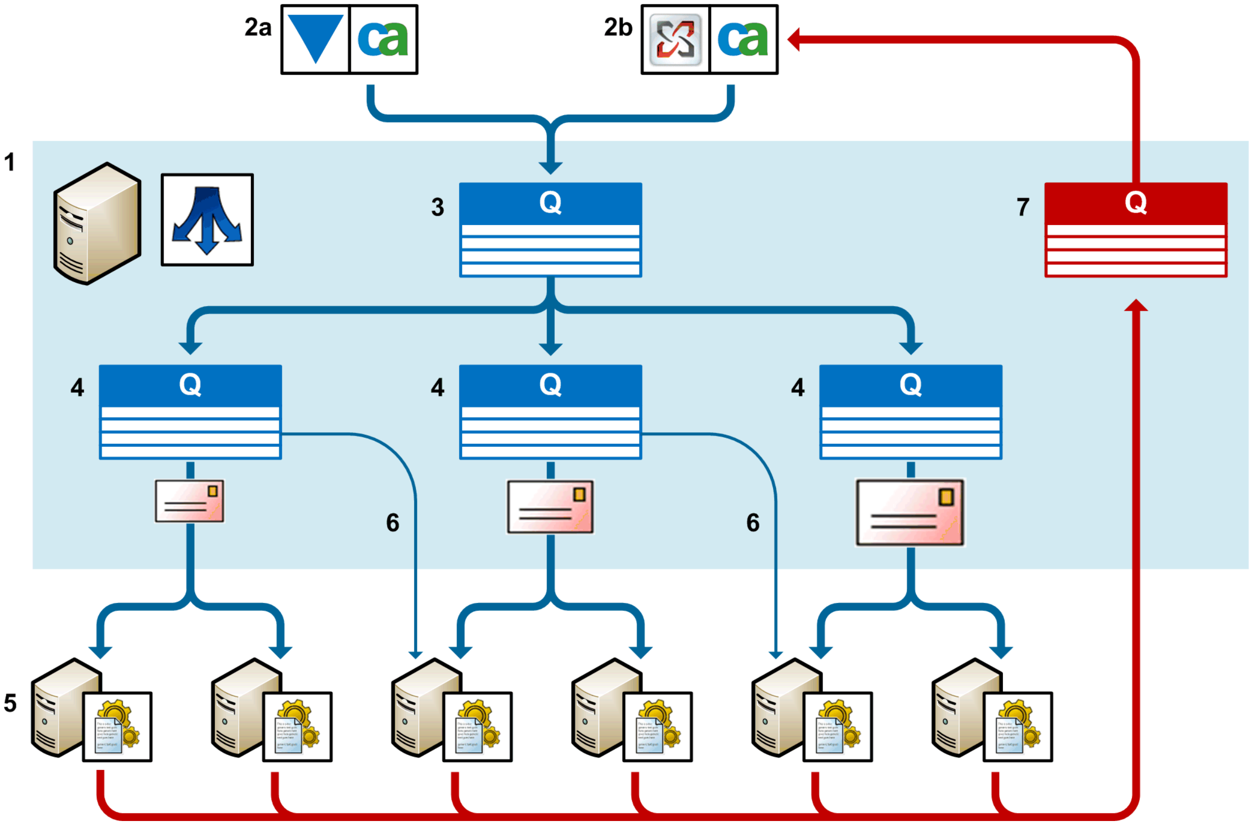

Policy engine hub architecture

A policy engine hub (**1**) can accept emails from various sources, including Event Import (**2a**) and an email server agent (**2b**).

Emails arriving at the hub are added to the input queue (**3**). The hub then assigns each email to an event queue (**4**). A hub can have multiple event queues (three in this example), configurable by e-mail size. Registry settings on the host machine (**1**) define the size band for each queue.

Each queue can be served by multiple policy engines (**5**). Registry settings on the host machine (**1**) specify the policy engines allocated to each queue. To minimize processing times, if a queue is empty then idle policy engines assigned to that queue are also permitted to poach messages from other queues (**6**), but only from queues with a smaller maximum size limit.

After a policy engine has successfully processed an email, the e-mail is passed back to the completion queue (**7**) on the hub before being finally returned to the source application (**2a** or **2b**).

More information:

Hub Event Queues (see page 367)

Hub Event Queues

For each hub, there is always a default queue that can hold messages of unlimited size. To allow fast-track processing of small messages, you need to create one or more additional, size-restricted queues. You do this by editing the registry.

Queue Settings

For each additional queue, you can specify:

- **The maximum size of queued messages**

  If a message exceeds this maximum size limit, the hub automatically assigns the message to the next appropriate queue.

  For example, two additional queues are defined: Small (for messages up to 10 KB) and Medium (up to 100 KB). If a 15 KB message arrives at the hub, it is immediately assigned to the Medium queue; conversely, if a 2 MB message arrives, it is immediately assigned to the default queue, which handles messages of unlimited size.

- **Dedicated policy engines available to process queued messages**

  You can specify separate lists of policy engines available for processing each queue. However, to minimize processing times, if a queue is empty then any idle policy engines assigned to that queue are also permitted to poach messages from other queues (but only from queues with a smaller maximum size limit).

Note: The default queue is not represented in the registry by a dedicated registry key and has no maximum size limit for queued messages; by contrast, any additional queues **are** represented by a dedicated registry key.

Monitoring the Queue Status

To monitor the status of individual event queues, check the relevant counters in the CA DataMinder Hub Queues performance object.

More information:

Modify the Hub Registry Values (see page 372)

Registry Flow Chart: Email Processing on the Hub

The diagram below shows how the crucial registry values for the policy engine hub and the Exchange or Domino server agent operate in a strict sequence, with the event finally passing to a policy engine.

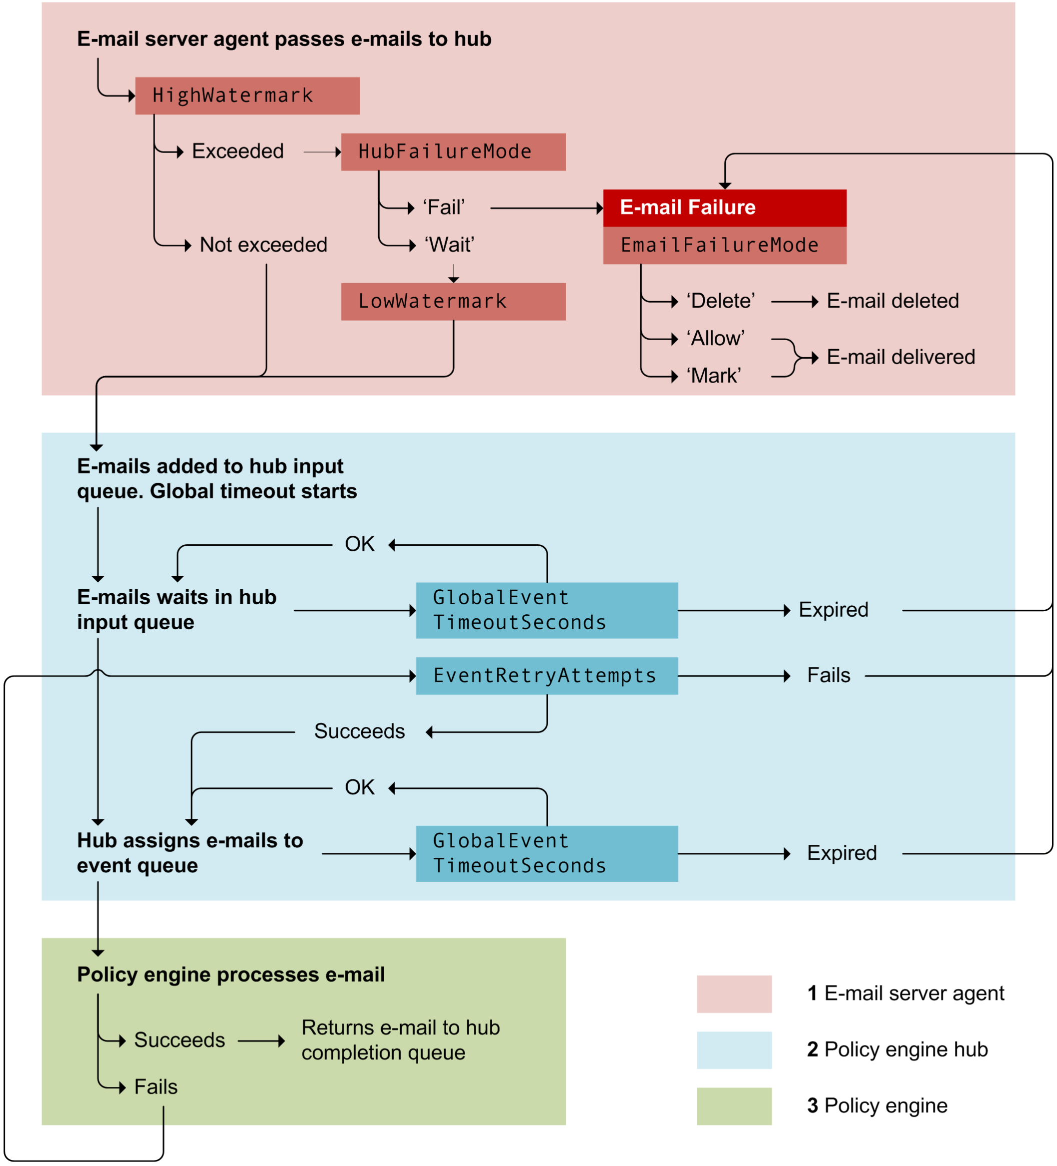
E-mail server agent passes e-mails to hub
HighWatermark
Exceeded
HubFailureMode
'Fail'
'Wait'
E-mail Failure
EmailFailureMode
Not exceeded
LowWatermark
'Delete'
E-mail deleted
'Allow'
'Mark'
E-mail delivered
E-mails added to hub input queue. Global timeout starts
OK
E-mails waits in hub input queue
GlobalEvent TimeoutSeconds
Expired
EventRetryAttempts
Fails
Succeeds
OK
Hub assigns e-mails to event queue
GlobalEvent TimeoutSeconds
Expired
Policy engine processes e-mail
Succeeds
Returns e-mail to hub completion queue
Fails
1 E-mail server agent
2 Policy engine hub
3 Policy engine

More information:

Modify the Hub Registry Values (see page 372)

Deploy the Policy Engine Hub

Deploying a policy engine hub requires the following tasks:

To deploy the policy engine hub

1. Deploy your policy engines.
2. Set up the event source.

   For example, to integrate with Exchange Server, you must deploy the CA DataMinder Exchange server agent. See the relevant Integration guidies for installation instructions.
3. Install the PE hub.

   A hub is installed automatically when you install a CA DataMinder server agent or Import Policy.
4. Configure the hub.
   a. Assign a security privilege to the PE domain user.
   b. Edit the hub registry values.

More information:

Deploying Policy Engines (see page 338)
Install the Policy Engine Hub (see page 370)
Configure the Policy Engine Hub (see page 370)

Hub Host Machine Requirements

The following are requirements for the machine hosting the policy engine hub:

PE domain user must be local administrator

The PE domain user must be a member of the local Administrators group on the machine hosting the policy engine hub. Confirm that this is so before installing the policy engine hub.

More information:

Specify a PE Domain User (see page 340)
CA DataMinder Installations on 64-bit Machines (see page 580)

Install the Policy Engine Hub

A policy engine hub is installed automatically when you install any of the following CA DataMinder agents. See the relevant chapters for installation instructions.

- **Email Server Agents**

  Domino Server Agent

  Exchange Server Agent

  IIS SMTP Agent

- **Archive Agents**

  EMC SourceOne

  Symantec Enterprise Vault

- **Other**

  ICAP Agent

  External Agent API

In addition, you can optionally install a remote Policy Engine Connector (a type of hub) when you install the File Scanning Agent.

Configure the Policy Engine Hub

After installing the policy engine hub, you must:

- Assign the ‘Log on as a batch job’ security privilege to the PE domain user on the host machine for the policy engine hub.
- Configure the policy engine hub by modifying the associated registry values on the host machine.

More information:

Registry Flow Chart: Email Processing on the Hub (see page 368)

Assign Security Privilege to the PE Domain User

The PE domain user requires the ‘Log on as a batch job’ security privilege. This permits policy engines on remote machines to access the policy engine hub. To assign this privilege:

1. Ensure that you are logged on with local administrator rights on the host machine for the policy engine hub.
2. On the host machine, open the Local Security Policy applet or, if this machine is a domain controller, open the Domain Controller Security Policy applet. Both applets are available in Administrative Tools.
3. Expand the Local Policies branch and select User Rights Assignment. This security area determines which users have logon privileges on the local computer.
4. Assign the ‘Log on as a batch job’ privilege to the PE domain user.

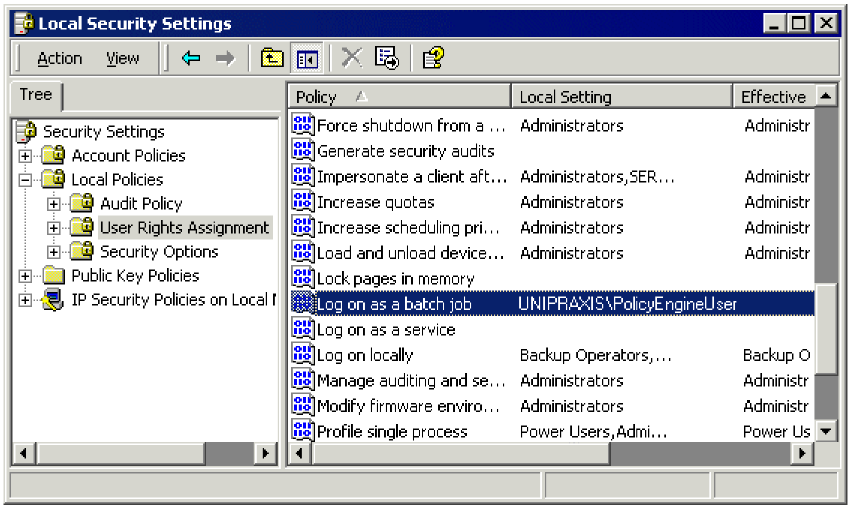

- Security Settings
  - Account Policies
  - Local Policies
    - Audit Policy
    - User Rights Assignment
    - Security Options
  - Public Key Policies
  - IP Security Policies on Local I

Policy	Local Setting	Effective
Force shutdown from a ...	Administrators	Administr
Generate security audits		
Impersonate a client aft...	Administrators,SER...	Administr
Increase quotas	Administrators	Administr
Increase scheduling pri...	Administrators	Administr
Load and unload device...	Administrators	Administr
Lock pages in memory		
Log on as a batch job	UNIPRAXIS\PolicyEngineUser	
Log on as a service		
Log on locally	Backup Operators,...	Backup O
Manage auditing and se...	Administrators	Administr
Modify firmware enviro...	Administrators	Administr
Profile single process	Power Users,Admi...	Power Us

Local Security Policy applet

In this applet, the left pane shows the Local Policies branch and the User Rights Assignment node. The 'Log on a batch job' policy, or privilege, is shown in the right pane.

More information:

Specify a PE Domain User (see page 340)

Modify the Hub Registry Values

To configure the policy engine hub, you need to modify values in the following registry key.

```
HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA
DataMinder\CurrentVersion\Policy Engine Hub
```

Changes to the Policy Engine Hubs key take immediate effect. Below this registry key there are various subkeys. The key structure is shown below:

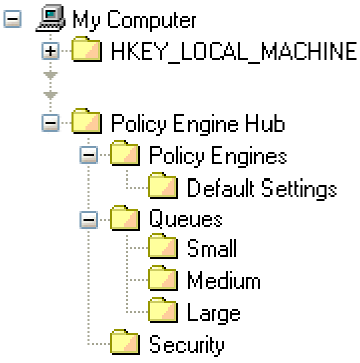

```
My Computer
  HKEY_LOCAL_MACHINE

  Policy Engine Hub
    Policy Engines
      Default Settings
    Queues
      Small
      Medium
      Large
    Security
```

Policy engine hub registry keys: Registry values in the Policy Engine Hub key and its subkeys are described in the following sections.

Important! If the hub is deployed with the Exchange 2007 or 2010 server agent on a 64-bit machine, the registry key is slightly different.

More information:

Policy Engine Hub Registry Values (see page 373)
CA DataMinder Installations on 64-bit Machines (see page 580)

Policy Engine Hub Registry Values

The table below lists the available registry values for policy engine hubs.

- Policy Engine Hub key (see page 374)
  `EventLoggingLevel`
  `EventRetryAttempts`
  `GlobalEventTimeoutSeconds`
  `HighWaterMarkMB`
  `HighWaterMarkEventCount`
  `LogFilePath`
  `LogMaxNumFiles`
  `LogMaxSizeBytes`
  `LowWaterMarkMB`
  `LowWaterMarkEventCount`
  `NoPEFailTimeoutSeconds`
  `OperationalLoggingLevel`
  `PECallTimeoutMilliseconds`
- DefaultSettings subkey (see page 378)
  `HeartbeatPeriodMilliseconds`
  `MetricsPeriodMilliseconds`
  `ReconnectTimeoutSeconds`
- Queues key (see page 380)
  `ActivePolicyEngines`
  `AdditionalQueues`
  `StandbyPolicyEngines`
- <Queue name> subkey (see page 381)
  `ActivePolicyEngines`
  `MaxSizeBytes`
  `StandbyPolicyEngines`
- Security key (see page 381)
  `NTNetworkDomain`
  `NTNetworkUser`

More information:

Modify the Hub Registry Values (see page 372)

Policy Engine Hub Key

The Policy Engine Hub registry key contains the following registry values:

EventLoggingLevel

Type: REG_DWORD

Data: Defaults to 2. This determines the level of logging for message processing. For example, you can configure the hub to only log errors or warning system messages.

Log entries are written to the wgnphub_<date>.log file, where <date> is the date and time when the log file was created; the file location is set by the LogFilePath registry value. The supported logging levels are:

1 - Errors only

2 - Errors and warnings

3 - Errors and warnings, plus informational and status messages

Note: Setting EventLoggingLevel=3 will cause the log file to grow extremely rapidly. This level of logging is provided for testing purposes only.

EventRetryAttempts

Type: REG_DWORD

Data: Defaults to 4. Determines how many times the hub attempts to pass an email to a policy engine before it is flagged as an ‘email failure’ and passed back to the Exchange or Domino server agent.

This **only** applies to email failures caused by a problem with the policy engine (such as a host machine crash) or with the email itself (that is, some unexpected condition that prevents the policy engine from analyzing the email).

It does not apply to email failures resulting from the time-out GlobalEventTimeoutSeconds expiring or because the HighWaterMarkEventCount or HighWaterMarkMB thresholds have been exceeded.

How the email server agent handles ‘email failures’ depends on its EMailFailureMode registry value.

Note: For Import Policy jobs, we recommend a value of 0, to stop the policy engine retrying failed events.

GlobalEventTimeoutSeconds

Type: REG_DWORD

Data: Defaults to 300. A time-out (in seconds) that specifies how long an email can stay in the event queue on the hub or policy engine before it is flagged as an ‘email failure’ and passed back to the source component (typically the Exchange server agent or, for Import Policy jobs, Event Import).

How the Exchange server agent handles ‘email failures’ depends on its EMailFailureMode registry value (see link above). For Event Import, the handling of import failures depends on the type of import operation.

Note: If Import Policy is installed, the default for this value changes to 21,600 seconds (6 hours), so that events are less likely to timeout.

HighWaterMarkMB

Type: REG_DWORD

Data: Defaults to 400. The maximum amount of memory (in MB) that can be allocated to the various hub event queues. If any queue lengthens and causes the allocated memory to rise above this maximum level, the hub temporarily suspends normal operations until the queue shortens (see LowWaterMarkMB).

While normal operations are suspended, the hub’s behavior depends on the HubFailureMode registry value. Either the hub delays new emails from the email server agent, or it returns them to the server agent as ‘email failures’.

How the server agent handles ‘email failures’ depends on its EMailFailureMode registry value (see link above).

Note: Note the following:

- Memory-based throttling operates in parallel with event-based throttling (for details, see HighWaterMarkEventCount). If either threshold is exceeded, hub operations are suspended.
- For Import Policy jobs, we recommend a low value (for example, 40MB) to ensure a steady stream of events.

HighWaterMarkEventCount

Type: REG_DWORD

Data: Defaults to 400. The maximum total number of events that can be allocated to the various hub event queues.

If any queue lengthens and causes the event count to rise above this maximum level, the hub temporarily suspends normal operations until the queue shortens (see LowWaterMarkEventCount below).

While normal operations are suspended, the hub's behavior depends on the HubFailureMode. Either the hub delays new emails from the Exchange or Domino server agent, or it returns them to the server agent as 'email failures' (whose handling is dependent on the EMailFailureMode; see link above)

Note: Note the following:

- Event-based throttling operates in parallel with memory-based throttling (see HighWaterMarkMB). If either threshold is exceeded, hub operations are suspended.
- For Import Policy jobs, we recommend a low value (for example, 40) to ensure a steady stream of events.

LogFilePath

Type: REG_SZ

Data: Defaults to empty. This specifies the folder you want to write log files to. The PE domain user must have write access to the specified folder.

If the path is not defined, the log file is saved in the default location.

In the current CA DataMinder release, log files are typically saved in CA's \data\log subfolder. On 32-bit machines, find this subfolder in the Windows All Users profile. On 64-bit machines, find this subfolder below the \ProgramData folder.

LogMaxNumFiles

Type: REG_DWORD

Data: Defaults to 10. This specifies the maximum number of log files. When the maximum number of log files exists and the maximum size of the latest is reached (see below), the oldest log file is deleted to enable a new one to be created.

LogMaxSizeBytes

Type: REG_SZ

Data: Defaults to 1,000,000. This specifies the maximum size (in bytes) for each log file. When the current log file reaches its maximum size, the policy engine hub creates a new log file. Log entries are written to a wgnphub_<date>.log file—for details see EventLoggingLevel above.

LowWaterMarkMB

Type: REG_DWORD

Data: Defaults to 200. The total amount of memory allocated to the various hub queues (in MB) that triggers a resumption in hub operations.

If normal hub operations are suspended because the HighWaterMarkMB has been exceeded (see above), they only resume when the queues shorten and allocated memory falls back below the LowWaterMarkMB amount.

Note: For Import Policy jobs, we recommend a low value (for example, 20 MB) to ensure a steady stream of events.

LowWaterMarkEventCount

Type: REG_DWORD

Data: Defaults to 300. The total number of events allocated to the various hub queues that triggers a resumption in hub operations.

If normal hub operations are suspended because the HighWaterMarkEventCount has been exceeded (see above), they only resume when the queues shorten and the event count falls back below the LowWaterMarkEventCount amount.

Note: For Import Policy jobs, we recommend a low value (for example, 20) to ensure a steady stream of events.

NoPEFailTimeoutSeconds

Type: REG_DWORD

Data: Defaults to 60. This specifies how long (in seconds) the hub waits after it detects there is no active policy engine available for a specific queue before it times out events in that queue (that is, flags them as email failures).

When this timeout expires, all events in the queue are immediately flagged as email failures. This overrides the GlobalEventTimeoutSeconds (see above).

How the Exchange or Domino server agent handles ‘email failures’ depends on the EMailFailureMode; see link above)

Note: *We recommend that you do not change the default timeout of 60 seconds.*

OperationalLoggingLevel

Type: REG_DWORD

Data: Defaults to 3. This determines the level of logging for hub operations. For example, typical log entries cover hub installation, creating or deleting queues, the failure or suspension of policy engines, and so on.

Log entries are written to the wgnphub_<date>.log file, where <date> is the date and time when the log file was created; file name and location details and supported logging levels are the same as for EventLoggingLevel (see above).

PECallTimeoutMilliseconds

Type: REG_DWORD

Data: Defaults to 10000. A time-out (in milliseconds) that specifies how long the hub will wait to connect to a policy engine for configuration purposes before it cancels the call and assumes that the policy engine is currently unavailable.

(More information:

Uninstall Policy Engine Hubs (see page 386)

Policy Engines Subkey

Below the Policy Engine Hub registry subkey (see the previous section), there is a Policy Engines subkey. This subkey contains no values; instead, it contains the DefaultSettings subkey and, optionally, a <Machine name> subkey.

More information:

DefaultSettings Subkey (see page 378)
<Machine name> Subkey (see page 379)

DefaultSettings Subkey

Below the Policy Engines registry subkey (see the previous section), there is a DefaultSettings subkey. Values in this subkey define the default configuration for all policy engines.

HeartbeatPeriodMilliseconds

Type: REG_DWORD

Data: Defaults to 40,000. This specifies how often (in milliseconds) the policy engine sends a heartbeat signal to the policy engine hub. If the hub does not receive three successive heartbeat signals, it infers there is a problem with the policy engine.

MetricsPeriodMilliseconds

Type: REG_DWORD

Data: Defaults to 40,000. This specifies how often (in milliseconds) the policy engine returns metrics to the policy engine hub. This value must be an integer multiple of HeartbeatPeriodMilliseconds.

ReconnectTimeoutSeconds

Type: REG_DWORD

Data: Defaults to 600. If a policy engine does not restart immediately when the hub tries to connect to it, this value specifies how long (in seconds) the hub waits between subsequent reconnection attempts.

Note: This timeout is only applicable if the hub fails in its initial attempts after startup to connect to a policy engine (for example, because the policy engine host machine is switched off).

<Machine name> Subkey

To override the default policy engine configuration, you can create a <Machine> subkey below the Policy Engines registry subkey. Note that <Machine> is the host machine for the policy engine you want to customize. The <Machine> subkey can contain customized versions of any registry value in the DefaultSettings subkey.

More information:

DefaultSettings Subkey (see page 378)

Queues Key

Below the Policy Engine Hub registry key, is the Queues subkey. It contains the following registry values, plus various subkeys, one for each message queue supported by the hub.

ActivePolicyEngines

Type: REG_SZ

Data: Defaults to <localhost>. This specifies a comma-separated list of names or IP addresses of machines hosting policy engines available to the *default queue*.

AdditionalQueues

Type: REG_SZ

Data: Defaults to null. This specifies a comma-separated list of additional message queues available to the policy engine hub. These queues are in addition to a default queue, which is always available. The example registry architecture diagram shows three additional queues: Small, Medium and Large.

Note: If you edit this registry value while the hub service is running, CA DataMinder automatically creates new registry subkeys for each additional queue. However, if you want to configure your hub in advance, you must manually create a subkey for each additional queue, and you must also add and configure the necessary registry values to this new subkey.

StandbyPolicyEngines

Type: REG_SZ

Data: Specifies a comma-separated list of names or IP addresses of machines, available to the default queue, and which can be used by the hub if an ‘active’ policy engine is unavailable.

More information:

Policy Engine Hub Architecture (see page 365)

<Queue name> Subkey

In the example registry architecture diagram, the hub supports three additional queues: Small, Medium and Large. There is a separate subkey for each of these queues. These queues are in addition to a default queue, which is always available and does not have its own registry subkey.

You must create these additional subkeys manually or, if the hub service is running when you edit the AdditionalQueues registry value, the subkey is created automatically. Each of these manually added subkeys contains the following registry values.

ActivePolicyEngines

Type: REG_SZ

Data: Defaults to <localhost>. This specifies a comma-separated list of names or IP addresses of machines hosting policy engines available to the *specified queue*.

StandbyPolicyEngines

Type: REG_SZ

Data: Specifies a comma-separated list of names or IP addresses of machines, available to the *specified queue*, and which can be used by the hub if an 'active' policy engine is unavailable.

MaxSizeBytes

Type: REG_DWORD

Data: Defaults to zero. This specifies the maximum size (in bytes) of message events that can be processed by the specified queue. If a message is too large for this queue, it is assigned to the next size queue. You must change the default value of MaxSizeBytes. If it remains set to zero, this queue can never process any messages!

More information:

Policy Engine Hub Architecture (see page 365)

Security Key

Below the Policy Engine Hub registry key, there is a Security subkey. You do not need to modify the values in this subkey because they are managed by the policy engine hub. But for reference, the values are:

NTNetworkDomain

Type: REG_SZ

Data: Domain name. Created automatically when you configure the PE domain user. *Do not modify this value directly.*

NTNetworkUser

Type: REG_SZ

Data: User name. This value is created automatically when you configure the PE domain user. *Do not modify this value directly.*

Hub Maintenance

If you need to shut down the policy engine hub (for example, while you upgrade the Exchange or Domino server agent), you must follow the recommended procedure to ensure that no emails are inadvertently deleted or transmitted without being monitored by CA DataMinder.

Stopping the Policy Engine Hub

1. You must suspend normal Exchange, IIS SMTP, or Domino operations before you stop the policy engine hub service. There are several ways to do this for Exchange 2003 and IIS SMTP, but we recommend that you stop Internet Information Services (IIS). This is because you cannot upgrade the email server agent .DLL file while IIS is running. For Exchange 2007 and 2010 we recommend that you stop the Microsoft Exchange Transport service.
2. Stop the CA DataMinder Policy Engine Hub service. You can now perform any necessary maintenance or upgrades on the host machine.

Restarting the Policy Engine Hub

You must restart the services in the reverse order to which they were stopped. That is, restart the CA DataMinder Policy Engine Hub service, then restart IIS (or the Microsoft Exchange Transport service).

Consequences If You Stop the Hub Before IIS

If you stop the policy engine hub before stopping IIS (when applicable), there is a risk that emails may be deleted or transmitted without being monitored by CA DataMinder, or that emails may be imported twice. These consequences can arise if the Exchange or Domino server agent passes emails to the hub while it is shutting down.

When the email server agent receives no response from the hub, it infers there has been an email failure and uses the EMailFailureMode registry value to determine how to handle the email: this value can be set to Delete, Allow, or Mark (see EMailFailureMode).

Alternatively, the timing of the hub shutdown may be such that an email is sent to a policy engine for processing immediately before the hub shuts down. The policy engine successfully processes the email but is unable to notify the hub. Consequently, the email is resubmitted to a policy engine when the hub restarts.

Monitor Policy Engine Hub Activity

There are various sources of diagnostic information when monitoring policy engine hubs. These are performance counters, log files and diagnostic files.

PE Hub Log Files

Note: If the hub is deployed with the Exchange 2007 or 2010 server agent on a 64-bit machine, note that the server agent and hub log files may be in different locations.

Policy Engine Hub

The policy engine hub service writes entries to the log file, wgnphub.log. These detail progress as each email is processed. This log file is in the same folder as the hub executable, wgnphub.exe, typically installed to the \System subfolder in the CA DataMinder installation folder on the Exchange or Domino server.

The hub writes entries to this log file using the PE domain user account. By default (on NTFS file systems), security for wgnphub.log has been configured to give the PE domain user Read and Write access to this file.

You configure the policy engine hub log files by editing the relevant registry values..

Exchange, IIS SMTP, and Domino Server Agents

For details about the Exchange, IIS SMTP, and Domino server agent log files, see Log files for email server agents .

More information:

CA DataMinder Installations on 64-bit Machines (see page 580)
Policy Engine Hub Registry Values (see page 373)

Hub Performance Counters

The policy engine hub includes three performance objects that are useful when diagnosing hub problems. In the Performance applet (accessible from Administrative Tools), you can add a range of useful performance objects. For each performance object, you can specify which counters and, where relevant, instances you want to view.

Note: On a 64-bit system, all CA DataMinder performance counters, including counters for a 32-bit policy engine hub, are supported in the default 64-bit version of the Performance applet. See the following section for details.

Policy Engine Performance Object

For policy engine performance object details, see 'Monitor Policy Engines' in the *Platform Deployment Guide*.

Hub Performance Objects

For policy engine hubs, the following performance objects are available:

CA DataMinder Hub

Available only as a single instance. This contains counters for the policy engine hub itself. For example, you can see the number of active, standby and connected policy engines, the number of pending events in the hub input queue, the number of event failures, and the total memory allocated to events in the queue.

CA DataMinder Hub Connections

Multiple instances available; one per policy engine. This performance object contains counters for hub connections to individual policy engines. To view counters for a connection to a specific policy engine, select its corresponding instance.

Available counters show statistics such as the number and rate at which events are being passed to the policy engine and the internal state of the policy engine, for example, ‘inactive’, ‘processing’ and ‘dead’ (if you encounter a problem with a policy engine, Technical Support may ask for details of its internal state).

CA DataMinder Hub Queues

Multiple instances available; one per event queue on the hub. For each instance, this performance object contains counters for individual policy engines. To view counters for a specific queue, select its corresponding instance.

Available counters show statistics such as the queue size band (in bytes), the number of active and standby policy engines available to process the queue, and the number of items assigned to the queue.

More information:

Performance Counters (see page 349)

Performance Counters Now Supported in 64-bit Perfmon.exe

On 64-bit systems, all CA DataMinder performance counters are now supported in the default 64-bit version of the Performance applet (perfmon.exe).

Previously on 64-bit systems, most CA DataMinder performance counters were only supported in a 32-bit version of the Performance applet. This anomaly particularly affected performance counters for the policy engine hub when the hub was installed on a 64-bit Exchange 2007 or 2010 server.

If you previously used the 32-bit Performance applet to monitor CA DataMinder components on a 64-bit system, you must switch to the 64-bit Performance applet after upgrading to CA DataMinder 14.1.

Background

On a 64-bit Windows operating system, there are two versions of perfmon.exe:

- A 64-bit version is available in the \Windows\System32 folder. This is the main System folder for the 64-bit operating system.

  Why 'System32'? The folder name, an apparent misnomer, is a legacy of the folder naming scheme in earlier Windows operating systems.

- A 32-bit version is available in the \Windows\SysWOW64 folder.

  Why ‘WOW64’? On a 64-bit Windows operating system, there is an emulation of a 32-bit operating system called ‘Windows on Windows 64’, or WOW64.

Uninstall Policy Engine Hubs

Important! If a policy engine hub is installed on the same computer as an Exchange, Domino, or Enterprise Vault server agent, you must uninstall the server agent before uninstalling the policy engine hub.

To uninstall a policy engine hub, you must uninstall its associated server agent; the hub is then uninstalled automatically. Use Add/Remove Programs to manually uninstall the Exchange or Domino server agents. This applet is part of the Control Panel.

1. In Add/Remove Programs, select CA DataMinder Integration Agents and click Change.
2. When the wizard starts, go to the Program Maintenance screen and choose Modify.

   Note: If you choose Remove, this removes all CA DataMinder components, not just the Exchange or Domino server agents.
3. In the Custom Setup screen, choose the Exchange Server Agent or Domino Server Agent, as required.
4. In the final wizard screen, click Install to begin the uninstallation.

IIS Restarts When Uninstalling Exchange Server Agent or IIS SMTP Agent

When uninstalling the Exchange server agent or IIS SMTP agent, the wizard stops Internet Information Services (IIS) before uninstalling the server agent and hub components. It then restarts IIS automatically when the uninstall is complete.

Note: IIS is installed automatically as part of an Exchange Server installation.

Chapter 16: Content Registration

The Content Registration screen is where you manage content agents. These agents can detect fingerprinted documents.

This section contains the following topics:

About Content Registration (Fingerprinting)

The Content Registration feature in CA DataMinder, also known as fingerprinting, enables you to take 'fingerprints' of sensitive documents that you want to protect. In effect, you can register the content of these documents so that they can be quickly recognized if a user tries to copy or send them or when CA DataMinder runs a file scan. CA DataMinder can then apply appropriate policy controls.

Fingerprinting is simple to roll out and does not require complex changes to your user policies. Instead of defining complex document classifications in the user policy, you can register the content of the files you want to protect.

Fingerprinting is also the best way to protect documents with highly specialized text content, such as source code, and files with little text content. For example, you can fingerprint CAD drawings, graphics saved in a spreadsheet, and multimedia files.

These fingerprints represent unique document signatures and are made available to CA DataMinder policy engines and endpoint agents. When CA DataMinder analyzes a file, it can quickly determine whether the file matches a known fingerprint and apply policy controls to that file. For example, it can block a fingerprinted document from being sent as an email attachment or copied to a USB device. It can even detect if a document or email contains extracts of text copied from a protected file.

Notes

- The ability to detect text extracts copied from a fingerprinted document is provided by Text Detection content agents.
- Content agents cannot reliably detect spreadsheets (see page 395) or printed files (see page 394).

Fingerprinting Components

Fingerprinting in CA DataMinder relies on content agents, content indexes and specialized policy triggers.

Content agents

Each *content agent* has a list of protected files. These are files stored on your network whose content has already been scanned and fingerprinted. You can have as many content agents as you need. For example, you may have separate agents to fingerprint documents owned by the Finance and HR teams.

Create your content agents before you roll out fingerprinting across your CA DataMinder enterprise. For details about creating content agents, see the *Platform Deployment Guide* or the Administration console online help.

Note: The Content Registration feature uses File Scanning Agent technology to scan files and generate fingerprints.

Content indexes

A *content index* is a list of fingerprints for all the files protected by an individual content agent. You must manually build the index after creating a content agent. You must then publish the index to the CMS to make the content agent available to your policy engines and endpoint agents.

If the list of files protected by a content agent changes, you must rebuild and republish it. If you build an index again, it contains the fingerprints for the original list of protected files plus the fingerprints of any new or modified files. This means that the rebuilt index can contain fingerprints for both new and old versions of the same file. To eliminate multiple versions of the same file from the index, purge and rebuild the index.

If the file list changes are substantial, or if you remove a document from the protected files list, we recommend that you purge the index and then rebuild and republish it.

To build an index, CA DataMinder runs a specialized FSA scanning job.

Content agent triggers

A *content agent trigger* uses content agents to identify protected files. When the trigger analyzes a file (for example, an email attachment), it generates a digital signature, or fingerprint, for that file. The trigger then compares that fingerprint with lists of known fingerprints. If the fingerprints match, a policy trigger fires. You must set up your user policies to use content agent triggers before your fingerprinted files are fully protected.

Requirements

Content Registration has the following requirement:

File Scanning Agent (FSA)

CA DataMinder uses the FSA to build the indexes for content agents.

The FSA must be installed on your network before you can roll out fingerprinting across your CA DataMinder enterprise.

Content Agent Types

When you create a new content agent, you must choose its type. CA DataMinder supports the following types of content agent:

File Detection

These content agents can detect protected files in their entirety.

For example, if a user attaches a protected file, wholly unchanged, to an email or copies it to a USB drive, the content agent detects it.

Text Detection

These content agents can detect emails or files containing text copied from, or based on, a protected document. There are two aspects to these agents.

First, you must specify the agent accuracy. This depends on how much detail is stored in each document fingerprint, which in turn affects the size of the content index associated with the agent.

Second, you must specify how sensitive the agent is when checking suspected emails or documents for protected content.

Accuracy

You can choose how accurately your content agent can identify text extracted from a fingerprinted document.

The available settings (sentence, paragraph, page and so on) determine the level of detail stored in the fingerprint of a protected document. Specifically, these settings determine the size of each analyzable section in a protected document. The agent then generates a fingerprint for that document based on the most significant phrases in each section.

Agent Sensitivity

You can specify how sensitive the agent is to the loss, or potential loss, of protected files and documents. There are two methods for specifying the agent's sensitivity.

- Detect documents or their derivatives

  The agent can check for variants of protected documents, such as an early draft of a sensitive report.

  The agent can also detect documents that closely resemble a protected document. These include documents that have been deliberately modified (for example, by changing key phrases or re-ordering sections) in an attempt to circumvent CA DataMinder policy triggers.

- Detect extracts of documents

  Alternatively, the agent can check for extracts copied word-for- word from a protected document. You must specify the minimum size of these extracts.

At one extreme, you can set up an agent to detect any sentence, or any significant phrase, copied from a protected document and pasted into an email or another file.

At the other extreme, you can set up an agent to only detect emails or files that contain significant, extended passages copied from a protected document, equivalent to a several paragraphs or a full page of text.

Plain Text Embedded File Agents Are Superseded

Note: CA DataMinder 12.5 originally included support for Plain Text Embedded File content agents. These agents have since been superseded by Text Detection content agents.

Plain Text Embedded File content agents could detect the complete body text of a registered plain-text document. For example, the agents could detect a document with sensitive plain-text content (such as source code) if a user embedded the document in another document or in the body of an email.

However, CA DataMinder no longer allows you to create new Plain Text Embedded File agents, although existing agents will continue to work. Instead, we recommend that you replace your Plain Text Embedded File agents with the new Text Detection agents.

Text Detection agents represent a significant improvement over the old Plain Text Embedded File content agents. Previously, users could easily circumvent Plain Text Embedded File content agents by adding a minor text change to a protected document. Now you can set up Text Detection agents to detect, for example, variants of a protected document or attempts by a user to copy an extract from a protected document into an email or an attachment.

Using Content Agents to Detect Printed Files

CA DataMinder content agents can usually detect when a user tries to print a protected file. When this happens, the content agent invokes a Data In Motion trigger to, for example, block the print job. However, there are some limitations depending on the file type:

Text files

These include Microsoft Word and PDF documents.

Content agents **can** detect when a user tries to print extracts from a protected text file. However, an agent can only do this if it is configured to detect emails or files containing at least two-thirds of the text content from a protected document.

In practical terms, this means that the agent's Evaluation properties must use the 'Detect documents or their derivatives' method with the percentage similarity set to 70% or **lower**.

Note: If you choose a percentage similarity higher than 70%, the agent will be unable to match a print job to a protected document because of the presence of extra data such as print headers.

Spreadsheets and tables

These include Microsoft Excel files and information stored in tables in Microsoft Word documents.

Content agents **cannot** detect when a user tries to print a protected spreadsheet. Likewise, they cannot detect when a user tries to print a table copied from a protected file.

Images

These include Microsoft Visio drawings and all other image files, including PNGs, GIFs, and JPEGs.

Content agents **cannot** detect when a user tries to print protected Visio drawings or other images.

Presentations

These include Microsoft PowerPoint files.

Content agents **can** detect when a user tries to print slides from a protected presentation. However, if the slides have a decorative background, such as an image, this can sometimes interfere with the agent's ability to detect the text content. In turn, this may mean that policy triggers do not fire.

Using Content Agents to Detect Spreadsheets

CA DataMinder content agents can usually detect when a user tries to copy or send a protected spreadsheet. For example, you can use content agents to block emails where the sender has attached a protected spreadsheet.

File Detection and Text Detection content agents can both detect protected spreadsheet files if they are unchanged. For example, they can detect when a user attaches the spreadsheet to an email.

Text Detection agents can also detect protected spreadsheets with minor edits (for example, a spreadsheet that contains some cells with updated values). In addition, they can detect when a user copies multiple rows from a protected spreadsheet into a new spreadsheet.

However, Text Detect agents cannot reliably detect when a user copies cell ranges from a protected spreadsheet into a different document format. For example, they may fail to detect when a user copies a range of cells into an email. This is particularly likely if the copied cells include hidden rows or columns because the pasted cell range omits the hidden cells and so no longer matches the spreadsheet fingerprint.

How to Set Up Content Agents

A registered content agent can quickly detect when a user tries to send or copy a protected file. To do this, it compares the file's digital fingerprint with the fingerprints of files that it is protecting.

To roll out content agents across your CA DataMinder enterprise

1. Create your content agents. Do this in the Administration console.
2. For each agent, you must specify which files it protects and its index type.

   For Text Detection agents, specify the agent accuracy and detection thresholds.

   Note: If you later change or extend the files protected by a specific content agent, you must rebuild the index and republish the content agent.
3. Build an index for each content agent.

   The index contains fingerprints of the files you want the agent to protect.
4. Publish the content agent.

   This process pushes the content index file onto the CMS and makes a fully functioning content agent available to your policy engines and endpoint agents.

   Note: Until an index has been built and the content agent has been published, you cannot use the agent in user policy. Any trigger that uses an unpublished content agent will be unable to detect fingerprinted files.
5. Assign content agents to triggers in your user policies.

More information:

Create a Content Agent (see page 396)
Specify the Protected Files (see page 399)
Build a Content Index (see page 400)
Publish a Content Agent (see page 402)
Set Up Content Agent Triggers (see page 402)

Create a Content Agent

You create content agents in the Administration console.

Follow these steps:

1. In the Administration console, expand the Content Registration branch.
2. Right-click the Agents folder and click Create Content Agent.

   The Select Agent Index Type dialog appears.

3. Select the agent type:

   File Detection

   These content agents can detect protected files in their entirety. For example, if a user copies a protected file, wholly unchanged, to a USB drive, the content agent detects it.

   Text Detection

   These content agents can detect emails or files containing text copied from, or based on, a protected document. These agents have two key properties.

   First, you can specify how much detail is stored in each document fingerprint. The level of detail affects the size of the content index associated with the agent.

   Second, you can specify how sensitive the agent is when searching suspected emails or documents for protected content.

   (Optional) You can designate a Text Detection agent as one created explicitly to reduce the number of false positives when used in conjunction with an ordinary content agent. For details, see Supplementary Content Agents to Reduce False Positives (see page 399).

4. Click OK.

   The Agent Properties dialog appears.

5. Supply the following agent details:

   General tab

   Specify the agent name and description and the Index Builder server. This server can be any server on your network hosting the CA DataMinder File Scanning Agent.

   (Optional. Text Detection agents only) Specify supplementary agents to reduce the number of false positives. These agents identify and exclude from processing text extracts that are deemed benign or acceptable, such as corporate disclaimers.

   Filter tab

   (Optional) Specify the folder and file options, such as whether to fingerprint hidden files or files in subfolders.

   Build tab

   (Text Detection agents only) Specify how much detail is stored in the fingerprint of a protected document. The more detailed the fingerprint, the more reliably an agent can recognize content originating from that document. However, a content index containing highly detailed fingerprints can be extremely large.

   For details, see What Level of Accuracy Do I Need? (see page 398)

Evaluation tab

(Text Detection agents only) Specify how sensitive the agent is to the loss, or potential loss, of protected files and documents. You can use two methods to specify the agent sensitivity:

- The agent can search for variants of protected documents or documents that closely resemble a protected document.
- Alternatively, the agent can search for extracts copied word for word from a protected document. You can specify the minimum size of these extracts.

6. Click OK.

   The new content agent is listed in Agents folder of the Content Registration branch.

More information:

What Level of Accuracy Do I Need? (see page 398)
Supplementary Agents to Reduce False Positives (see page 399)

What Level of Accuracy Do I Need?

(Applies to Text Detection content agents only)

The accuracy of an agent depends on the level of detail in the document fingerpints. If you include more detail in a document's fingerprint, the agent is more accurate. That is, the agent can recognize content originating from that document more reliably.

At one extreme, you can generate a fingerprint that covers every **sentence** in a protected document. Agents based on highly detailed document fingerprints like these are the most accurate.

At the other extreme, you can generate a fingerprint that slices the document into sections analgous to a **page**. This fingerprint only covers the most significant phrases within each page-sized section. Agents based on these fingerprints are less accurate, but have far smaller indexes.

Note: A content index containing highly detailed (sentence-level) fingerprints can be very large indeed. In fact, these indexes can be *eight* times larger than comparable indexes containing the least detailed (page-level) fingerprints. Such large indexes can be a particular problem if you enable content agent triggers on your CA DataMinder endpoint machines, because each endpoint machine maintains a copy of each agent's content index.

Supplementary Agents to Reduce False Positives

(Applies to Text Detection content agents only)

A subsidiary content agent operates in conjunction with a parent agent to reduce the number of false positives detected by the parent content agent. The subsidiary agent identifies and excludes from processing any text that is deemed benign or acceptable, such as corporate disclaimers.

For example, if you have fingerprinted a series of confidential reports, each of which includes a corporate disclaimer, you do not want your content agent to positively flag every email or file containing this disclaimer. You can prevent this happening by specifying a subsidiary agent that has been configured to detect these disclaimers. When the subsidiary agent detects a disclaimer, the associated text is excluded from normal processing by the parent content agent.

You can designate one or more supplementary content agents when you create a Text Detection content agent. You can then create an index of digital fingerprints for each supplementary agent in the normal way. For example, you may want to create a text document containing your corporate disclaimer and generate a fingerprint of this document.

Specify the Protected Files

After you create a content agent, identify the files and folders that you want the content agent to protect.

Typically, a system administrator works in association with a security auditor to identify sensitive files that need fingerprinting. In this context, a security auditor is someone within your organization with responsibility for preventing data loss.

Note: If you later change or extend the files protected by a specific content agent, you must rebuild the index and republish the content agent.

Follow these steps:

1. In the Administration console, expand the Content Registration branch.
2. Right-click the agent and click Add Location.

   The Add Location dialog appears.
3. Click Add Files or Add Folders to select the items you want to protect.

   Note: CA DataMinder can fingerprint files contained within zipped files.
4. Click OK.

   The files and folders that you selected are listed in the Protected Files and Folders pane.

Build a Content Index

Before you can use a content agent, you must create a content index. A content index identifies the files protected by a content agent. Each content agent has its own content index.

The index contains digital fingerprints of the files that you want the content agent to protect. These fingerprints represent unique document signatures and are made available to CA DataMinder policy engines and endpoint agents. CA DataMinder generates the index by running a scanning job to collate the digital fingerprints.

In technical terms, CA DataMinder generates the index by running a specialized FSA scanning job to collate the digital fingerprints. It then creates the index file on the FSA host server. When you publish the content agent (see the next section), the index file is replicated to CA DataMinder policy engines and endpoint agents.

Follow these steps:

1. In the Administration console, expand the Content Registration branch.
2. Select the content agent and click the Start Build button.

   A scanning job starts. To cancel the scanning job, click the Stop Build button.

   When the scanning job completes, the Status field in the Content Registration Agent pane updates the number of available index files.

Rebuild a Content Index

If you change or extend the files protected by a specific content agent, you must rebuild the index and republish the agent.

If you build an index again, it contains the fingerprints for the original list of protected files plus the fingerprints of any new or modified files. This means that the rebuilt index can contain fingerprints for both new and old versions of the same file. To eliminate multiple versions of the same file from the index, purge and rebuild the index.

The rebuild procedure is the same as when you build an index for the first time.

Follow these steps:

1. (Optional) Purge the content index.

   If you have made substantial changes to the files protected by a content agent, we recommend that you first purge the index to remove all digital fingerprints and then rebuild it.

2. Select the agent in the Content Registration screen.
3. Click the ▶ Start Build button.

   When the build process is complete, the index contains the original list of protected files plus any new or modified files.

4. Re-publish the updated content agent.

To purge a content index

If you have made substantial changes to the files protected by a content agent, we recommend that you purge the index to remove all digital fingerprints and then rebuild it.

Follow these steps:

1. Select the agent in the Content Registration screen.
2. Click the 🗑 Purge button.

   You now need to rebuild the index.

Publish a Content Agent

After you have built the content index, you must publish the content agent. This process moves the content index from the scanned computer onto the CMS. This process also makes the content agent available to your CA DataMinder policy engines and endpoint agents.

Note: Until an index has been built and the content agent has been published, you cannot assign the agent to triggers in your user policies. If a trigger uses an *unpublished* content agent, the trigger is unable to detect fingerprinted files.

Follow these steps:

1. In the Administration console, expand the Content Registration branch.
2. Select the content agent and click the Publish button.

   When the publishing job completes, the Status field in the Content Registration Agent pane displays 'Published'.

Set Up Content Agent Triggers

After you have published a content agent, you assign the agent to triggers in your user policies.

These triggers fire when CA DataMinder detects a file that matches a fingerprint in the content index associated with your content agent.

For example, you can set up triggers to block a protected document from being sent as an email attachment or copied to a USB device.

In technical terms, the content agent generates a fingerprint of the file and then compares this fingerprint with those in its content index. If the fingerprints match, CA DataMinder infers that the file is protected.

Follow these steps:

1. Log on to the Administration console using an account that has the 'Policies: Edit policy' administrative privilege.
2. Expand the User Administration branch.
3. Right-click a user or group and click Edit Policy.

   The User Policy Editor screen appears.
4. Select the trigger that you want to assign a content agent to.

   For email triggers, you can only assign content agents to a Content Agent trigger.

   For Data In Motion and Data At Rest triggers, you can assign content agents to any trigger.

5. Edit the Trigger Name and Policy Class settings as required.

   Reviewers can search for 'protected document events' in the iConsole or Data Management console, using these trigger names and policy classes as search filters.

6. Edit the Use Content Agents For Files? trigger setting.

   Set the value to 'Use content agents to analyze text content'.

7. Edit the Which Content Agents? trigger setting.

   Add the content agents that you want to associate with this trigger.

8. Edit other trigger settings as required.

   For example, you can assign smart tags or a severity rating to the trigger.

9. Edit the Control Action setting.

   Select the control action that you want. For example, to block protected files from being copied or sent, select a 'Block' control action. If you want users to encrypt protected files before copying or sending them, select an 'Enforce Encryption' control action.

10. Repeat steps 3 through 7 for any other triggers that you want to assign the content agent to.

11. Click File, Save to save the policy changes.

    The Policy Change Summary dialog appears.

12. Click Save.

    Your content agent is now assigned to a policy trigger. The user policy is now configured to detect any documents protected by the content agent.

Chapter 17: Content Services

The current CA DataMinder release uses content database technology provided by Autonomy IDOL. You can install the IDOL content database included with CA DataMinder or you can use an existing IDOL content database.

Overview

CA DataMinder Content Services enable reviewers to run content searches. These searches use pattern-matching technology to retrieve events based on their text content. Content searches are available in the iConsole and Data Management console.

Content searches offer several advantages over standard CA DataMinder searches. You can specify the usual search criteria (event type, user or group, when the event was captured, and so on). But you can also search for specific text content and sort the search results by relevance. You can include logical operators in your search expression to zero in on key documents and eliminate irrelevant results. You can even search for documents with common themes or concepts.

To enable content searches, you must first index CA DataMinder events into a content database using the content indexer. This utility extracts CA DataMinder events from the CMS and submits them to the content database. The content database then processes these events, storing them as indexed, text-searchable documents. When a reviewer runs a content search, the content database analyzes the indexed documents and returns all documents that match the search criteria.

IDOL Content Indexing Not Intended To Support High Capture Rate

Content indexing in the current version of CA DataMinder uses Autonomy IDOL technology.

Content indexing based on Autonomy IDOL technology enhances the violation review capabilities for customers who need to:

- Capture a sub-set of an organization's overall information content
- Retain this information content for relatively short periods (for example, one year).

Content indexing using Autonomy IDOL technology is not intended to support high capture rates that would result in substantial volumes of data in the CMS database with long retention periods. CA DataMinder deployments that require Autonomy IDOL installations of significant scale are best served by leveraging the CA DataMinder Content Connector API. Using this API, third party content indexing solutions (including a separately licensed Autonomy IDOL) can consume CA DataMinder data directly and can be scaled independently to meet such requirements.

Content Services Components

CA DataMinder Content Services comprises the following components.

Content Database

The content database is the engine or system that contains indexed versions of events captured by CA DataMinder agents. Events are stored in the content database as text-searchable documents.

The current CA DataMinder release uses content database technology provided by Autonomy IDOL. You can install the IDOL content database included with CA DataMinder or you can use an existing IDOL content database.

Note: Although this guide refers to a single IDOL content database, under the hood this single database actually comprises multiple databases. These multiple databases are similar to conventional database partitions, splitting the overall content database into sections based on indexing timestamps.

CA DataMinder IDOL Connector

The CA DataMinder IDOL Connector enables the CA DataMinder content indexer and content proxy server to communicate with an IDOL content database.

All communication between the CA DataMinder components and the IDOL content database passes through the IDOL connector framework.

IDOL Connector Framework

The IDOL connector framework extracts the text content from files and emails submitted by the CA DataMinder content indexer and forwards the text content to the content database.

When you install the CA DataMinder IDOL Connector, you can also install the IDOL connector framework included with CA DataMinder or you can use an existing IDOL connector framework.

Content Indexer

The content indexer submits events from the CMS to a content database, where the events are stored as indexed documents. The content indexer comprises a service and console:

- The content indexer service monitors your CMS for newly captured data and indexes specified data into your content database.
- The content indexer console lets you run indexing jobs. For usage instructions, see the content indexer console online help.

Content Proxy Server

The content proxy server enables the iConsole or Data Management console to access a content database when running a content search.

Each content proxy server is tied to a single parent CMS. Multiple proxy servers can share a parent CMS.

Content Searches

Content searches use intelligent pattern-matching technology to analyze the text content of indexed events in a content database and retrieve events that match the search criteria.

Reviewers run content searches from the iConsole or Data Management console.

Note: Content searches only support the Management Group security model. If a reviewer has a different security model and runs a content search, the search returns zero events.

Content Services Requirements

Note the following requirements for CA DataMinder Content Services.

IDOL Content Database and Connector Framework

The IDOL content database and IDOL connector framework that ship with CA DataMinder are available in the Content_IDOL.msi installation package.

Content_IDOL.msi supports 64-bit versions of these operating systems:

- Windows Server 2008 R2
- Windows Server 2012

Important! If you install the IDOL connector framework included with CA DataMinder, you must install it on the same server as the content indexer and content proxy server.

Content Indexer and Content Proxy Server

You can install these components on any CA DataMinder server.

These components are included in the server.msi and server_x64.msi installation packages.

Server.msi supports 32-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008

Server_x64.msi supports 64-bit versions of these operating systems:

- Windows Server 2003 (see note 1)
- Windows Server 2008 (see note 1)
- Windows Server 2008 R2
- Windows Server 2012

Note 1: We have not tested these operating systems with the current versions of server.msi and server_x64.msi.

Deployment Architecture

The following diagram shows the deployment architecture for CA DataMinder Content Services. For the current release, CA DataMinder uses content database technology provided by Autonomy IDOL.

A single CA DataMinder content indexer handles data *ingestion* into a content database. The CA DataMinder content proxy server handles data *retrieval* from the content database when a content search runs. In both cases, the connection between CA DataMinder and the content database is handled by the IDOL connector framework.

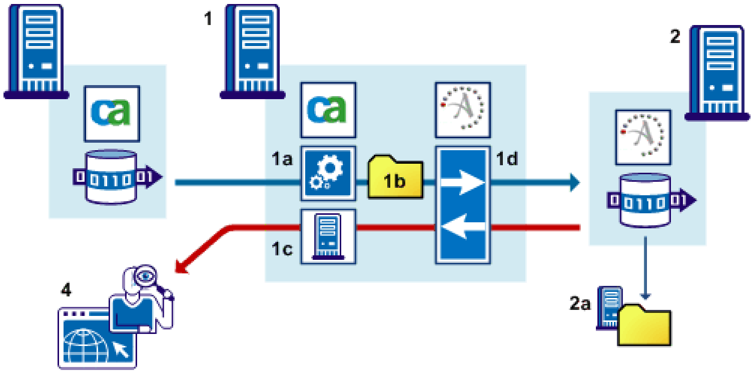

The Content Services server (**1**) hosts the CA DataMinder content indexer (**1a**), the transfer folder (**1b**), and the CA DataMinder content proxy server (**1c**). It also hosts the IDOL connector framework (**1d**).

The content database (**2**) normally runs on a dedicated server. Indexed event data is stored in a data folder (**2a**), which can be local or remote.

The content indexer extracts CA DataMinder emails and file events from the CMS database (**3**). These events are temporarily held in the transfer folder before being picked up by the connector framework. The connector framework submits these CA DataMinder events to the content database where their text content is analyzed and indexed.

When a CA DataMinder reviewer (**4**) runs a content search in the iConsole or Data Management console, the search query is channeled through the content proxy server and the connector framework to the content database. The search results are returned to CA DataMinder using the same mechanism.

How to Deploy Content Services

To deploy CA DataMinder Content Services, follow these steps:

1. Install a new IDOL content database. You can install the content database on any 64-bit Windows server.

   Alternatively, you can use an existing IDOL content database.

2. Install the CA DataMinder IDOL Connector. This feature connects CA DataMinder to the IDOL connector framework.

   You must install the connector framework on the same server as the content indexer and content proxy server. Alternatively, you can use an existing IDOL connector framework.

3. Install the following CA DataMinder components:

   - Content indexer server and console.
   - Content proxy server.

   You can install these components on any CA DataMinder server.

4. Start indexing events into the content database using the content indexer console.

   See the online help for usage instructions,

5. Install the content search onto your iConsole front-end servers.

   - See the iConsole Standard Searches chapter of the *Platform Deployment Guide* for installation details.
   - See the iConsole online help for usage instructions.

   If content searches are explicitly included in your license agreement, they are available automatically in the Data Management console. See the online help for usage instructions.

Installing the Connector Framework and Content Database

You install the IDOL connector framework and content database using the CA DataMinder IDOL Connector installer.

You must install the connector framework on the same server as the content indexer and content proxy server.

In a typical deployment, you install the content database on a separate, dedicated server. However, you can install the content database on the same server as the connector framework if preferred.

Install the Content Database Only

Using the CA DataMinder IDOL Connector installer, you can only install a content database onto the local server. You must therefore run the installer on the server that you want to host the content database.

To install the content database only

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose Content Services and then click Install.

   This launches the CA DataMinder Content Services installation wizard in a separate window.

4. In the content services server installation wizard, navigate to the Custom Setup screen.

5. Click the Autonomy IDOL Content Database feature only. Then click Next.

6. In the Content Database Configuration page, specify the following items:

   Service Port

   > Specify the port number that the content database listens on.

   Data Folder

   > Specify the folder used to hold the content database files. These files contain the indexed text content of CA DataMinder events. You must ensure that the target server has sufficient free disk space to store your indexed events.
   >
   > In a typical deployment, this folder is on the same server as the content database. But you can specify a remote location. If the folder is on a remote computer, specify a network file share using the universal naming convention (UNC) or mapped network drive.

   Service Account

   > (Optional) If you specified a remote folder location, specify a domain account with write access to the remote folder. The content database uses this domain account to store and retrieve indexed events.

7. Continue to the final wizard screen and click Install.

8. (Applicable only if you plan to install the content indexer on a different server to the content database) Manually specify the content indexer host server.

   See the reference below for details.

More information:

Manually Specify the Content Indexer Host Server (see page 417)

Install the Connector Framework Only

Using the CA DataMinder IDOL Connector installer, you can only install the connector framework onto the local server. You must therefore run the installer on the server that you want to host the content indexer and content proxy server.

To install the IDOL connector framework only

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose Content Services and then click Install.

   This launches the CA DataMinder Content Services installation wizard in a separate window.

4. In the content services server installation wizard, navigate to the Custom Setup screen.

5. Click the CA DataMinder IDOL Connector feature only. Then click Next.
6. In the Autonomy IDOL Content Framework Selection page, specify the following items and then click Next.

   Install New Connector Framework

   Click this option if you want to install a new IDOL connector framework service.

   Use Existing Connector Framework

   Click this option if you want to use an existing IDOL connector framework service or if you want to use IDOL distributed action handlers (DAHs) to process indexed data.

   Server

   Enter the fully qualified domain name (FQDN) or IP address of the server hosting the connector framework or DAH.

   Service Port

   Specify the port number that the remote connector framework service or DAH listens on.

   Transfer Folder

   Specify the folder used to store files and emails waiting to be indexed.

   This folder must be accessible to the content framework service or DAH. If you specified a remote content framework service or DAH, specify a network file share using the universal naming convention (UNC) or a mapped network drive.

7. In the Content Database Selection page, identify the server and service port for your existing content database.

   Host Server

   Enter the fully qualified domain name (FQDN) or IP address of the server hosting your content database.

   Service Port

   Specify the port number that the content database listens on.

8. Continue to the final wizard screen and click Install.

Install the Connector Framework and Content Database Together

If required, you can install the IDOL content database and IDOL connector framework on the same server. Use the CA DataMinder IDOL Connector installer to install these components.

To install the connector framework and content database together

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.

3. In the Advanced Install Options screen, choose Content Services and then click Install.

   This launches the CA DataMinder Content Services installation wizard in a separate window.

4. In the content services server installation wizard, navigate to the Custom Setup screen.

5. Click the Autonomy IDOL Content Database feature *and* the CA DataMinder IDOL Connector feature. Then click Next.

6. In the Autonomy IDOL Content Framework Selection page, specify the following items and then click Next.

   Install New Connector Framework

   Click this option if you want to install a new IDOL connector framework service.

   Use Existing Connector Framework

   Click this option if you want to use an existing IDOL connector framework service or if you want to use IDOL distributed action handlers (DAHs) to process indexed data.

   Server

   Enter the fully qualified domain name (FQDN) or IP address of the server hosting the connector framework or DAH.

   Service Port

   Specify the port number that the remote connector framework service or DAH listens on.

   Transfer Folder

   Specify the folder used to store files and emails waiting to be indexed.

   This folder must be accessible to the content framework service or DAH. If you specified a remote content framework service or DAH, specify a network file share using the universal naming convention (UNC) or a mapped network drive.

7. In the Content Database Configuration page, specify the following items:

   Service Port

   Specify the port number that the content database listens on.

   Data Folder

   Specify the folder used to hold the content database files. These files contain the indexed text content of CA DataMinder events. You must ensure that the target server has sufficient free disk space to store your indexed events.

   In a typical deployment, this folder is on the same server as the content database. But you can specify a remote location. If the folder is on a remote computer, specify a network file share using the universal naming convention (UNC) or mapped network drive.

   Service Account

   (Optional) If you specified a remote folder location, specify a domain account with write access to the remote folder. The content database uses this domain account to store and retrieve indexed events.

8. Continue to the final wizard screen and click Install.

Installing the Content Indexer and Proxy Server

You must install the content indexer and content proxy server on the same computer as the IDOL connector framework. You install these components using the CA DataMinder server installation wizard.

This section provides two sets of installation instructions. Follow the first set of instructions if you are installing these components on a 'new' computer (not already running the CA DataMinder infrastructure). Follow the second set of instructions if you want to install these components on an existing CA DataMinder server.

Note: You can only connect one content indexer to your content database. Do not deploy multiple content indexers.

To install the content indexer and content proxy server on a 'new' CA DataMinder computer

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.
3. In the Advanced Install Options screen, choose CA DataMinder Platform and then click Install.

   This launches the CA DataMinder server installation wizard in a separate window.

4. In the server installation wizard, navigate to the Custom Setup screen.
5. Click the Content Services feature and the following subfeatures:

   Content Indexer Console

   The console enables you to specify which events to index from the CMS into a content database.

   Content Indexer Server

   This subfeature installs the indexing service. Content indexing jobs use this service to index events from the CMS into a content database.

   Content Proxy Server

   This subfeature installs the proxy server service.The CMS uses this service to connect to the content database when a reviewer runs a content search.

6. In the Server Type screen, click Utility Machine.
7. In the Connectivity screen, enter the name or IP address of the parent server. This can be the CMS or a gateway server.
8. In the Data Location screen, specify the name and network location of the data folder.

   This folder contains configuration files for indexing jobs, .evl files for 'index failures' waiting to be resubmitted, and log files. You can accept the default location or specify a different location.
9. In the Database Type, select the database engine to use on the utility machine.

   For these content services components, we recommend that you use the default engine, Microsoft Jet.

   If you prefer to use a SQL Server or Oracle database, you must provide further details; see 'Installing a Utility Machine (see page 52)'.
10. In the Service Accounts screen, specify the logon account for the local CA DataMinder infrastructure service. This default acount is LocalSystem.

    Important! If you chose to specify a remote data folder in step 8, you must change the credentials of the infrastructure account.
11. Continue to the final wizard screen and click Install.
12. (Applicable only if the content indexer is installed on a different server to the content database) Manually specify the content indexer host server.

    See the next section for details.

To install the content indexer and content proxy server on an existing CA DataMinder server

If you want to install these components on an existing CA DataMinder CMS or gateway, you must modify the existing installation.

1. In the Add or Remove Programs applet, click CA DataMinder Server.
2. Click the Change button.

   The CA DataMinder installation wizard starts.
3. Go to the Program Maintenance screen and click Modify.

   The Custom Setup screen appears.
4. Click the Content Services feature and the following subfeatures:

   - Content Indexer Console
   - Content Indexer Server
   - Content Proxy Server

   See the previous section for descriptions of these features.

5. Continue to the final wizard screen and click Install.
6. (Applicable only if the content indexer is installed on a different server to the content database) Manually specify the content indexer host server.

   See the next section for details.

More information:

Manually Specify the Content Indexer Host Server (see page 417)

Manually Specify the Content Indexer Host Server

If the content indexer is installed on a different server to the content database, you must manually configure the content database to identify the content indexer host server.

Follow these steps:

1. On the content database host server, stop the CA DataMinder Content Database service.
2. Edit the content.cfg file.

   Find this file on the server hosting the content database. The file is in the \Content\IDOL folder below the CA DataMinder installation folder.
3. In the [IndexNotify] section, change the Host value 'localhost' to the name or IP address of the content indexer host server.
4. Save the changes to content.cfg file.
5. Restart the CA DataMinder Content Database service.

Optional Post-Install Configuration

(Optional) To control content search concurrency, edit the following registry key.

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
    \CurrentVersion\Content Services
```

Within this registry key, edit the following registry value:

SearchConcurrency

> **Type:** REG_DWORD
>
> Specifies the number of search requests that can execute concurrently. You can override this setting without having to restart the service.
>
> **Default:** 10

Configuring the IDOL Content Database

The following section describes an optional configuration for the IDOL content database.

More information:

Disable Search Result Estimates (see page 418)
Extend the Query Timeout (see page 419)

Disable Search Result Estimates

When you run a content search in the iConsole, the results screen displays the total number of matching events,

By default, this total is an *estimate*, calculated using an Autonomy IDOL cumulative predictive algorithm. IDOL uses this algorithm to optimize search performance and minimize search times. But if necessary, you can disable this predictive algorithm so that the iConsole results screen displays the *actual* number of matching events.

To disable estimated totals of matching events, you edit the content database configuration file and extend the query timeout. However, if you do disable estimated totals, content searches may take longer to complete.

Note: Do not confuse the total number of matching events with the total number of search results. By default, iConsole searches are capped. This means, for example, that a search may only return 100 results even if there are more events than this in the CMS database that match the search criteria. If necessary, you can override the search capping to display all matching results. See the iConsole online help for details.

To disable estimates of total matching events

1. Stop any content indexing jobs that are currently running.
2. Stop the CA DataMinder Content Database service.

   While this service is stopped, you cannot run content searches.

   This service runs on the server hosting the content database.
3. Edit the content.cfg configuration file.

   Find this file on the server hosting the content database. The file is in the \Content\IDOL folder below the CA DataMinder installation folder.
4. Add the following line to the [Server] section of content.cfg.

   ```
   TotalResultsPredictionThreshold=0
   ```
5. Restart the CA DataMinder Content Database service.
6. Extend the query timeout (see page 419) for content searches.

Extend the Query Timeout

By default, the query timeout for content searches defaults to 30 seconds. In some situations, you may need to extend this timeout.

For example, content searches take longer to return results if you disable estimated results totals. Also, if you use the idolMgr utility to list content databases by age, the query may timeout if you have a large number of databases in your IDOL content database.

In these situations, we recommend that you extend the query timeout.

To extend the query timeout

You must edit the registry on the server hosting the IDOL connector framework and content database.

1. Go to the following registry key:

   On 64-bit systems

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA
   DataMinder\CurrentVersion
       \Content Services\IDOL Provider
   ```

   On 32-bit systems

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder\CurrentVersion
       \Content Services\IDOL Provider
   ```

2. Increase the QueryTimeoutSecs registry value.

   Because content databases vary in size, we cannot give specific recommendations for the new timeout value. Investigate different timeouts until you discover the optimal timeout for your content database.

Retrying Failed Indexing Jobs

The content indexer can generate Event Link Files (EVLs) to help you deal with events that the indexer fails to process into the content database.

- Click the Retry Failed Events option button on the Job settings Summary tab to configure a retry job.
- Click Generate EVL Files on Failure to activate the generation of EVL files. This corresponds to setting the OUTPUTFLAGS=1 parameter in the OPT file.

More information:

OUTPUTFLAGS—Generate EVL Files (see page 428)

How Failed Batch Jobs Are Categorized

Set the OUTPUTFLAGS=1 parameter to configure standard jobs to generate EVLs. If a standard job fails to fully index events, the events are categorized into subfolders of the EVLFAILDEST directory.

1. The subfolders are grouped by their eventtimestamp. This way you can reference and process events of similar time periods more efficiently.

   Use the OPT parameter EVLSPLIT to control the granularity of the grouping.

2. The next subfolder is named either ‘Retry’ or ‘No retry’. The folder names indicate whether the failed event is likely to succeed on subsequent submissions.
3. If the event failed for a common reason, the EVL is stored in an appropriately labeled subfolder, for example, "Document Too Large" or "Provider Failure".

More information:

OUTPUTFLAGS—Generate EVL Files (see page 428)
EVLFAILDEST—Specify the EVL Directory (see page 428)
EVLSPLIT—Sort EVL Folders by Eventimestamp (see page 428)

How Failed Retry Jobs Are Categorized

Set the OUTPUTFLAGS=1 parameter to configure batch jobs to generate EVLs. If a retry job fails to fully index events, the events are written to subfolders of the EVLFAILDEST directory.

- By default, retry jobs recurse through subfolders of EVLRETRYSRC and reprocess previously failed events into the content database.

  Use the parameter EVLRECURSE to control whether to drill-down into dubfolders.

- Retry jobs do not process any EVLs under the ‘no retry’ folder.

More information:

OUTPUTFLAGS—Generate EVL Files (see page 428)
EVLFAILDEST—Specify the EVL Directory (see page 428)
EVLRECURSE—Retry Failed Subfolders (see page 429)

How Retries Are Tracked

Due to unforeseen circumstances, the content indexer may not process some events. The EVLs of these failed events are not placed in the ‘no retry’ folder because they are scheduled to be retried. If the event fails permanently, you want to prevent retry jobs from processing the same failed event repeatedly.

The content indexer prevents this endless loop by keeping a record of how often it retried to process an EVL. Set the OPT parameter EVLRETRIES to control how often jobs are reprocessed.

1. Each time a retry job reprocesses an EVL, it increments a counter (‘Retries=’) which is stored in the EVL.
2. If the job retries an EVL, and it fails, and the retry counter exceeds the EVLRETRIES value, then the EVL is written to the ‘No retry’ folder under the subfolder ‘Struck out’. When the EVL is written to the ‘struck out’ folder, the failure folder structure persists.
3. Retry jobs do not target the ‘No retry’ folder and the failing job is not reprocessed.

EVL File Format

EVL files are labeled <Servername>-<Eventtimestamp>-<eventuid>.evl. For example, 'MONITRAX-20070718-1125053.evl'.

Inside the EVL file, a retry value stores the number of times the EVL has been processed by a retry job. The retry jobs update this counter everytime the retry an EVL.

```
[Content Indexer]
Retries=0x1
```

Folder Structure

The top level folder is the destination specified in the .opt file. The EVL folder structure is created as follows:

EVL creation date sub folder (year/month)

- No retry folder
  - Struck-out folder
    - Batch failed folder
    - RTS Output Send Failed folder
    - Memory Exception folder

        - External Timeout folder
        - Other

          ‘E105A’ folder, etc.
    - Unknown file format folder
    - Cannot Access Content folder
    - Process Read Error folder
- Retry folder
    - Batch failed folder
    - RTS output send failed folder
    - Memory Exception folder
    - External Timeout folder
    - Other

      'E105A’ folder, etc.

Common Event Failures

These tables contain exammples of common event failures and their failed folder mappings.

Non-Retryable Event Failure	Folder
Unsupported mime-type format	Unknown file format
File is password protected or encrypted	Cannot Access Content
File is empty or corrupt	Process Read Error
XML Unsupported encoding	Process Read Error

Retryable Event Failure	Folder
FAST Memory Exceptions	Memory Exception
Document cannot be sent to RTS	RTS Output Send Failed
FAST External Process timeout	External Timeout
Batch submission failure	Batch failed

Purging the IDOL Content Database

Although this guide refers to a single IDOL content database, under the hood this single database actually comprises multiple databases. These multiple databases are similar to conventional database partitions, splitting the overall content database into sections based on indexing timestamps.

By default, CA DataMinder instructs IDOL to create a new database every seven days, although you can change this interval. When new documents are indexed, they are always stored in the latest database. Older databases remain in the IDOL content database and are included in content searches until you purge them.

Splitting the IDOL content database into multiple databases provides various benefits. These benefits include faster content searches and easier database maintenance.

Change the New Database Interval

You can change the interval at which new databases are created in the IDOL content database.

You can change this interval even if an indexing job is currently running. The change becomes effective at midnight.

To Change the New Database Interval

1. Edit the registry on the server hosting the IDOL connector framework and content database.
2. Go to the following registry key:
   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
       \CurrentVersion\Content Services\IDOL Provider
   ```
3. Modify the NewDatabaseIntervalDays registry value.

   The default data is 7, meaning that IDOL starts a new database every 7 days. Enter a new interval.

   For example, set NewDatabaseIntervalDays to 30 if you only index a small amount of data each month.

Restore a Backed Up Database

Before you purge an individual database, we recommend that you back up the entire content database. This section describes how to restore the backup.

To restore a database backup

In a browser, type this command to restore the entire content database:

```
http://<IDOL_server>:<index_port>/DREINITIAL?<backup_folder>
```

Where:

<IDOL_server>

Identifies the server hosting your IDOL content database.

<index_port>

Specifies the index port for the content database. The default port number is 9001.

<backup_folder>

Specifies the path to the folder containing the database backup.

Example

Run this command on the content database host server to restore the content database that you previously backed up:

```
http://localhost:9001/DREINITIAL?C:\MyBackup
```

More information:

Extend the Query Timeout (see page 419)

Purge Databases by Age

Older content databases remain in the IDOL content database until you purge them (that is, delete them). Use the IDOL database management utility, IdolMgr, to purge content databases based on their age.

For example, you can schedule regular IdolMgr jobs to purge databases that are more than six months old. This allows you to maintain a rolling six-month history of indexed documents.

You typically schedule the IdolMgr to run on the CA DataMinder content indexer server or content proxy server. You run IdolMgr from a command line. The command syntax is shown below.

Syntax

```
IdolMgr [-databases[=days|date]] [-purge=days|date]
        [-commitchanges] [-noprompt] [-verbose]
        [-server=server[:port]] [-license=string] [-help]
```

-databases[=*n*|*date*]

Lists the IDOL databases. If you specify the optional parameter, IdolMgr only lists databases that are older than *n* days or older than the specified *date*.

Format: A time period of *n* days, or a date in the format *dd/mm/yy*.

Example: To list only databases containing documents more than 7 days old:

```
IdolMgr -databases=7
```

Note: If your IDOL content database is very large, IdolMgr may time out before it can finish listing the databases. If this happens, you must extend the query timeout (see page 419).

-purge=*n*|*date*

Deletes IDOL databases that are older than *n* days or older than the given date and which contain no documents more recent than that. IdolMgr prompts you for confirmation before it deletes a database. Use this command with the -commitchanges option.

Example: Use the following command to view the list of databases that are older than 60 days and contain no documents more recent than 60 days, without actually deleting any databases yet:

```
IdolMgr -purge=60
```

Format: A time period of *n* days, or a date in the format *dd/mm/yy*.

-commitchanges

(Optional) Commit IDOL database changes.

Important: If you do *not* specify this parameter, IdolMgr only executes searches and other actions and reports the results, but it does not actually commit the changes to the IDOL server databases.

Example: Use the following command to delete databases on the IDOL server that are older than 60 days and contain no documents more recent than 60 days:

```
IdolMgr -purge=60 -commitchanges
```

-noprompt

Executes the command without prompting you for confirmation. The command still reports the results of the action.

Example: This command deletes all databases that are older than 60 days without prompting for confirmation:

```
IdolMgr -purge=60 -commitchanges -noprompt
```

-verbose

Displays additional progress and status information for the command. The output includes the commands sent by IdolMgr to the IDOL server plus extended error information.

Example: This command displays additional progress information while deleting all databases that are older than 60 days:

```
IdolMgr -verbose -purge=60 -commitchanges
```

-server=*name[:port]*

(Optional) Specifies the name of the IDOL server and, optionally, the service port number. If you do not specify the server details, IdolMgr attempts to use the default values from the registry. If the defaults are not specified either, IdolMgr does not continue.

Examples:

```
IdolMgr -server=localhost
IdolMgr -server=localhost:9002
```

Default: The data set for the IDOLServerName and IDOLServerServicePort registry values. Find these values in the follow registry key:

On 64-bit systems

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA
DataMinder\CurrentVersion
    \Content Services\IDOL Provider
```

On 32-bit systems

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder\CurrentVersion
    \Content Services\IDOL Provider
```

-license=*string*

Specifies the IDOL server license string to connect to the IDOL server. If you do not specify the license string, IdolMgr attempts to read it from the registry value OverrideKey (see the previous parameter for registry key details).

If you do not specify the license string parameter or the OverrideKey registry value, IdolMgr uses the default CA OEM license string.

Example:

```
IdolMgr -databases=7 -license=UHFyiffLKtljKe7JuM3DA74EAI3zMPsN/khhg2i
```

Default: CA OEM license string.

Example

IdolMgr runs as a scheduled task once a month to delete databases that are older than 6 months (180 days). To do this, schedule the following command on the content indexer server or content proxy server:

```
IdolMgr –purge:180 –commitchanges –noprompt
```

Note: IdolMgr contains several other command options to allow you to examine your databases and decide what management regime is best for you. We recommend that you try out these other commands before committing to a purge regime.

OPT File Parameters

The OPT file controls settings for the Content Indexer. You find the OPT file in ProgramData\CA\CA DataMinder\data\Indexer. The OPT file parameters mentioned in this document are capitalized.

Important: Changing job settings using the Content Indexer Console reverts manual changes to parameters in the OPT file. Configure the job first using the Content Indexer Console, and manually modify the remaining settings in the OPT file afterwards. Restart the CA DataMinder Content Indexer service to load the changes.

OUTPUTFLAGS—Generate EVL Files

Use the OUTPUTFLAGS option to configure jobs to generate EVLs. If a retry or standard job fails to fully index events, the events are written to subfolders of the EVLFAILDEST directory.

Note: This option corresponds to clicking Generate EVL Files on Failure in the Content Indexer Console.

This option accepts the following values:

0

Specifies that no EVL files are generated.

1

Specifies that EVL files are generated.

EVLFAILDEST—Specify the EVL Directory

Use EVLFAILDEST to define the path where EVL files are saved. The files are created in new subdirectories of the provided EVLFAILDEST.

EVLRETRYSRC—Specify the Retry Directory

Use EVLRETRYSRC to define the path where directories of retry jobs are stored. If the EVLRECURSE option is set, retry jobs drill down into subfolders of EVLRETRYSRC when they are looking for EVLs to process.

EVLSPLIT—Sort EVL Folders by Eventimestamp

Use the EVLSPLIT option to specify the subfolder groupingWhen an EVL file is created, it is written to a subfolder of EVLFAILDEST. The subfolder are grouped by their eventtimestamp.

This option accepts the following values:

0

Specifies that there is no eventtimestamp grouping.

1

Specifies that subfolders use the year (YYYY) as timestamp grouping.

2

Specifies that subfolders use year and month (YYYY-MM) as timestamp grouping. This value is the default.

3

Specifies that subfolders use the year, month, and day (YYYY-MM-DD) as timestamp grouping

Example: Year-Month

If you set EVLSPLIT to 2, the folders are grouped by year and month timestamps. The Eventtimestamps are rounded down to the nearest month. For example, all events that fail in August 2011 are written to a folder named '2011-08'.

EVLRECURSE—Retry Failed Subfolders

Use the EVLRECURSE option to specify whether a retry job drills down into subfolders of EVLRETRYSRC when it is looking for EVLs to process.

This option accepts the following values:

0

Specifies that retry jobs do not recurse nor drill down.

1

Specifies that retry jobs recurse and drill down into subfolders.

EVLRETRIES—Specify the Number of Retries

Use EVLRETRIES to define the maximum number of times an EVL is retried before it is written to the ‘no retry’ folder.

Note: All EVL files that are not in the ‘no retry’ folder and that are targeted by a retry job are tried at least once.

This option accepts the following values:

-1

Specifies that there is no retry limit. All EVL files are generated with a retry counter of 0.

0

Specifies that there is no retry limit. EVL files cannot be written to the ‘no retry’ folder.

>0

Specifies the number of times that an EVL can be retried.

TIMESLICE—Configure Batch Time Period

Events are indexed in chronological order by event creation time (updatetimestamp) and in batches. Use the TIMESLICE option to configure the period of time that a batch spans.

Default: 60 minutes (one hour)

Troubleshooting the Content Indexer

Documents Submitted to Content Database Time Out

Symptom:

Documents that I submitted to the content database time out when I use the IDOL content connector.

Solution:

Timeouts occur when indexing a document into a content database takes longer than a preconfigured upper limit. Indexing can take longer, for example, because the document is large, or the server is too busy to index the document fast enough.

You can improve performance by performing one or all of the following solutions.

- Extend the index timeout limit.

  a. Edit the registry on the server hosting the IDOL connector framework and content database.

  b. Go to the following registry key:

  ```
  HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
      \CurrentVersion\Content Services\IDOL Provider
  ```

  c. Modify the IndexTimeoutMins registry value.

  Because content databases vary in size, we cannot give specific recommendations for the new timeout value. Investigate different timeouts until you discover the optimal timeout for your content database.

- Reduce the number of threads configured in the Content Indexer job. The default number of threads is 2 but you can specify values up to 26. Decrement the number of threads until you find the optimal value.

- Improve the performance of the IDOL server machine.

- Configure the Content Indexer job to generate EVL's on failure and re-ingest the failed events using a Retry Job at a quieter period.

Content Indexer Fails to Initialize the Content Connector

Symptom:

After a 'CA DataMinder Content Database' service restart, Content Indexer jobs fail to start and reports "Failed to initialize the Content Connector".

Reason:

The 'CA DataMinder Connector Framework' service has lost communication with the Content Database.

Solution:

Restart the 'CA DataMinder Connector Framework' and then the 'CA DataMinder Content Indexer' service.

How do I change the location EVLs are written to?

Symptom:

I need to change the default location where the EVL files are stored.

Solution:

The EVL destination can be changed by manually editing the OPT file of the CI job. Update the EVLFAILDEST parameter to refer to the location of choice.

No Events Found in Time Slices

Symptom:

With full logging turned on, the log of a CI job contains repeated I1027 informational messages saying that no events are found in time slices. The log shows that no events were indexed between the messages.

Reason:

The CMS contains time periods when no events are captured.

Solution:

Increase the TIMESLICE setting of the job to reduce the number of time slices that fail to find events.

The default time slice value in the OPT file is 60 minutes (1 hour). Increase the size of a time slice until the performance of the content indexer improves. For example, try setting it to 1440 minutes (a day).

More information:

TIMESLICE—Configure Batch Time Period (see page 430)

IDOL server times out after upgrade

Symptom:

After repairing or reinstalling the CA DataMinder IDOL Connector, my indexing job stalls and documents are timed out before they can be added to the content database.

The index.log file on the IDOL server contains repeated messages in this format:
```
"Database (<database name>) is readonly. Document <file name> not indexed."
```

Reason:

Timeouts occur when indexing a document into a content database takes longer than a preconfigured upper limit. In this case, the time-outs occur because the target database has been flagged as Read Only in the content.cfg configuration file. This change to content.cfg is caused by the following sequence of events:

1. When you repair or reinstall the CA DataMinder IDOL Connector, content.cfg gets overwritten with a new version. The new version of content.cfg does not list the constituent databases in the content database.
2. When the CA DataMinder Content Database service next starts, it scans all the constituent databases in the content database. If a constituent database exists but has no entry in the content.cfg configuration file, the service creates an entry and marks the constituent database as Read Only.
3. While the CA DataMinder Content Database service is running, the CA DataMinder IDOL Connector updates content.cfg and flags the constituent databases as Read-Write.

   However, the Connector cannot force the service to reread content.cfg.

Solution:

The CA DataMinder Content Database service only reads the content.cfg configuration file when the service starts. Therefore you must follow these steps:

1. Restart the CA DataMinder Content Database service
2. Restart any indexing jobs that are currently running.

More information:

TIMESLICE—Configure Batch Time Period (see page 430)

Chapter 18: EML Conversion Utilities: Cnv2email and BB2email

This section contains the following topics:

About Cnv2email and BB2email

This section introduces the EML conversion utilities, Cnv2email.exe and BB2email.exe. These utilities convert various types of message data into EML files (that is, Internet emails) that can be subsequently processed by policy engines and imported into the CMS. This section also describes how to configure policy engines to detect information stored in custom x-headers within the resulting EML files.

- **Cnv2email.exe** converts IM conversations saved in CNV files to EML files.

  Note that CNV files are generated by the IM Import utility, IMFrontEnd.exe.

- **BB2email.exe** converts Bloomberg messages to EML files.

  Note that Bloomberg messages are emails sent using Bloomberg terminals.

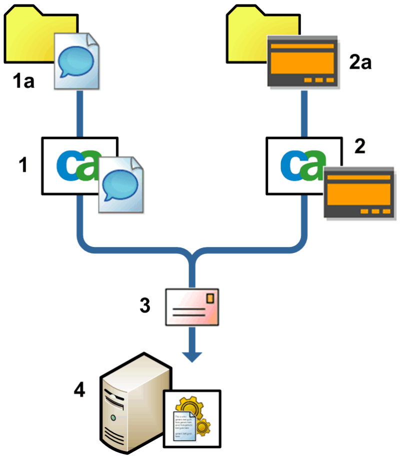

Cnv2email.exe and BB2email.exe

1. Cnv2email.exe retrieves CNV files (**1a**) and converts them to EML e-mails (**3**).
2. BB2email.exe retrieves Bloomberg messages from XML dump files (**2a**) and converts them to EML emails (**3**).
3. Custom x-headers in the EML files contain details about the embedded IM conversation or Bloomberg message.
4. Policy engines detect the x-headers and use email triggers to process the EML files (for example, to apply smart tags).

Deployment Architecture

The diagram below summarizes the deployment architecture for the EML conversion utilities, Cnv2email.exe and BB2email.exe.

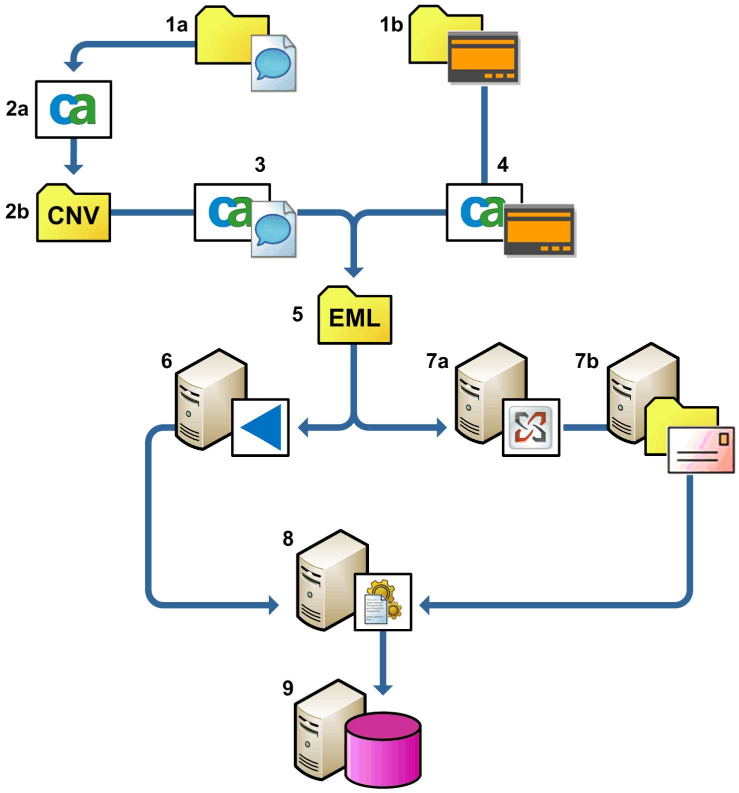

EML utilities: deployment architecture

Dump files contain IM data (**1a**) and Bloomberg emails (**1b**).

IMFrontEnd.exe (**2a**) extracts IM conversations from the dump files and saves them as .CNV files (**2b**). .

Cnv2email.exe (**3**) converts the CNV files to EML files (**5**), and outputs them either to a source folder for Event Import (**6**) or an Exchange pickup folder (**7a**) for ingestion into an existing email archive system (**7b**).

BB2email.exe (**4**) extracts Bloomberg emails from XML dump files (**1b**) and converts them to EML files (**5**). As with Cnv2email.exe, these EML files can be written to a source folder for Event Import (**6**) or an Exchange pickup folder (**7a**).

After processing by Event Import, EML files can be forwarded to a policy engine (**8**) for processing (for example, for classification or to apply smart tags) before being saved on the CMS (**9**). You must run an Import Policy job to do this. Archived EML emails (**7b**) can also be forwarded to a policy engine if the archive is integrated with CA DataMinder.

Cnv2email.exe Utility

The Cnv2email.exe utility can convert IM conversations saved in CNV files to Internet emails, that is, EML files. CNV files contain IM conversations extracted from dump files. EML is the file format for Internet e-mails.

If required, lengthy IM conversations can be subdivided into chapters so they are easier to review.

The resulting EML files can then be imported directly into the CMS, or they can be sent to an Exchange mailbox (via the pickup folder) to await retrieval by a third party archiving solution.

In both cases, you can apply policy to the IM conversations, either by importing them as an Import Policy job or by integrating CA DataMinder with your e-mail archive solution and applying policy before the EML emails are archived.

After an EML email has been processed by a policy engine, it is stored on the CMS as an ‘instant message’ event and can be searched for (in the iConsole standard search) by specifying the IM Network. (Note that even EML emails ingested from an Exchange mailbox retain an ‘instant message’ event type.)

Also, the participants in each IM conversation are stored as *recipients* in the resulting EML email (providing that the participants’ email addresses were available in the original IM dump file); conversely, the email sender is a configurable parameter (see EML.From) and determines which policy is applied to the EML e-mail.

IM Conversion Process

The conversion process involves the following steps:

1. The IM Import utility, IMFrontEnd.exe, extracts IM conversations from dump files and converts them into CNV files.
2. Cnv2email.exe converts the resulting CNV files into EML files, chapterized if necessary.

   Conversion operations are configured using parameters in a configuration file, cnv2email.ini.
3. These EML files can then be:
   - Imported directly into the CMS by running an Event Import job or Import Policy job, or
   - Forwarded to an Exchange journal mailbox. You typically choose this method if your existing email archive solution extracts emails from Exchange journal mailboxes. You can then import the EML files from the journal mailbox into the CMS using the appropriate CA DataMinder archive integration feature.

Install Cnv2email.exe

The process of converting IM conversations to Internet emails uses Cnv2email.exe and IMFrontEnd.exe.

To install Cnv2email.exe and IMFrontEnd.exe

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.
3. In the Advanced Install Options screen, choose Server Agents and then click Install.

   This launches the Integration Agents installation wizard in a separate window.

4. In the Integration Agents installation wizard, navigate to the Customer Setup screen.
5. In the Custom Setup screen, choose Bloomberg and IM Support.

   Cnv2email.exe and IMFrontEnd.exe are installed as part of this feature.

6. In the final wizard screen, click Install to start the file transfer.

   Cnv2email.exe and IMFrontEnd.exe are installed to the \Import subfolder in the CA DataMinder installation folder.

   A sample configuration file, Cnv2email.ini, is installed to the \Import\Templates subfolder in the CA DataMinder installation folder.

7. Configure Cnv2email.ini to suit your network environment and conversion requirements.

Run a CNV Conversion Job

You typically run Cnv2email.exe from a command line. You configure conversion operations by specifying parameters in the configuration file, cnv2email.ini. This configuration file must be in the same folder as Cnv2email.exe and IMFrontEnd.exe.

In addition, be aware that:

- Cnv2email.exe and IMFrontEnd.exe must be configured to use a shared folder for file interchanges. For details, see the File.Pathspec parameter.
- Cnv2email.exe can run in continuous mode or once-only mode. To set up continuous unattended CNV conversion operations, you must set File.ContinuousInput=Yes.
- You can output EML emails to:
  - A ‘standard’ folder for subsequent retrieval by an Event Import or Import Policy job, or
  - An Exchange pickup folder, from where they are subsequently moved to the specified target mailbox.

  For details, see the FileOut.Directory parameter.
- Chapter-splitting parameters enable you to split lengthy IM conversations into chapters based on a regular timeout, or the number of comments submitted, or both. For details, see IM.ChapterTimeoutMins and IM.ChapterMaxMessages parameters.

BB2email.exe Utility

The BB2email.exe utility can convert Bloomberg emails (often referred to simply as ‘Bloombergs’) to Internet e-mails (that is, EML files; this is the file format for Internet e-mails).

The resulting EML files can then be imported directly into the CMS, or they can be sent to an Exchange mailbox (via the pickup folder) to await retrieval by a third party archiving solution.

In both cases, you can apply policy to the Bloomberg emails, either by importing them as an Import Policy job or by integrating CA DataMinder with your e-mail archive solution and applying policy before the EML emails are archived.

Policy engines can recognize a converted EML file as ‘Bloomberg’ via the x-headers that BB2email.exe includes in each EML file. BB2email.exe preserves all user details in custom x-headers enabling them to be used as a basis for applying policy.

After an EML email has been processed by a policy engine, it is stored on the CMS as an ‘embedded Bloomberg email’ event and can be searched for (in the iConsole standard search) by specifying ‘Bloomberg Events’.

Bloomberg Email Conversion Process

The conversion process involves the following steps:

1. BB2email.exe extracts Bloomberg emails from XML dump files (along with any attachments) and converts them into EML files.

   Conversion operations are configured using parameters in a configuration file, BB2email.ini.

2. These EML files can then be:

   - Imported directly into the CMS by running an Event Import job or Import Policy job, or
   - Forwarded to an Exchange journal mailbox. You typically choose this method if your existing email archive solution extracts emails from Exchange journal mailboxes. You can then import the EML files from the journal mailbox into the CMS using the appropriate CA DataMinder archive integration feature (see Third party integration).

Install BB2email.exe

The process of converting Bloomberg emails to Internet emails uses BB2email.exe.

To install BB2email.exe

1. Find setup.exe in the root of your CA DataMinder distribution image. Run setup.exe to launch the CA DataMinder installation wizard.

   The Installation Type screen opens.

2. Click Advanced Installation.
3. In the Advanced Install Options screen, choose Server Agents and then click Install.

   This launches the CA DataMinder Integration Agents installation wizard in a separate window.

4. In the Integration Agents installation wizard, navigate to the Customer Setup screen.
5. In the Custom Setup screen, choose Bloomberg and IM Support.

   BB2email.exe is installed as part of this feature.

6. In the final wizard screen, click Install to start the file transfer.

   BB2email.exe is installed to the \Import subfolder in the CA DataMinder installation folder.

   A sample configuration file, BB2email.ini, is installed to the \Import\Templates subfolder in the CA DataMinder installation folder.

7. Configure BB2email.ini to suit your network environment and conversion requirements.

Set Up BB2email.exe to Ingest Attachments

You can configure BB2email.exe to ingest Bloomberg email attachments from the XML file. To do this:

1. Specify the folder to search for attachments using the Attachment.Folder parameter.
2. You can then configure BB2email.exe to:

   - **Search subdirectories** of the location specified in the Attachment.IncludeSubdirs parameter to look for further attachments.
   - **Delete the attachment** after the email has been successfully imported using the Attachment.Delete parameter.
   - **Fail the import process** if an email’s attachment cannot be accessed. To do this, use the Attachment.RejectOnFailure parameter.

Run a Bloomberg Email Conversion Job

You typically run BB2email.exe from a command line. You configure conversion operations by specifying parameters in the configuration file, BB2email.ini. This configuration file must be in the same folder as BB2email.exe.

In addition, be aware that:

- BB2email.exe can run in once-only mode, or continuous mode. To set up continuous unattended XML conversion operations, you must set File.ContinuousInput=Yes.
- You can output EML emails to:
    - A ‘standard’ folder for subsequent retrieval by an Event Import or Import Policy job, or
    - An Exchange pickup folder, from where they are subsequently moved to the specified target mailbox. For details, see FileOut.Folder.

Configure EML Conversion Jobs

You configure jobs to convert CNV files and Bloomberg messages to EML files by specifying parameters in configuration files.

To configure EML conversion operations

1. Create a configuration file.

   - For **Cnv2email.exe**, name the file **cnv2email.ini** and save it in the same folder as cnv2email.exe and imfrontend.exe.
   - For **BB2email.exe**, name the file **BB2email.ini** and save it in the same folder as BB2email.exe.

2. Specify the conversion parameters in the configuration file.

The available parameters are:

```
Attachment.Delete
Attachment.Folder
Attachment.IncludeSubdirs
Attachment.RejectOnFailure
EML.From
EML.TargetMailbox
EML.To
File.ContinuousInput
File.DeleteAfterImport
File.IncludeSubdirs
File.Pathspec
FileOut.Directory
FileOut.Folder
FileOut.MaxFilesPerFolder
IM.ChapterMaxMessages
IM.ChapterTimeoutMins
Log.Level
Log.MaxNumberFiles
Log.MaxSizeInBytes
```

More information:

EML Conversion Parameters (see page 445)

EML Conversion Parameters

These parameters can be used to configure CNV2email.exe and BB2email.exe conversion jobs.

Attachment.Folder

(BB2email.exe only) This parameter has no default value. It can specify either a shortcut, or a fully-qualified path to a folder containing attachment files for imported conversations and emails. The syntax is:

```
Attachment.Folder=<folder path>
```

This parameter can specify:

- A folder containing the source attachment files for imported conversations and emails.
- A folder containing subfolders that contain attachment files. The name of the subfolder must match the name of the dump file to which these attachments belong.
- The source attachment files, plus any subfolders, if Attachment.IncludeSubdirs=Yes. We do not recommend this configuration, as it is the least efficient option.

Note: Do not use quote marks around pathnames, even if they contain spaces.

Note: There are no attachment parameters for Cnv2email.exe. This is because any IM attachments will already have been included in the source CNV files by IMFrontEnd.exe.

Attachment.IncludeSubdirs

(BB2email.exe only) Defaults to No. This parameter determines whether subfolders within the Attachment.Folder will be included when BB2email.exe is searching for attachment data files. The syntax is:

```
Attachment.IncludeSubdirs=<Yes or No>
```

Attachment.Delete

(BB2email.exe only) Defaults to No. This parameter specifies whether to delete email attachments from the source folder after they have been imported. The syntax is:

```
Attachment.Delete=<Yes or No>
```

If set to Yes, attachments are deleted; if set to No, they are not deleted.

Note: Attachment source folders are defined by Attachment.Folder.

Attachment.RejectOnFailure

(BB2email.exe only) Defaults to No. An attachment file can sometimes fail to be imported, for example if an attachment referenced in the log file being imported cannot be found. The syntax is:

```
Attachment.RejectOnFailure=<Yes or No>
```

This parameter specifies whether the BB2email.exe process ignores the attachment and continues with the import process, or fails this log file and moves on to the next. If set to:

Yes

BB2email.exe fails the attachment's log file and moves it to the \Failed subfolder.

No

BB2email.exe ignores the failed attachment and continues to import the current log file.

EML.From

(Cnv2email.exe only) This parameter has no default value. It specifies the email address that gets written to the From: field in each EML email. The syntax is:

```
EML.From=<email address>
```

This parameter is only needed if you intend to run an Import Policy job in order to apply policy to these converted IM conversations before importing them into the CMS.

Typically, EML.from is set to an email address that corresponds to a custom CA DataMinder user account whose policy has been specifically tailored to process IM conversations. (When you run an Import Policy job, the policy engine uses this From: address to determine which policy to apply. Specifically, it maps the address to an existing CA DataMinder user account and applies outgoing email triggers in that user's policy to the EML email.)

EML.SubjectPrefix

(BB2email.exe only) This parameter has no default value. It specifies an optional prefix that gets appended in front of the Subject to allow these emails to be easily identified. The syntax is:

```
EML.SubjectPrefix=Bloomberg EMAIL:
```

EML.TargetMailbox

(BB2email.exe only) This parameter has no default value. It specifies a target SMTP email address. The syntax is:

```
EML.TargetMailbox=<email address>
```

This parameter is only needed if you intend to forward converted EML emails to an Exchange journal mailbox. If so, this address determines the target mailbox when ingesting the emails into Exchange.

EML.To

(Cnv2email.exe only) This parameter has no default value. It specifies a target SMTP email address. The syntax is:

```
EML.To=<email address>
```

This parameter is only needed if you intend to forward converted EML emails to an Exchange journal mailbox. If so, this address determines the target mailbox when ingesting the emails into Exchange.

EML.UseCorporateAddresses

(BB2email.exe only) Defaults to No. This paramerter specifies whether to use corporate addresses for the sender and recipient instead of Bloomberg addresses. The syntax is:

```
EML.UseCorporateAddresses=<Yes or No>
```

If set to Yes, BB2email.exe uses corporate addresses, if available. If set to No, BB2email.exe uses the Bloomberg address.

If set to Yes but no corporate address is present, BB2email.exe uses the Bloomberg address.

File.Pathspec

This mandatory parameter has no default value. It specifies a fully-qualified path to the folder containing the source files that you want to convert. The syntax is:

```
File.Pathspec=<file path>
```

For **Cnv2email.exe,** this folder must contain the source CNV files that you want to convert, for example C:\IM_msgs. The folder must correspond to the IM Import output folder, that is, the folder specified by the Import Policy parameter DirFinalDest.

For **BB2email.exe**, this folder must contain the source XML dump files containing the Bloomberg emails that you want to convert, for example C:\BB_msgs.

Note: You only need to specify a path to the source folder. You do not need to append *.cnv or *.xml to the path. Cnv2email.exe will automatically detect CNV files, and BB2email.exe will automatically detect XML files, in the source folder.

File.IncludeSubdirs

Defaults to No. This parameter specifies whether to search for matching CNV or XML files in subfolders below the source folder specified by File.Pathspec. The syntax is:

```
File.IncludeSubdirs=Yes or No
```

File.DeleteAfterImport

Defaults to No. This parameter specifies whether to delete the source CNV or XML files after processing. The syntax is:

```
File.DeleteAfterImport=Yes, No, or Always
```

If this parameter is set to:

Yes

The source CNV or XML file is deleted after it has been successfully processed.

No

Source CNV or XML files are never deleted. If used with File.ContinuousInput=Yes, then the conversion job will fail to start, because you cannot use these parameters together in such a way.

Always

Source CNV or XML files are always deleted, whether the conversion succeeds or not, and even if File.ContinuousInput=Yes.

File.ContinuousInput

Defaults to No. This parameter specifies whether Cnv2email.exe or BB2email.exe repeatedly scans for and converts CNV or XML files specified by File.Pathspec, or whether it shuts down after the input folders and files have been processed. The syntax is:

```
File.ContinuousInput=Yes or No
```

If this parameter is set to:

Yes

The source folder is continuously monitored and any CNV or XML files written to there are processed automatically. Typically:

Cnv2email.exe: Choose **Yes** if IM Import is also set to run continuously. That is, the Import Policy parameter RunOnce is set to No.

BB2email.exe: Choose **Yes** if Event Import is also set to run continuously. That is, the Event Import parameter File.ContinuousInput is set to Yes.

Note also:

- CNV or XML files are always deleted after being successfully processed. That is, the File.DeleteAfterImport parameter is ignored.
- Any CNV or XML files that cannot be converted are moved to a \Failed subfolder.
- The File.IncludeSubdirs parameter is ignored. In continuous mode, Cnv2email.exe and BB2email.exe do not check for or process files in subfolders below the source folder.

No

After any CNV or XML files in the source folder have been processed, Cnv2email.exe and BB2email.exe quit.

- The File.DeleteAfterImport parameter is adhered to. That is, CNV and XML files can be either deleted or retained after processing, as required.
- The File.IncludeSubdirs parameter is adhered to (if set). That is, CNV and XML files in any subfolders below the source folder are processed as well.

FileOut.Directory

(Cnv2email.exe only) This parameter has no default value. It specifies the target output folder for the EML files. The syntax is:

```
FileOut.Directory=<folder path>
```

Event Import or Import Policy can retrieve the EML files from this folder and import them directly into CA DataMinder.

Alternatively, if you want to automate integration with your Exchange-based archive solution, set this parameter to be your Exchange pickup folder. When the EML emails generated by Cnv2email.exe are written to the pickup folder, they are automatically processed by Exchange and forwarded to the mailbox specified by the EML.To parameter.

FileOut.Folder

(BB2email.exe only) This parameter has no default value. It specifies the target output folder for the EML files. The syntax is:

```
FileOut.Folder=<folder path>
```

Event Import or Import Policy can retrieve the EML files from this folder and import them directly into CA DataMinder.

Alternatively, if you want to automate integration with your Exchange-based archive solution, set this parameter to be your Exchange pickup folder. When the EML emails generated by BB2email.exe are written to the pickup folder, they are automatically processed by Exchange and forwarded to the mailbox specified by the EML.TargetMailbox parameter.

FileOut.MaxFilesPerFolder

This optional parameter specifies the maximum number of EML emails that can be written to each output subfolder. The syntax is:

```
FileOut.MaxFilesPerFolder=<number of emails>
```

If you include this parameter in your Cnv2email.ini or BB2email.ini configuration file, EML emails are written to a series of subfolders below the output folder specified by FileOut.Directory or FileOut.Folder. When the maximum number of EML emails has been written to the current subfolder, a new subfolder is created (numbered sequentially) and all subsequent EML emails are written to that folder until the maximum is reached once more, when the process repeats. The output subfolders are structured as follows:

1 Set by the FileOut.Directory or FileOut.Folder parameter

2 GUID generated by Cnv2email.exe or BB2email.exe

3 Subfolders numbered sequentially

Important! Do not use FileOut.MaxFilesPerFolder if Cnv2email.exe or BB2email.exe are configured to write EML files to an Exchange pickup folder (see FileOut.Folder or FileOut.Directory). EML files are **not** retrieved from subfolders below an Exchange pickup folder.

IM.ChapterTimeoutMins

(Cnv2email.exe only) Defaults to 60 minutes. This ‘conversation-splitting’ parameter enables Cnv2email.exe to split IM conversations into separate chapters based on a regular timeout interval, starting from the first message in the first chapter of the conversation. The syntax is:

```
IM.ChapterTimeoutMins=<number of minutes>
```

For example, to create a new chapter after every two hours, set this parameter to:

```
IM.ChapterTimeoutMins=120
```

Note: A single configuration file can contain both the ‘timeout’ and ‘maximum message count’ conversation-splitting parameters. Cnv2email.exe monitors both parameters simultaneously and creates a new chapter as soon as either meets its criteria.

IM.ChapterMaxMessages

(For Cnv2email.exe only) Defaults to 50 messages. This ‘conversation-splitting’ parameter enables Cnv2email.exe to split IM conversations into separate chapters based on the number of messages or comments submitted. The syntax is:

```
IM.ChapterMaxMessages=<number of messages>
```

For example, to create a new chapter after every 100 messages, set this parameter to:

```
IM.ChapterMaxMessages=100
```

Note: A single configuration file can contain both the ‘timeout’ and ‘maximum message count’ conversation-splitting parameters. Cnv2email.exe monitors both parameters simultaneously and creates a new chapter as soon as either meets its criteria.

Log.Level

Defaults to 2. This determines the level of logging for the conversion process. For example, you can configure Cnv2email.exe or BB2email.exe to only log errors or important system messages. The syntax is:

```
Log.Level=<number>
```

The supported logging levels are:

0 No logging

1 Errors only

2 Errors and warnings

3 Errors and warnings, plus informational messages

4 Errors and warnings, plus informational messages, and trace

Note: Setting Log.Level=4 will cause the log file to grow extremely rapidly. This level of logging is provided for testing purposes only.

Log entries are written to Cnv2email_<date>.log or BB2email_<date>.log, where <date> is the date and time when the log file was created.

Log.MaxSizeInBytes

Defaults to 100,000. This specifies the maximum size (in bytes) for each log file. The syntax is:

```
Log.MaxSizeInBytes=<number of bytes>
```

When the current log file reaches its maximum size, Cnv2email.exe or BB2email.exe creates a new log file.

Log.MaxNumberFiles

Defaults to 100. This specifies the maximum number of log files. The syntax is:

```
Log.MaxNumberFiles=<number of files>
```

When the maximum number of log files exist, and the most recent log file reaches the maximum log file size (see Log.MaxSizeInBytes), the oldest log file is deleted to enable a new one to be created.

Embedded Content Details Saved in EML x-headers

EML utilities Cnv2email.exe and BB2email.exe generate **embedded content emails**. These are EML files that contain, respectively, IM conversations and Bloomberg messages. This section explains how CA DataMinder policy engines handle embedded content emails to ensure that the event type and, for IM conversations, the IM network is correctly stored on the CMS.

Custom x-headers in EML Emails Contain Event Details

When IM conversations and Bloomberg messages are converted to EML emails, CA DataMinder policy engines need to distinguish between ‘genuine’ emails sent using Outlook or Notes and ‘embedded content’ emails that are really IM conversations or Bloomberg messages.

However, details such as the event type, IM network and event participants cannot be saved as conventional email attributes. To overcome this problem, Cnv2email.exe and BB2email.exe store these details as bespoke x-headers in the EML email. Policy engines automatically detect these x-headers when processing the EML email and store the relevant details with the event on the CMS.

Configure Policy Engines to Detect Embedded Content Emails

Policy engines need to distinguish between ‘genuine’ emails and ‘embedded content’ emails. This is accomplished through a setting in the machine policy. The **Embedded Message Identification** setting in the \Policy Engine folder enables policy engines to detect embedded content emails and set the event type as ‘embedded IM’, ‘Bloomberg’, or ‘eFax’. For IM conversations, this setting can also be used to extract or set the IM network.

Important! You must configure **outgoing email triggers** in the user policy to capture or apply smart tags to embedded IM conversations or Bloomberg messages!

Automatic IM and Bloomberg Detection

By default, policy engines can detect ‘IM’ EML emails generated by Cnv2email.exe and ‘Bloomberg’ EML emails generated by BB2email.exe. If you only need policy engines to detect and process these event types, *you do not need to modify the default values for the Embedded Message Identification policy setting!*

Detecting Embedded IM Conversations

By default, machine policy is configured to automatically identify IM conversations embedded in EML emails generated by Cnv2email.exe and to set the event type to ‘embedded IM’ and extract the IM network from the relevant x-headers. The Embedded Message Identification policy setting defaults to:

```
IM:where "X-ORCH-Network" ExtractIMNetwork "[\"%all%\"]"
```

Important! Do not change this value! The "[\"%all%\"]" extracts the relevant text string from the X-ORCH-Network x-header.

Detecting Embedded Bloomberg Messages

By default, machine policy is configured to automatically identify Bloomberg messages embedded in EML emails generated by BB2email.exe and set the event type to ‘Bloomberg’. The Embedded Message Identification policy setting defaults to:

```
BB:where "X-ORCH-MessageType" is "Bloomberg"
```

Important! Do not change this value!

Detect Other Embedded Content

If required, you can configure policy engines to detect other forms of embedded content such as eFaxes or IM conversations embedded in EML files that were generated by third party applications. To do this, you must add additional values to the **Embedded Message Identification** machine policy setting.

More information:

Conversion Expression Syntax (see page 299)

Syntax for ‘Embedded Message Identification’ Policy Setting

If you want policy engines to detect other forms of embedded content in EML emails (such as eFaxes or IM conversations converted by third party applications), you must add additional values to the Embedded Message Identification policy setting, using this syntax:

```
<TYPE>:where "<X-header>"
[IS or CONTAINS "<Header value>"]
[<IM action> "<Network>"]
```

where:

<TYPE>

Determines the event type. It can be IM or BB (for Bloomberg messages). In addition, this can be set to EFAX to detect eFaxes (see next section).

<X-header>

Is the name of the x-header that the policy engine must look for when processing the EML email. You must enclose the x-header name in double quotes.

IS or CONTAINS <Header value>

Validates the x-header. These operators test whether the actual text value of the x-header is, or contains, the text defined by <Header value>. You must enclose the text in double quotes.

Note: These operators are not needed to process 'IM' EML files generated by Cnv2email.exe. Policy engines can automatically detect the IM network embedded in these EML files.

<IM action>

Is only needed when processing 'IM' EML files. It can be either SetIMNetwork or ExtractIMNetwork.

<Network>

Is only needed when processing 'IM' EML files. If <IM action> is set to:

- SetIMNetwork, then <Network> specifies the IM network that you want to assign to the resulting 'embedded IM' event. You must enclose the IM network in double quotes.
- ExtractIMNetwork, then <Network> extracts the specified text from the x-header and saves this text as the IM network when generating an 'embedded IM' event. Note that you can use a **conversion expression** to modify this extracted text before storing the IM network. You must enclose the specified text or conversion expression in double quotes.

These conversion expressions use the same syntax as the conversion expressions for importing LDAP attributes. They can parse, extract and (if necessary) remove or substitute any characters in the IM network.

More information:

Conversion Expression Syntax (see page 299)

Detect an Embedded eFax

eFaxes are faxes sent or received as email attachments using an online fax service. Policy engines do not automatically detect eFaxes embedded in EML emails, but you can add an extra value to the Embedded Message Identification policy setting that enables them to do so. The required syntax is:

```
EFAX:where "<x-header>" IS "<Name>"
```

For example:

```
EFAX:where "UPRXS-eFax" IS "efax"
```

where:

EFAX

Sets the event type to eFax.

<X-header>

Is the name of the x-header that identifies the EML email as containing an eFax. The policy engine looks for this x-header when processing the email. You must enclose the x-header name in double quotes.

IS <Name>

Validates the x-header. These operators test whether the actual text value of the eFax x-header matches the text specified by <Name>. You must enclose the text in double quotes.

Note: You can also use CONTAINS instead of IS.

Chapter 19: Binary Text Extractor

This section contains the following topics:

Overview

The Binary Text Extractor (BTE) is a configurable utility that can extract the text content from document types that are not normally supported by CA DataMinder. Policy engines and endpoint agents can then apply CA DataMinder policy to these files as normal.

For example, if you need to analyze information stored in proprietary or industry file formats, or even in executable files, you can configure the BTE to extract the text content from these file types.

How Does the Binary Text Extractor Work?

You configure the BTE to extract text from specific types of file. If CA DataMinder is unable to analyze a file using standard methods, it calls the BTE. If the BTE recognizes the file type, it extracts the text and passes it to a CA DataMinder policy engine or endpoint agent for analysis. CA DataMinder then applies policy to the file as normal.

A configuration file, BinaryTextorConfig.xml, specifies which file types the BTE can process. For each file type, the configuration file specifies the 'magic number' and the extracted text:

Magic Number

The magic number is a file signature (a text string or hexadecimal string) that identifies the file type that you want the BTE to support.

If the BTE detects that a file contains a magic number listed in BinaryTextorConfig.xml, the BTE proceeds to extract the text content.

Extracted Text

When you specify the text that you want the BTE to extract, you specify the encoding system, the character range, and the minimum string length. The BTE then extracts any text strings that match these requirements.

The encoding system defines how individual text characters are represented in the files that you want to analyze. The BTE supports ASCII, UTF-8, Little Endian UTF-16 (more commonly known as Unicode) and Big Endian UTF16.

The character set defines the range of characters that are eligible for extraction. You typically ignore non-printing characters (such as paragraph markers) and only extract printing characters.

The minimum string length defines the shortest word that you want the BTE to extract. For example, if you want CA DataMinder policies to detect files that contain the word 'Unipraxis', the BTE only needs to extract strings with a minimum length of 9 characters.

How to Configure the Binary Text Extractor

Configuration details for the BTE are saved in BinaryTextorConfig.xml. This is a CA DataMinder system definition file. To configure BinaryTextorConfig.xml, follow these steps:

1. Export the default version of BinaryTextorConfig.xml from the CMS.
2. Edit BinaryTextorConfig.xml to specify the file types that you want CA DataMinder to analyze.
3. Install your customized version of BinaryTextorConfig.xml back onto the CMS.

   CA DataMinder then replicates your customized version of BinaryTextorConfig.xml automatically to each CA DataMinder policy engine and endpoint computer.

If you need to disable the BTE, you must replace the existing configuration file on the CMS with a 'zero file type' file configuration file.

These steps are described in the following sections.

More information:

Export the Default Configuration File (see page 459)
Edit the Default Configuration File (see page 460)
Install a Customized Configuration File (see page 460)
Disable the Binary Text Extractor (see page 460)

Export the Default Configuration File

The default BinaryTextorConfig.xml is a system definition file. It is saved in the \System subfolder of the CA DataMinder installation folder on the CMS. However, we recommend that you use CA DataMinder System File Explorer to export the file for editing.

To export the default BinaryTextorConfig.xml

1. Log on to the Administration console using an account that has the 'Admin: Manage System Files' administrative privilege.
2. Click Tools, Edit System Files.

   The CA DataMinder System File Explorer displays.
3. Browse to the Public, System, Definitions folder.
4. Right-click BinaryTextorConfig.xml and click Export.

Edit the Default Configuration File

Use your preferred XML editor to customize the default BinaryTextorConfig.xml.

The following sections contain details about the XML schema and an example configuration file.

More information:

Example BTE Configuration File (see page 463)
Schema Notes (see page 465)
XML Schema for BTE Configuration File (see page 469)

Install a Customized Configuration File

After you have edited the BTE configuration file, you must reinstall it on the CMS. CA DataMinder then automatically distributes this file to all policy engines and endpoint computers.

To install your customized BinaryTextorConfig.xml

1. Log on to the Administration console using an account that has the 'Admin: Manage System Files' administrative privilege.
2. Click Tools, Install System Definition File.

   The Install System Definition File dialog appears.
3. Click 'Binary Text Extractor Configuration File' in the File Type list.
4. Browse to the BinaryTextorConfig.xml file that you edited.
5. Return to the Install System Definition File dialog and click Install.

   CA DataMinder automatically replicates BinaryTextorConfig.xml to all policy engines and endpoint computers.

Disable the Binary Text Extractor

If you need to disable the BTE, you must create a configuration file with zero *FileType* elements and then install this file on the CMS. Details about creating a 'zero file type' configuration file are included in the Schema Notes (see page 465).

BTE Requirements on FSA Servers

You may need to install the BinaryTextorConfig.xml configuration file manually on a File Scanning Agent (FSA) server.

The BTE is installed automatically on servers hosting the FSA. This enables you to apply CA DataMinder policy to scanned files that have non-standard formats.

When the FSA scans a file, it extracts the text content and passes the content to a remote policy engine for analysis. Consequently, the FSA requires a local copy of the BinaryTextorConfig.xml configuration file in order to scan and extract text from files with non-standard formats. .

Normally, CA DataMinder replicates this configuration file automatically to CA DataMinder policy engines and endpoint computers after you install the file on the CMS. However, CA DataMinder can only replicate BinaryTextorConfig.xml to an FSA host server if the CA DataMinder infrastructure is installed on the server. If the CA DataMinder infrastructure is not installed, you must manually copy BinaryTextorConfig.xml to the FSA server and edit the registry to specify the location of this file.

Do I need to install the BTE configuration file on an FSA server?

You must install BinaryTextorConfig.xml manually if the FSA uses the Remote Policy Engine Connector. (The Connector is a special policy engine hub that an FSA can use to distribute scanned items to remote policy engines for processing. The Connector does not use the CA DataMinder infrastructure.)

You do not need to install BinaryTextorConfig.xml manually if the FSA passes scanned items directly to a local CA DataMinder policy engine. (A CA DataMinder infrastructure is installed on these FSA servers when you install the policy engine.)

Full details about the FSA are in the *Stored Data Integration Guide*.

To manually install the BTE configuration file on an FSA host server

Repeat the following instructions on each FSA server where the FSA needs the BTE to extract text content from files.

1. Export your customized BinaryTextorConfig.xml from the CMS.

   The procedure is the same as exporting the default configuration file.

2. Copy BinaryTextorConfig.xml to the FSA host server.

   You can choose any target folder.

3. Create the following registry key on the FSA host server:

   On 32-bit servers:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
     \CurrentVersion\UserProcess\BinaryTextExtractor
   ```

   On 64-bit servers:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA DataMinder
     \CurrentVersion\UserProcess\BinaryTextExtractor
   ```

4. Within this registry key, add the following registry value:

   ConfigFile

   Type: REG_SZ

   Data: Specify the full path to BinaryTextorConfig.xml on the local FSA server. For example:

   ```
   C:\BTE\BinaryTextorConfig.xml
   ```

Important! If you update BinaryTextorConfig.xml in the future, repeat the above instructions on each affected FSA server.

More information:

Export the Default Configuration File (see page 459)

Example BTE Configuration File

The example BinaryTextorConfig.xml file specifies three file formats:

Executable files

First, the configuration file instructs the Binary Text Extractor to detect executables and .DLLs. These files all start with an ASCII magic number of *"MZ".*

In this example, two <Encoding> elements specify that the Binary Text Extractor extracts any ASCII or Unicode text string of 6 characters or longer. The <CharSet> element specifies the range of characters that are eligible for extraction. This example simply excludes the non-printing characters. CA DataMinder then applies policy to the extracted text.

.AFF files

Second, the configuration file instructs the Binary Text Extractor to detect .AFF files. The magic number for these files comprises an ASCII string *"AFF Rev."* plus an eight character hexadecimal number. The ASCII string and hexadecimal string do not occur concurrently. The configuration file uses two <MagicNumber> elements to define the magic number format.

The *Encoding* element specifies that the Binary Text Extractor extracts any ASCII text string of 5 characters or longer from these .AFF files. As above, the <CharSet> element excludes strings containing non-printing characters. CA DataMinder then applies policy to the extracted text.

.DLIS files

Third, the configuration file instructs the Binary Text Extractor to detect .DLIS files. The magic number for these files comprises an ASCII string in the format *"V?.??RECORD"*. You cannot use wildcards to specify magic numbers. The configuration file therefore uses a combination of three <MagicNumber> elements to define the magic number format.

The *Encoding* element specifies that the Binary Text Extractor extracts any ASCII text string of 5 characters or longer from these .DLIS files. As above, the <CharSet> element excludes strings containing non-printing characters. CA DataMinder then applies policy to the extracted text.

Example Configuration File

```
<?xml version="1.0" encoding="utf-8" ?>
<UniversalBinaryTextor>
    <!-- Executable "MZ"  -->
    <FileType name="Executable/DLL">
      <MagicNumber value="MZ" type="ascii-string" offSet="0" />
        <Encoding name="ASCII" minLength="6">
          <CharSet start="0x20" end="0x7E" />
        </Encoding>
        <Encoding name="UTF16_LITTLEENDIAN" minLength="6">
          <CharSet start="0x20" end="0xFF" />
      </Encoding>
    </FileType>
  <!-- AFF Rev. X.6  -->
  <FileType name="Advanced File Format">
    <MagicNumber value="AFF Rev." type="ascii-string" offSet="56" />
    <MagicNumber value="FFFFFFFF" type="hex-string" offSet="240" />
      <Encoding name="ASCII" minLength="5">
       <CharSet start="0x20" end="0x7E" />
      </Encoding>
  </FileType>
  <!-- Digital Log Interchange Standard (DLIS)
       http://w3.energistics.org/rp66/v1/Toc/main.html   -->
  <FileType name="DLIS">
    <MagicNumber value="RECORD" type="ascii-string" offSet="9" />
    <MagicNumber value="V" type="ascii-string" offSet="4" />
    <MagicNumber value="." type="ascii-string" offSet="6" />
      <Encoding name="ASCII" minLength="5">
        <CharSet start="0x20" end="0x07F" />
      </Encoding>
    </FileType>
  </UniversalBinaryTextor>
```

Note: See the following Schema Notes for guidelines on writing a custom BTE configuration file.

Schema Notes

This section describes the XML elements and attributes that you need to include in a BinaryTextorConfig.xml configuration file.

<UniveralBinaryTextor>

This root element contains zero or more <FileType> elements.

You can use a 'zero file type' configuration file to disable the BTE.

Zero File Type Configuration Files

If the <UniversalBinaryTextor> element contains zero <FileType> elements, the BTE is effectively disabled. A 'zero file type' configuration file is shown below:

```
<?xml version="1.0" encoding="utf-8" ?>
<UniversalBinaryTextor>
  <!--This is an empty configuration file-->
</UniversalBinaryTextor>
```

Note: You cannot use an empty BinaryTextorConfig.xml file to disable the BTE. You must use a 'zero file type' version of BinaryTextorConfig.xml.

<FileType>

This element specifies the type of file that you want the BTE to process. It has the following optional attribute:

name

The name is only used in log files. Use a descriptive name that identifies the file type. For example:

```
<FileType name="AFF Files for use in the Oil Industry">
```

Each <FileType> element can contain any number of <MagicNumber> and <Encoding> sub-elements.

<MagicNumber>

This element specifies the magic number, or file signature, of the file type that you want the BTE to process. You can include multiple <MagicNumber> elements for each <FileType> element.

If the file's magic number does not match the magic number specified in the configuration file, the BTE does not process the file.

This element has the following attributes:

value

This attribute specifies the actual magic number (or part of the magic number) used by the file type.

type

This attribute specifies whether the magic number is a text string or hexadecimal string:

```
type="hex-string"
type="ascii-string"
```

For example, F8DE627B6 and A1A1A1 are valid hexadecimal magic numbers.

Likewise, 'abcPK£' is interpreted as a single byte of ASCII and is valid magic number. But 'Ωω' is not a valid magic number (because Ω and ω are not valid ASCII characters).

offset

This attribute specifies the location of the magic number within the file. The location is specified as a character offset, where zero specifies the first character, 1 specifies the second character, and so on.

Examples

- This example matches files that begin with the hexadecimal 'EFBBBF' magic number.

  ```
  <MagicNumber value="EFBBBF" type="hex-string" offSet="0"/>
  ```

- This example matches any DLL or executable. These files begin with an 'MZ' magic number:

  ```
  <MagicNumber value="MZ" type="ascii-string" offSet="0" />
  ```

<Encoding>

This element specifies the encoding system. The encoding system defines how individual text characters are represented in the files that you want the BTE to process.

The BTE supports four encodings: ASCII, UTF-8, Little Endian UTF-16 (more commonly known as Unicode) and Big Endian UTF16.

The <Encoding> element has these attributes:

name

This attribute specifies the encoding system type. The supported values are:

```
ASCII
UTF8
UTF16-LITTLEENDIAN
UTF16-BIGENDIAN
```

minLength

You can specify the shortest word that you want the BTE to extract. For example, if you want CA DataMinder policies to detect files that contain the word 'Unipraxis', the BTE only needs to extract strings with a minimum length of 9 characters.

This attribute specifies how long a character string must be before the BTE considers it to be a valid string. This example specifies a 6-character ASCII string as the shortest word that you want the BTE to extract:

```
<Encoding name="ASCII" minLength="6">
```

For each <Encoding> element, you must also specify which character ranges are valid. You define valid character ranges in the <CharSet> sub-elements.

<CharSet>

The character set defines the range of characters that are eligible for extraction. You typically ignore non-printing characters (such as paragraph markers) and only extract printing characters.

The <CharSet> elements determine which characters are considered valid constituents of strings. One or more elements are required.

You identify characters by their Unicode code point. You can identify a valid range by specifying the first and last characters in the range, or you can specify a Unicode block.

start, end

These attributes specify the first and last Unicode code points in the character range. Code points are expressed in hexadecimal.

This example specifies the range of printable ASCII characters:

```
<CharSet start="0x20" end="0x7F" />
```

This example specifies the Latin alphabet in lower case:

```
<CharSet start="97" end="122" />
```

blockName

Block names are aliases for character ranges. You can find a list of valid block names at:

http://www.unicode.org/Public/UNIDATA/Blocks.txt

For example, this element is equivalent to the Arabic character range 0x0600..0x06FF:

```
<CharSet blockName="Arabic" />
```

Note: The BTE ignores case, spaces, hyphens, and underscores when checking block names. For example, the BTE interprets these block names as being the same:

```
<CharSet blockName="Basic Latin" />
<CharSet blockName="BasicLatin" />
<CharSet blockName="basic latin" />
```

XML Schema for BTE Configuration File

The Binary Text Extractor configuration file, BinaryTextorConfig.xml, uses the following schema:

```
<?xml version="1.0" encoding="utf-8" ?>
<xs:schema attributeFormDefault="unqualified" elementFormDefault="qualified"
 xmlns:xs="http://www.w3.org/2001/XMLSchema">
<!--  FileType Element Type   -->
<xs:complexType name="FileTypeType">
  <xs:sequence>
    <xs:sequence minOccurs="1" maxOccurs="unbounded">
      <xs:element name="MagicNumber" type="MagicNumberType" />
    </xs:sequence>
    <xs:sequence minOccurs="1" maxOccurs="unbounded">
      <xs:element name="Encoding" type="EncodingType" />
    </xs:sequence>
  </xs:sequence>

  <!--  The magic number can be a hex string or a text string  -->
  <xs:attribute name="name" type="xs:string" />
  </xs:complexType>
  <!--  Top Level Element  -->
  <xs:element name="UniversalBinaryTextor">
    <xs:complexType>
      <xs:sequence>
        <xs:sequence maxOccurs="unbounded" minOccurs="0">
          <xs:element name="FileType" type="FileTypeType" />
        </xs:sequence>
      </xs:sequence>
    </xs:complexType>
  </xs:element>

  <!--  CharSet Element Type  -->
  <xs:complexType name="CharSetType">
    <!--  If both start and end are populated, these are used instead of 'name'  -->
    <!--  'start' and 'end' are in hex, prefixed with 0x eg, 0xF007  -->
    <xs:attribute name="start" type="xs:string" use="optional" />
    <xs:attribute name="end" type="xs:string" use="optional" />
    <!--  CharSet name can be any Block Name from
http://www.unicode.org/Public/UNIDATA/Blocks.txt
     Case, spaces, hyphens and underbars are ignored when comparing block names  -->
    <xs:attribute name="blockName" type="xs:string" use="optional" />
  </xs:complexType>
  <!--  Supported Encodings  -->
  <xs:simpleType name="EncodingEnumType">
    <xs:restriction base="xs:string">
      <xs:enumeration value="ASCII" />
      <xs:enumeration value="UTF8" />
      <xs:enumeration value="UTF16_LITTLEENDIAN" />
      <xs:enumeration value="UTF16_BIGENDIAN" />
    </xs:restriction>
  </xs:simpleType>

  <!--  Encoding Element Type  -->
  <xs:complexType name="EncodingType">
    <xs:sequence minOccurs="1" maxOccurs="unbounded">
      <xs:element name="CharSet" type="CharSetType" />
    </xs:sequence>
    <xs:attribute name="name" type="EncodingEnumType" use="required" />
    <xs:attribute name="minLength" type="xs:unsignedByte" use="required" />
  </xs:complexType>
  <xs:complexType name="MagicNumberType">
    <xs:attribute name="value" type="xs:string" />
    <xs:attribute name="type" type="MagicNumberTypeEnum" />
    <xs:attribute name="offSet" type="xs:integer" />
  </xs:complexType>
  <xs:simpleType name="MagicNumberTypeEnum">
    <xs:restriction base="xs:string">
      <xs:enumeration value="ascii-string" />
      <xs:enumeration value="hex-string" />
    </xs:restriction>
  </xs:simpleType>
</xs:schema>
```

Chapter 20: Universal Extractor

This section introduces the CA DataMinder Universal Extractor utility. The Universal Extractor (UE) is the generic name for various data extraction modules. Each module is designed to extract information from the CMS, typically event metadata, in a specific format.

The following sections describe how to define and run UE jobs, and includes the full XML schema used to define extraction jobs. It then focuses on the first UE module to be released, the XML metadata extractor. It describes how to configure the XML metadata extractor and how to monitor extraction progress.

This section contains the following topics:

UE Requirements

The UE requirements are as follows:

- **Installation:** UE extraction jobs require WgnTask.exe. This utility is installed automatically when you install a CMS.
- **UE database requirements:** For UE extraction jobs, the supported CMS databases are:
  - Microsoft SQL Server 2005 or 2008
  - Oracle 10g (10.2.0.4) and 11g (11.1.0.7) or later . Note the privilege requirements for the Oracle primary user. See the *Database Guide*; search the index for ‘privileges for Oracle users’.

More information:

Run an Extraction Job (see page 472)

Run an Extraction Job

UE extraction jobs are invoked from a command line using the WgnTask.exe utility. This enables you to schedule jobs to run at regular intervals. You configure extraction jobs with an XML job definition document called when you run an extraction job. The command syntax is:

```
WgnTask <job_file> [options]
```

where:

<job_file>

Is the path and file name for the XML job definition document. Available job parameters are listed in XML metadata extractor.

[options]

Can be:

-l

Sets the logging level. The supported logging levels are:

- **1** Errors only. This is the default level.
- **2** Errors and warnings.
- **3** Errors and warnings, plus informational and status messages.

-e

Displays error codes on-screen when an extraction job runs.

-v

Displays verbose (full detail) output on-screen when an extraction job runs.

More information:

XML Schema for UE Job Definitions (see page 472)
XML Metadata Extractor (see page 476)

XML Schema for UE Job Definitions

The XML schema for all UE job definition documents is shown below and on the following page. A job definition document identifies which UE module to run and specifies the job configuration parameters, including the database stored procedures responsible for selecting events for extraction.

Note: Do not confuse the *job definition* schema with the schema for *XML output files*.

More information:

Example UE Job Definition (see page 475)
Format for XML Output Files (see page 477)

UE Job Definition Schema

```
<?xml version="1.0" encoding="UTF-8" ?>
<xsd:schema FormDefault="qualified" xmlns:xsd="http://www.w3.org/2001/XMLSchema"
version="1.0">
  <xsd:annotation>
    <xsd:documentation>
      #*******************************************************
      # Header: Copyright, version source control details
      #*******************************************************
    </xsd:documentation>
  </xsd:annotation>
  <xsd: name="WgnTask">
    <xsd:complexType>
      <xsd:sequence>
        <xsd: minOccurs="0" maxOccurs="1" ref="schedule" />
        <xsd: minOccurs="1" maxOccurs="1" ref="configuration" />
      </xsd:sequence>
      <xsd:attribute name="module_name"  use="required" type="xsd:string" />
      <xsd:attribute name="guid"         use="required" type="xsd:string" />
      <xsd:attribute name="title"        use="optional" type="xsd:string" />
      <xsd:attribute name="description"  use="optional" type="xsd:string" />
      <xsd:attribute name="log_max_file_size"
                                         use="optional" type="xsd:unsignedLong" />
      <xsd:attribute name="log_max_num_files"
                                         use="required" type="xsd:unsignedLong" />
    </xsd:complexType>
  </xsd:>
  <xsd: name="schedule">
    <xsd:complexType>
      <xsd:attribute name="interval"    use="required" type="xsd:unsignedLong" />
      <xsd:attribute name="task_server"  use="optional" type="xsd:string" />
      <xsd:attribute name="parallel"     use="optional" type="xsd:boolean" />
    </xsd:complexType>
  </xsd:>
  <xsd: name="configuration">
    <xsd:complexType>
      <xsd:sequence>
        <xsd: minOccurs="0" maxOccurs="1"          ref="connection" />
        <xsd: minOccurs="0" maxOccurs="unbounded"  ref="item" />
      </xsd:sequence>
    </xsd:complexType>
  </xsd:>
  <xsd: name="item">
    <xsd:complexType>
      <xsd:attribute name="name"   use="required"  type="xsd:string" />
      <xsd:attribute name="value"  use="required"  type="xsd:string" />
    </xsd:complexType>
  </xsd:>
  <xsd: name="connection">
    <xsd:complexType>
      <xsd:attribute name="server"    use="required"  type="xsd:string" />
      <xsd:attribute name="user"      use="optional"  type="xsd:string" />
      <xsd:attribute name="password"  use="optional"  type="xsd:string" />
    </xsd:complexType>
  </xsd:>
```

Example UE Job Definition

This is an example job definition document for the XML metadata extractor.

```
<?xml version="1.0">
<WgnTask module_name="WgnUE_XML"
 guid="521fdd03-d5c5-11d4-b613-000102027bbb"
 title="CMS-HARDY metadata extraction"
 description="Extracts event metadata from CMS-HARDY"
 log_max_file_size"750"
 log_max_num_files"15" />
  <!--
  #*******************************************************
  # Header: Copyright, version source control details
  #*******************************************************
  -->
  <schedule interval="60" />
  <configuration>
    <connection server="CMS-HARDY" user="UEAdmin" password="MyUEPW"/>
    <item name="RemotePrefix"value="" />
    <item name="MaxRecords"value="1000" />
    <item name="RetryPeriod"value="1440" />
    <item name="MaxFileSize"value="0" />
    <item name="DataIntegrity"value="true" />
    <item name="ExportDirectory"value="C:\temp\UE" />
    <item name="ExportPrefix"value="UEUpdates-" />
    <item name="AckDirectory"value="C:\temp\UE" />
    <item name="AckPrefix"value="UEAck-" />
  </configuration>
</WgnTask>
```

XML Metadata Extractor

The XML metadata extractor is a UE module that extracts metadata associated with CA DataMinder events stored on the CMS whose actual message data is stored in a third-party archive. The extracted metedata is written to XML output files.

This enables organizations to import the extracted metadata into their archive and store it with the actual message data. In effect, it enables organizations to designate their archive solution as a definitive system of record for all event data.

Extracting event metadata to XML output files is a two-step process:

1. Create an XML job definition document. This contains the configuration parameters and defines which events are eligible for extraction.
2. Schedule the job. Because the extraction jobs are run from a command line, you can use any compatible scheduler to run regular extraction jobs.

Note: Instructions for importing the metadata from the XML output files into your archive are beyond the scope of this Deployment guide.

More information:

XML Schema for UE Job Definitions (see page 472)
Example UE Job Definition (see page 475)
Extraction Job Parameters (see page 478)
Run an Extraction Job (see page 472)

Extracted Events

By default, the selection criteria specify that to be eligible for extraction, an event must:

- Match the RemoteID prefix requirement (the RemoteID is an identifier used to associate events stored on the CMS with their actual message data held in a third party archive).
- Have at least one associated trigger, issue or audit record. This ensures that only 'significant' events are extracted to the archive.
- Not have been previously extracted or, if it **has** been extracted previously:
  - It was not successfully committed (for example, because a problem prevented the XML files from being written), or
  - No acknowledgement was received from the target archive within the specified retry period, or
  - It has since been updated (for example, the expiry date has been changed), or
  - New issues have been created for that event, or its event history has been updated.

The database SP selects the events that are to be extracted. For Oracle CMSs, event selection is incorporated into the Wgn_V_UE_EventExport database view; for SQL Server CMSs, event selection is handled by a user-defined function, WgnUEExportEvents.

If you want to change the selection criteria (for example, to export all events regardless of triggers, issues and audit records), you can do this by modifying and reinstalling the UE stored procedures and database views. For guidance on reinstalling modified SPs and database views, please contact CA Technical Support (see page 21).

How Event Metadata Is Selected for Extraction

The XML metadata extractor uses a database view, Wgn_V_UE_EventExport, to specify which items of metadata will be extracted and the structure of the resulting XML output file. Note that the view syntax varies according to whether the CMS uses an Oracle or SQL Server database.

To add or remove attributes or s from the metedata extraction, you need to modify and reinstall this database view. For guidance on modifying and reinstalling this view, please contact CA Technical Support.

Format for XML Output Files

The schema for the XML output files is available in a separate technical note, available from CA Technical Support.

Archive Acknowledgements

You can configure the UE job to request acknowledgements from the target archive when it has successfully received and stored the metadata for the specified events. You can request bulk acknowledgements (that is, all acknowledgements for the current operation are written to a single file), or separate acknowledgements for each event. For example, you may want to specify individual acknowledgements when setting up and testing the UE, then bulk acknowledgements when you run ‘real’ extraction jobs.

Monitor Metadata Extraction Jobs

When an extraction job is running, errors, warnings and other messages are written to a log file. You can also specify that error codes and other information is displayed on-screen.

You can specify the logging level when you run an extraction job (see above). Log file names take the format: wgntask_<date>.log, where <date> is the date and time when the log file was created.

Extraction Job Parameters

For the XML metadata extractor, the job definition can include these categories of parameter:

- Job parameters (see page 479)
- Archive parameter (see page 480)
- Operation parameters (see page 480)
- Connection parameters (see page 482)
- File parameters (see page 482)
- Stored procedure parameters (see page 483)

Job Parameters

These parameters configure the basic job details.

module_name

This parameter is mandatory!

This parameter identifies the WgnTask.exe module. For UE metadata extraction jobs, it **must** be set to:

WgnUE_XML

guid

This parameter is mandatory!

This parameter identifies the WgnTask.exe module. For UE metadata extraction jobs, it **must** be set to:

521fdd03-d5c5-11d4-b613-000102027bbb

title

Use this parameter to define a job title. Although optional, we recommend that you supply a title. This title is incorporated into the file name of any resulting ‘Tasks’ log files.

description

Use this parameter to specify a brief description of the UE metadata extraction job. This description is written to the job definition document.

log_max_num_files

This specifies the maximum number of log files. When the maximum number of log files exists and the maximum size of the latest is reached (see below), the oldest log file is deleted to enable a new one to be created.

If this parameter is omitted, the job defaults to a maximum of 10 log files.

log_max_file_size

This specifies the maximum size (in KB) for each log file. When the current log file reaches its maximum size, the job creates a new log file.

If this parameter is omitted, the job defaults to a maximum file size of 1,000 KB (1 MB).

More information:

Log File Types (see page 504)

Archive Parameter

This parameter identifies the target archive and determines which events stored in the CMS are eligible for extraction.

RemotePrefix

This parameter is mandatory!

Use this parameter to identify the target archive and filter events on the CMS so that only those stored in the target archive are eligible for extraction.

This parameter specifies a prefix on the event RemoteID (this is an identifier used to associate events stored on the CMS with their actual message data held in a third party archive). If you do not specify a prefix, all events on the CMS are eligible for extraction. The default prefix is blank.

Operation Parameters

These parameters configure the extraction operation. For example, they specify the maximum size for files of extracted data, and the maximum number of events per file. They also specify whether data receipt acknowledgements from the target system are required.

DataIntegrity

Defaults to False. If set to True, this parameter specifies that the archive must supply acknowledgements that it has received extracted events.

BatchSize

Defaults to 0. This parameter specifies the maximum number of events for which XML Metadata is to be extracted in one request. It provides scalability, allowing a large value of MaxRecords without overloading the database resources. This parameter is independent of MaxFileSize.

BulkAcknowledgement

Defaults to True. This parameter specifies whether bulk acknowledgements are required. If set to True, all events extracted at a particular time are acknowledged in bulk. If False, events are acknowledged individually. This parameter can be overridden by the acknowledgement file.

MaxRecords

Defaults to 1000. This defines the maximum number of records (that is, events) that can be extracted in a single operation.

MaxFileSize

Defaults to 0. If this parameter is set to:

- A non-zero value, it defines the maximum size (in Kb) for the XML output files. When an output file reaches this limit, a new file is created. Each file name has a numeric suffix, indicating its sequential position in the series for the current operation.
- Zero, there is no file size limit and all extracted events are written to a single output file.

RetryPeriod

Defaults to 1440 (equivalent to 24 hours). This defines the period (in minutes) after which extracted events that have not been acknowledged are retried.

StartDate

Defaults to empty. This parameter specifies the ‘event timestamp’ date before which events will not be considered for export. If specified as an integer, it specifies a relative offset in days. For example, a value of -28 indicates a start date of 4 weeks ago.

EndDate

Defaults to empty. This parameter specifies the ‘event timestamp’ date after which events will not be considered for export. If specified as an integer, it specifies a relative offset in days. For example, a value of -7 indicates an end date of 1 week ago.

More information:

Archive Acknowledgements (see page 478)

Connection Parameters

These parameters are optional.

By default, the UE job uses single sign-on (see About single sign-on) to identify the connection details, but you can use these parameters to specify a server and CA DataMinder user.

Server

Specify the server name of the CMS you want to connect to for the UE job.

User

Specify the CA DataMinder user you want to connect to the CMS as.

This user must have the administrative privilege Admin: Disable management group security set to True in order to search for events outside of their management group.

Password

Identify the password for the user specified in the User parameter, above.

File Parameters

These parameters specify the output folders and file name prefixes for output and acknowledgement files.

AckPrefix

Defaults to ‘UEAck-’. This specifies the prefix for acknowledgement file names. This prefix enables the UE to identify acknowledgement files in a folder containing other files.

AckDirectory

Defaults to C:\temp\UE. This specifies the target folder for acknowledgement files.

ExportPrefix

Defaults is ‘UEUpdate-’. This specifies the prefix for XML output file names.

ExportDirectory

Defaults to C:\temp\UE. This specifies the target folder for XML output files.

Stored Procedure Parameters

These parameters specify the required database stored procedure (SP). It is unlikely that you will need to change these parameters. They are listed for testing purposes and to allow advanced UE configuration.

SPPackage

For Oracle CMSs only. This defines the package that the SP belong to. It defaults to WgnUE.

ExportSPName

Defaults to GetExportEventsAsXML. This specifies the SP that extracts the metadata as XML.

CommitSPName

Defaults to CommitExport. This specifies the SP that commits the extraction after the XML files have been successfully written.

AckSPName

Defaults to AcknowledgeExport. This specifies the SP used to acknowledge events individually.

BulkAckSPName

Defaults to BulkAcknowledgeExport. This specifies the SP used to acknowledge events in bulk.

Chapter 21: Mapping Events to Users

This section summarizes the various methods how CA DataMinder maps events to users.

This section contains the following topics:

Introduction

CA DataMinder needs to associate captured events with participants in order to allow reviewers to search for events by user. CA DataMinder also requires a mechanism to determine which user policy to apply to these events.

For emails and IM conversations, this mapping is typically based on email addresses. This guide therefore also describes how CA DataMinder maps email addresses onto user accounts, and lists which CA DataMinder features use address mapping. Finally, it includes some FAQs about address handling in CA DataMinder.

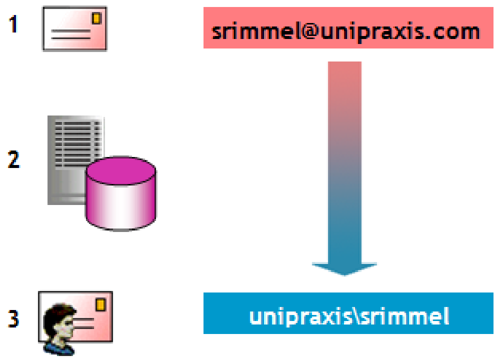

When mapping email addresses to user names, CA DataMinder intercepts or imports an email and extracts the addresses of the sender and recipients (1). Linked tables in the CMS database (2) enable CA DataMinder to map the email address of each participant to an existing CA DataMinder user (3).

Mapping Emails to Users

The method used to associate captured emails and IM conversations with users depends on the capture source.

Email Server Agents

CA DataMinder uses the following process to map emails captured by an Exchange or Domino server agent to users.

Searching by user

The sender and recipient addresses are all stored in the address table. If a reviewer subsequently searches for that email, CA DataMinder maps these addresses onto CA DataMinder users when the search runs.

Which user policy is applied?

The policy engine (PE) first determines whether the sender's email address matches an internal address pattern.

- If it does and the sender can be mapped onto an existing CA DataMinder user account, the PE applies that user's policy.
- If no matching CA DataMinder user can be found, or if the sender is deemed to be external, the PE applies a default policy (see page 488).

More information:

Default Policies for Emails From Unrecognized Senders (see page 488)

Email Client Agents

CA DataMinder uses the following process to map emails captured by Outlook or Notes client agents to users.

Searching by user

The sender and recipient addresses are stored in the address table. During subsequent searches for emails, these addresses are mapped onto CA DataMinder user accounts.

Which user policy is applied?

The client agent associates the user's Windows logon credentials with a matching CA DataMinder user account and applies that user's policy. The synchronization between Windows accounts and CA DataMinder accounts is created in the CMS database during an Account Import job or, less commonly, by using Microsoft Windows user authentication to automatically generate new user accounts.

Imported Emails

Emails can be imported onto the CMS as part of an Event Import job or Import Policy job.

Searching by user

The sender and recipient email addresses are stored in the address table. During subsequent searches for emails, these addresses are mapped onto CA DataMinder user accounts.

Which user policy is applied?

For Import Policy jobs, the policy engine maps the sender's address onto a CA DataMinder user account and applies that user's policy (that is, the sender's policy). If this mapping fails, the policy engine applies a default policy (see page 488).

More information:

Default Policies for Emails From Unrecognized Senders (see page 488)
NBA-Captured Emails (see page 487)

NBA-Captured Emails

For emails captured by or imported from the NBA, the mechanisms for assigning participants and applying policy is the always the same, regardless of whether the NBA is in active or passive mode.

Searching by user

The sender and recipient email addresses are stored as event participants.

Which user policy is applied?

First, the policy engine tries to map the sender's email address to a CA DataMinder user account. However, because the NBA typically targets messages sent from private Webmail accounts, in many cases it will not be possible to map the sender's address to a CA DataMinder account. In these cases, the policy applies a default policy (see page 488) to emails from unrecognized senders.

More information:

Default Policies for Emails From Unrecognized Senders (see page 488)
Imported Emails (see page 487)
Address Mapping Procedure (see page 497)

Default Policies for Emails From Unrecognized Senders

Policy engines (PEs) use settings in their machine policy to determine which policy to apply. Briefly, there are three key settings:

Internal Email Address Pattern

This setting specifies a list of internal email address patterns. When a PE receives an email for processing, it first compares the sender's address against these internal address patterns:

- Email is internal: If a match is confirmed, the email is deemed internal. The PE then checks whether the sender's address matches an existing CA DataMinder user's address. If it does, the PE applies that user's policy; if not, the PE applies the Unknown Internal Sender's policy (see below).
- Email is external: If the sender's address does not match an internal address pattern, the policy engine infers the email is external and applies the External Sender's policy (see below).

Unknown Internal Sender

Policy engines use this setting to apply policy to internal emails sent from unrecognized users within your organization. It defaults to 'UnknownInternalSender'; this user account is created automatically when you install a new CMS.

External Sender

Policy engines use this setting to apply policy to external emails. That is, emails sent from someone outside your organization. It defaults to 'ExternalSender'; this user account is created automatically when you install a new CMS.

For details on Unknown Internal Senders or External Senders see Configuring the Local Machine Policy (see page 343) section in the Policy Engines chapter of this guide

More information:

Email Server Agents (see page 486)
Imported Emails (see page 487)
NBA-Captured Emails (see page 487)
Mapping IM Conversations to Users (see page 489)
Configure the Local Machine Policy (see page 343)

Mapping Application Events to Users

Application events are captured by the Application endpoint agent.

Searching by user

The Application Monitor endpont agent maps the user's Windows logon details onto a CA DataMinder user account and stores that user as the event participant.

Which user policy is applied?

The Application Monitor endpont agent associates the user's Windows logon credentials with a matching CA DataMinder user account and apply that user's policy. The synchronization between Windows accounts and CA DataMinder accounts is created in the CMS database during an Account Import job or, less commonly, by using Microsoft Windows user authentication to automatically generate new user accounts.

Mapping IM Conversations to Users

IM conversations can be imported into the CMS embedded within EML files (Internet emails). The import process involves the following steps:

1. Use IMFrontEnd.exe to extract the IM data from dump files and generate CNV files.
2. Use the Cnv2EML.exe utility to extract the IM conversations from these CNV files and embed them within EML files.
3. Use Event Import to import the EML files.

The following methods are used to store IM participants and determine which policy is applied.

Searching by user

Providing that the participants' email addresses were available in the original IM dump file, the addresses for each participant in an IM conversation are stored as recipients in the EML email.

Which user policy is applied?

The Cnv2EML.exe parameter EML.From determines which policy is applied to the EML email. This parameter specifies an email address that the policy engine (PE) then compares against a list of internal address patterns. If the address is deemed internal and it can be mapped onto an existing CA DataMinder user account, the PE applies that user account's policy. If no matching CA DataMinder user can be found, or if the sender is deemed to be external, default policies (see page 488) handle emails from unrecognized senders.

More information:

Default Policies for Emails From Unrecognized Senders (see page 488)

Mapping File Events to Users

The method used to associate file events with CA DataMinder user accounts depends on the capture source.

Event Import

You can configure import jobs to associate imported files with specific CA DataMinder users.

Searching by user

To enable reviewers to search for imported files by user, this Event Import parameter associates a single event participant with each imported file. This means each file in this import is associated with the same participant.

```
ImpFile.AssociatedParticipant
```

If this parameter is omitted or cannot be resolved, the file has unspecified participants (see page 494) so CA DataMinder associates it with the source machine.

Which user policy is applied?

This Event Import parameter determines which policy is applied:

```
ImpFile.PolicyParticipant
```

If this parameter is omitted or cannot be resolved, the policy engine applies a default policy (see page 493).

Important: This parameter must be specified and the email address must exist for a File Import Policy job. If the parameter is omitted, all files will be excluded and the default policy is not applied.

In both cases, these parameters specify email addresses. Linked tables in the CMS database then enable CA DataMinderto map these email addresses (see page 497) onto existing CA DataMinder user accounts.

More information:

Searching for Files If No Participants Are Specified (see page 494)
Default Policies for File Events (see page 493)
Address Mapping Procedure (see page 497)

File Scanning Agent (FSA)

You can use the FSA Job Definition Wizard to associate scanned files with specific CA DataMinder users. Fill in the following fields in the second to last screen of the wizard.

Which user policy is applied?

This is determined by the following options in the wizard.

Use default policy for files

If this option is selected, the default policy for files is applied to scanned files.

Use policy of specified user

If this option is selected, the specified user policy is applied to scanned files.

If this scanning job element cannot be resolved, the policy engine applies a default policy.

Searching by user

(Optional) To enable reviewers to search for scanned files by user, you must associate a single event participant with each scanned file. This means, each file is associated with the same participant. This is determined by the following option in the wizard.

Associated Participant

This option defines a valid email address of a CA DataMinder user.

If this scanning job element cannot be resolved, the file has unspecified participants (see page 494) so CA DataMinder associates it with the source machine.

As with Event Import, linked tables in the CMS database then enable CA DataMinder to map these email addresses onto existing CA DataMinder user accounts.

More information:

Searching for Files If No Participants Are Specified (see page 494)
Default Policies for File Events (see page 493)
Address Mapping Procedure (see page 497)

Network Boundary Agent (NBA)

The NBA stores source and destination machine details when it captures files being sent across the Internet boundary. These files can include downloads, uploads, FTP transfers, and email attachments. The mechanism for associating file events with users depends on whether the NBA is running in active or passive mode.

Socket Output Mode

In socket output mode, the NBA outputs data to policy engines via a socket connection. When the NBA passes captured files to policy engines for processing:

Searching by user

For files captured by the NBA, the NBA passes the IP addresses of the source and destination machines to the policy engine. Both IP addresses are then stored as event participants.

To ensure that these file events are subsequently searchable by user in the iConsole or Data Management console, CA DataMinder administrators must add these machine addresses to users' address lists in the Administration console.

Which user policy is applied?

For all files captured by the NBA, the policy engine always applies the default policy (see page 493) for files.

Disk Output Mode

In disk output mode, the NBA outputs captured files to the local disk. These files are subsequently imported onto the CMS. You can configure Event Import to associate NBA-captured files with specific CA DataMinder user accounts:

Searching by a single user

To enable reviewers to search for imported files by user, this Event Import parameter associates a single event participant with each imported file (that is, the same participant is associated with each file):

```
ImpFile.AssociatedParticipant
```

If this parameter is omitted or cannot be resolved, CA DataMinder treats it as a file with unspecified participants (see page 494).

Searching by individual users

You can associate NBA-captured files with individual users. To do this, you can configure Event Import to extract the source and destination machine IP addresses from each imported file and save these IP addresses as event participants. The relevant import parameter is:

```
ImpFile.ParticipantsFromNBAFilename
```

To ensure that these file events are subsequently searchable by individual user in the iConsole or Data Management console, CA DataMinder administrators must add the these machine addresses to users' address lists in the Administration console.

Which user policy is applied?

This Event Import parameter determines which policy is applied:

```
ImpFile.PolicyParticipant
```

If this parameter is omitted or cannot be resolved, the policy engine applies the default policy (see page 493) for files.

In all cases, these parameters specify email addresses, or pseudo addresses in the case of ImpFile.ParticipantsFromNBAFilename. Linked tables in the CMS database then enable CA DataMinder to map these email addresses onto existing CA DataMinder user accounts.

More information:

Searching for Files If No Participants Are Specified (see page 494)
Default Policies for File Events (see page 493)
Address Mapping Procedure (see page 497)
Features That Use Email Address Mapping (see page 495)

Default Policies for File Events

Policy engines use the following setting in their machine policy to determine which policy to apply to scanned, captured or imported files if no other means are available to determine the policy participant. For example, this is necessary if an Import Policy job for FSA scanning job omits to specify the policy participant, or if the specified user account does not exist.

Default Policy For Files

This setting specifies a CA DataMinder user account. Policy engines use this setting to apply the specified user’s policy to file events. The setting defaults to 'DefaultFileUser'; this user account is created automatically when you install a new CMS.

Note: The user account specified for this setting is not stored as an event participant. For example, ‘DefaultFileUser’ is not stored as an event participant.

Important: If the policy participant is specified incorrectly and cannot be resolved for a File Import Policy job or an FSA job, policy cannot be applied to any imported or scanned files!

More information:

Event Import (see page 490)
File Scanning Agent (FSA) (see page 491)
Network Boundary Agent (NBA) (see page 492)

Searching for Files If No Participants Are Specified

Typically, a file import job or FSA scanning job explicitly specifies the event participant. But in addition, if Event Import or the FSA are importing or scanning local files, CA DataMinder automatically associates these file events with the source machine. This ensures that reviewers can search for these file events even if Row Level Security (RLS) is enabled. However, you must add the source machine ID to the address list of an appropriate CA DataMinder user.

How does RLS affect file searches?

RLS is implemented by default when you install a CMS database. RLS ensures that a reviewer can only see events associated with users in their management group when searching the CMS database. However, RLS also means all events must have a participant, otherwise reviewers cannot see them. But file events are problematic in this respect because, unlike emails or IM conversations, there may be no obvious user (that is, no obvious participant) associated with a file. This is particularly true of files scanned or imported from a network server rather than an employee's workstation.

Note: Reviewers can override RLS restrictions and search for any events if they have the 'Admin: Disable management group filtering' administrative privilege. See the Administration console help for details; search for 'privileges'.

Source machine ID stored as an event participant

If RLS is enabled, a file event must have a participant to be searchable. Therefore, as a fail-safe mechanism CA DataMinder automatically associates file events with the source machine if Event Import or the FSA are importing or scanning local files (but see the note below). This ensures that each file has a participant, even if the import job or scanning job omit to explicitly identify one. In turn, this allows reviewers to search for these files and determine which machine they originate from (providing these machine IDs have been associated with a CA DataMinder user account).

In technical terms, CA DataMinder automatically stores the source machine ID (this is the server hosting Event Import or the FSA) in the address table in the CMS database. For example:

```
cn=UX-MILAN-W2K3, cn=computers or cn=10.130.2.28, cn=computers
```

More information:

Event Import (see page 490)
File Scanning Agent (FSA) (see page 491)
Network Boundary Agent (NBA) (see page 492)

Mapping Email Addresses to Users

The following sections describe how CA DataMinder maps email addresses onto user accounts and the CA DataMinder features that use this mapping.

Features That Use Email Address Mapping

CA DataMinder needs to map email addresses onto individual users for 'multiple participant' events. Specifically, address mapping is used by the following CA DataMinder features to associate imported emails and IM conversations, and emails captured on an Exchange or Domino server, with specific CA DataMinder users:

Policy engines

Before a policy engine (PE) can apply policy triggers to an intercepted email, it needs to map the sender's email address to a CA DataMinder user. The mapping identifies the email owner and determines which policy to apply.

If the PE is unable to map an email address to an existing CA DataMinder user, the following machine policy settings determine which policy is applied:

- Retention Period for Unused Policies
- Unknown Internal Sender
- External Sender
- Internal Email Address Pattern

User attribute data lookup

Before CA DataMinder can evaluate control triggers based on user attribute (userattr) lookup commands, it must map the recipients of an outgoing e-mail (or the sender of an incoming email) onto CA DataMinder users. It can then evaluate the lookup command, comparing the attributes of the recipients (or the sender of an incoming email) against the test criteria.

If the lookup command is unable to map a recipient onto an existing CA DataMinder user, the command typically evaluates to False so the trigger does not activate. For full details about user attribute lookup commands, see the Administration console online help; search the index for 'data lookup, User Attribute lookup'.

Event Import

Unlike in previous versions of CA DataMinder, Event Import does not assign emails or IM conversations directly to owners. Instead, it identifies ‘event participants’ and associates an e-mail address with each participant. Under normal conditions, address mapping is not required while an import job is running. Instead, it is used subsequently to associate imported events with specific CA DataMinder users during event searches.

However, address mapping *is* used during an import job if a ‘user attribute’ filter is specified. This enables import jobs to exclude or only include all e-mail or IM conversations associated with CA DataMinder users who have specific account attributes.

Import Policy

Import Policy provides a mechanism for applying policy triggers to imported emails directly before they are stored in the CMS. For import policy jobs, address mapping is not used during the import phase. Instead, address mapping is used by the policy engines to determine which policy to apply.

More information;

Configure the Local Machine Policy (see page 343)

Address Mapping Procedure

To minimize storage and optimize performance, events and users are stored separately in the CMS database. The diagram below shows, in simplified form, how tables in the CMS database permit an email to be associated CA DataMinder users. In this example, an imported email sent from lsteel@unipraxis.com to srimmel@unipraxis.com is associated with two user accounts, UNIPRAXIS\lsteel and UNIPRAXIS\rimmel.

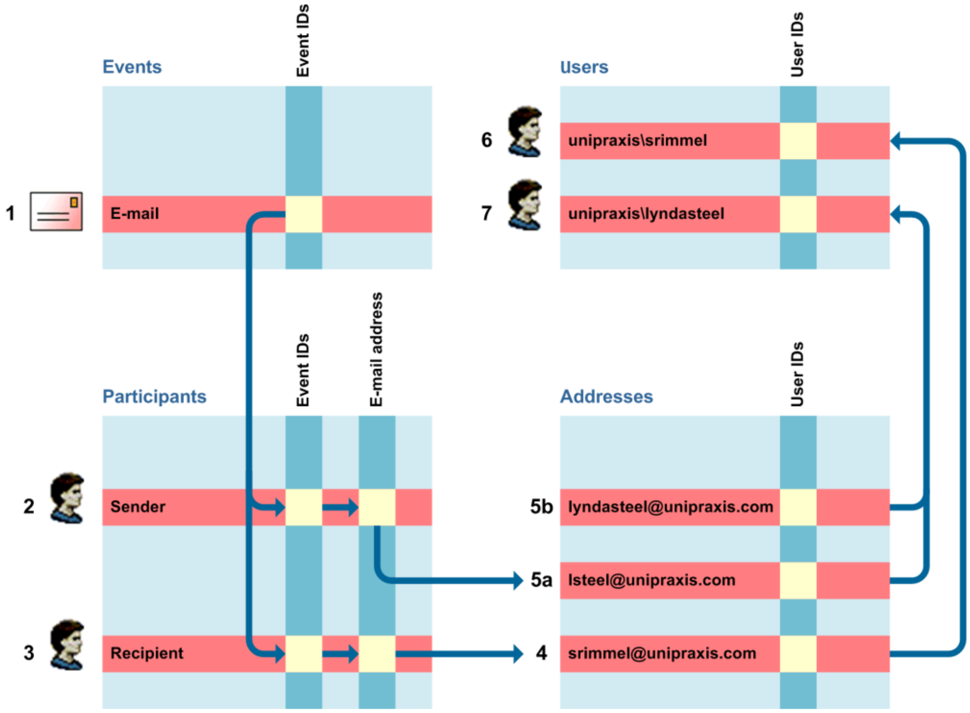

CMS database tables: Mapping events to CA DataMinder users

1 Rows in the events table of the CMS database contain individual events, each with a unique ID, linking events to participants. This example highlights an email event.

2 and **3** Rows in the participants table identify all event participants. Each participant can only be associated, via the Event ID, with a single event and a single email address. In this example, there are two email participants, a Sender and Recipient.

4 The addresses table lists all email and IM addresses known to CA DataMinder. In this example, the email recipient has the address srimmel@unipraxis.com.

5a The email sender has the address lsteel@unipraxis.com.

5b The addresses table lists all primary addresses and aliases. In this example, lyndasteel@unipraxis.com. is an alias for the primary address lsteel@unipraxis.com.

6 The users table associates CA DataMinder users, via a unique user ID, with one or more email addresses (or aliases). In this example, the user UNIPRAXIS\lsteel is associated with two addresses (**5a** and **5b**).

7 The user UNIPRAXIS\srimmel is associated via their user ID with the email address srimmel@unipraxis.com.

Chapter 22: Backing Up and Restoring the CMS

We recommend that you make a full back up of your CA DataMinder database on the CMS at least once per week, and incremental backups on a daily basis. This section gives an overview of the backup and restore procedures for the CMS, for both Microsoft SQL Server and Oracle database engines.

Regardless of which database engine you use, the backup procedure involves two general tasks, in addition to various database-specific tasks.

This section contains the following topics:

Database Backup Tasks

For details about backing up your SQL Server or Oracle database, see the *Database Guide*.

General Backup Tasks

When you back up the CMS, you must perform these backup tasks irrespective of the type of database engine used on the CMS.

Back Up the Data Folder

The Data folder holds all the configuration data and captured data used by your CA DataMinder installation.

You need to incorporate this folder into your existing backup regime. We recommend you back up this folder a minimum of once a week, but a daily backup is preferential. By default, when you install CA DataMinder this folder is added as a Data subfolder in the installation folder, but you can rename it and locate it anywhere suitable on your network.

There are two subfolders inside the Data folder which require special handling in the backup:

- **cache:** This subfolder stores cache files of all processed data. It does **not** need to be backed up.
- **e:** This subfolder stores all captured data. We recommend that you carry out a partial back up of this subfolder—see below.

The e subfolder needs special handling because of its potentially large size. You can either back up the entire folder, or back up selected events by date range using the dated folder names. Typically, you may only need to back up BLOB files associated with events from the last 2 to 7 days.

It is good practice to have a separate strategy in place for backing up BLOB files associated with events more than a week old. For example, to complete a full backup of the e subfolder once a month. Assuming that the daily backup collects data from the last seven days, then a monthly backup means that no more than 1 full restore and at most 4 daily backups are required to restore the event folder to any day within the last month.

Back Up the Master Encryption Key

The CMS uses a password-protected key to provide highly secure data management. The master encryption key is stored in the registry.

The master encryption key is stored in the registry. If you need to restore the CMS, you will need to restore the key. For this purpose, CA DataMinder provides a data management utility for exporting and re-importing these keys.

To back up the master encryption key

After installing your CMS, run the data management utility, wgnmgmt.exe, on the CMS machine to export a password-protected file containing the necessary registry details. Find this utility in the \Support folder on your CA DataMinder distribution media. The command line syntax is:

```
wgnmgmt e <file name> <password>
```

Restoring a CMS

These instructions describe how to restore your CMS to a point-in-time. For example, if your CMS host server suffers a hardware failure then you will need to restore the CMS.

1. **Stop the CA DataMinder infrastructure**

   Using the Computer Management utility in Windows, expand the Services group and stop the service ‘CA DataMinder Infrastructure’.

   Or you can stop the infrastructure service from a command line:
   ```
   net stop wgninfra
   ```

2. **Reimport the data management registry key**

   This data management registry key was used by the original installation. Run wgnmgmt.exe on the CMS host serverm, using this command syntax:
   ```
   wgnmgmt i <file name> <password>
   ```

   Where <file name> and <password> are the backup file and password you specified when you backed up the master encryption key.

3. **Restore the CA DataMinder data folder**

   Using your normal data-recovery procedures, restore the CA DataMinder data folder to your CMS host server. You can restore it to any suitable location on your network, and give it any name. You simply specify the name and location when you reinstall CA DataMinder (step 5).

4. **Restore the CMS database**

   The procedure depends on your Database Engine:

   - For Oracle databases, create a recovery job for your CA DataMinder database using the Recovery wizard. See your Oracle documentation for details.
   - For SQL Server databases, restore your CA DataMinder database using the SQL Server Database Restore feature. For details, see your SQL Server documentation.

   Note: If you are restoring after a complete system failure, recreate a login for CA DataMinder to use. See the *Database Guide* for details; search for ‘Oracle accounts’ and ‘SQL Server login properties’.

5. **Reinstall CA DataMinder on the CMS server**

   Use the CA DataMinder installation wizard:

   - In the Database Type screen, enter the name and password for the login that CA DataMinder uses to manage the CMS database.
   - In the Data Folder screen, specify the name and location of the data folder that you previously recovered in step 3.

More information:

Back Up the Master Encryption Key (see page 501)

Chapter 23: Log Files

This section introduces the CA DataMinder log files. These are maintained for all product components to record significant activity or events. It provides a summary of all supported log types, including an overview of Policy Incident logs. It also identifies those logs generated by the CA DataMinder infrastructure and logs generated directly by other CA DataMinder components or utilities.

This section also describes how to configure CA DataMinder to write log entries to external logs, including the Windows Application log and Syslog servers such as ArcSight.

This section contains the following topics:

About Log Files

Infrastructure-based logs can be viewed in the Administration console and are configured in the local machine policy. Non-infrastructure logs are typically configured by editing the registry or a .ini configuration file. Certain non-infrastructure logs are also viewable in the Administration console.

The table below summarizes the log files supported by CA DataMinder. See the following sections for details.

Log type	**File name**	**Infrastructure? ***
Archive integration	wgnemcs1_*.log wgnsev_*.log zdsretrieval_*.log	
Account Import	ldap_*.log	Yes
Activity	activity_*.log	Yes
Command Output	command_*.log	Yes
Event Import	evtimport_*.log	
File Scanning Agent	wgnfsa_*.log	

Log type	File name	Infrastructure? *
iConsole	iconsole_*.log	
Policy Engine Hub	wgnphub_*.log	
Policy Incidents	policyincident_*.log	Yes
Quarantine Manager	wgnqmgr_*.log	
Replication	repl_*.log	Yes
Socket API	wgnsagent_*.log	
System	stderr_*.log	Yes
Tasks	task_*.log	
User administration	useradmin_*.log	Yes

** Log files maintained by CA DataMinder infrastructure*

More information:

Viewing Log Files (see page 509)
Configure Log Files (see page 511)

Log File Names

Log file names indicate the type of log, and incorporate the date and time when the file was created. For example, activity_200903170945.log is an Activity log created on 17 March 2009 09:45.

Log File Types

CA DataMinder agents and utilities support the following types of log file. Logs maintained by the CA DataMinder infrastructure are identified accordingly.

Archive integration logs

These logs are not infrastructure-based. They record the progress of message processing operations by archive integration agents. Log file names take this format:

EMC SourceOne integration: wgnemcs1_<date>.log

Symantec Enterprise Vault integration: wgnsev_<date>.log

Zantaz Digital Safe integration: zdsretrieval_<date>.log

Account Import logs

These are infrastructure-based logs. They record the outcome of any operations using Account Import. Log entries typically include changes to the user or machine hierarchy, such as the addition of new users, groups or client machines.

Log file names take the format: ldap_<date>.log.

Activity logs

These are infrastructure-based logs. They record general activity by all machines. For example, each time users or machines log in or out, and each time policies are created or updated.

Log file names take the format: activity_<date>.log.

You can also record user account changes in the User Administration log.

Command Output logs

When running a wgninfra -exec command, for example to carry out a policy integrity check, the returned status messages are written to a Command Output log file.

Log file names take the format: command_<date>.log.

Event Import logs

These logs are **not** infrastructure-based, but are viewable in the Administration console.

They record the outcome of Event Import operations, including import failures and any system errors (for example, when a user cannot be created).

Log file names take the format :evtimport_<instance>_<date>.log

Where <instance> identifies the service instance associated with the Event Import job.

File Scanning Agent logs

These logs are **not** infrastructure-based, but are viewable in the Administration console.

They record the outcome of FSA scanning jobs, such as details of replaced files, connections to the scan database, and when jobs started and completed. FSA log files are saved on the machine hosting the FSA.

Log file names take the format: wgnfsa_<date>.log.

iConsole logs

These logs are **not** infrastructure-based, but are viewable in the Administration console.

They record the outcome of iConsole operations, including details of any errors.

Log file names take the format: iconsole_<date>.log.

Quarantine Manager logs

There are two types of QM log messages.

The Quarantine Manager (QM) writes messages to the Activity log and also to its own log saved on the QM host machine. This QM log is *not* infrastructure-based.

Log file names take the format: wgnqmgr_<date>.log.

Activity log messages generally record the outcome of quarantine operations, while the QM's own log provides more diagnostic details.

Policy Engine Hub logs

These logs are **not** infrastructure-based.

They detail progress as each event is processed by the PE hub. The log file location is configurable but defaults to the same folder as the hub executable, wgnphub.exe.

Log file names take the format: wgnphub_<date>.log.

Policy Incidents logs

These are infrastructure-based logs. They record the outcome of user policy processing. Each time a policy incident is replicated to the CMS, an entry is written to a log file. Log entries identify the associated user and include a URL to view the incident in the iConsole.

Log file names take the format: policyincident_<date>.log.

You can configure this log file using machine policy settings.

Replication logs

These are infrastructure-based logs. They record any database changes that were made on a remote machine and copied to the local machine. These typically include captured data objects, changes to a machine or user policy, and changes to user accounts and user groups. These changes are recorded in the replication log on each machine.

Log file names take the format: repl_<date>.log.

You can configure this log file using machine policy settings.

Socket API logs

These logs are **not** infrastructure-based.

These record the processing results for messages passed to the Socket API.

Log file names take the format: WgnSAgent_<date>.log

System logs

These are infrastructure-based logs. They record any infrastructure errors that occur while the CA DataMinder service is running. Under normal conditions, this log file is empty.

Log file names take the format: stderr_<date>.log.

You can configure this log file using machine policy settings.

Note: Any errors detected when the CA DataMinder service starts up are written to the file wgninfra.out. This file is in the \data\log subfolder with the conventional log files.

Tasks logs

These logs are *not* infrastructure-based, but are viewable in the Administration console. They record the outcome of jobs run using the WgnTask.exe utility.

Log file names take the format: task_<title>_<date>.log

Where <title> is an optional identifier based on the title parameter in the associated job definition document. For example, a task based on the Universal Extractor's XML metadata extractor may generate log files in this format:

```
task_Extracted XML metadata_200903170945.log
```

User Administration logs

These are infrastructure-based logs. They record any changes made to user accounts or groups. These typically include changes to user accounts and user groups.

Log file names take the format: useradmin_<date>.log.

You can configure this log file using machine policy settings.

More information:

Configure Log Files (see page 511)
iConsole Deployment (see page 105)
Quarantine Manager (see page 225)
PE Hub Log Files (see page 383)
About Policy Incident Logs (see page 508)

About Policy Incident Logs

These logs are created on the CMS only. They record the outcome each time a user policy trigger fires. The log includes both event-level and trigger-level entries.

A key feature of these log entries is that they include an event URL to display the incident in the iConsole. Administrators can use this URL to view the incident (for example, a captured email) in the iConsole, plus any attachments and a summary of the policy that was applied.

Event-level Messages

Only one of these messages is logged for each event, regardless of how many triggers the event causes to fire. They are structured as follows:

```
<Associated user>
<Message ID>
<User action>
<Policy outcome>
<Event severity>
<Machine name>
<Event ID>
<Event URL>
```

where:

<Associated user>

Is the primary participant of an event, for example, the sender of an outgoing e-mail.

For details about how CA DataMinder assigns participants to files events (such as files scanned by the FSA or captured by the NBA), see the Event Participants technical note, available from CA Technical Support.

<Message ID>

Is a code that identifies the message type (event-level or trigger-level) and severity.

<User action>

Describes what the user did (for example, ‘The user sent an email’) or the event type (such as ‘Scanned file’) .

<Policy outcome>

Summarizes the outcome of policy processing. For example, CA DataMinder blocked the email or warned the sender.

<Event severity>

Indicates which severity band the event is assigned to (Low, Medium or High).

<Machine name>

Indicates the source machine. For example, this could be the machine from which an email was sent.

<Event ID>

Uniquely identifies a captured or imported event in the CMS database.

<Event URL>

Provides a URL to display the event in the iConsole. Users can browse to this URL to view the event in the iConsole.

Trigger-level Messages

For each trigger that fires, a log message records these details:

```
<Associated user>
<Message ID>
<Trigger name>
<Event severity>
<Event ID>
<Event URL>
```

where:

<Trigger name>

Identifies which policy trigger activated.

Other message details are the same as for error-level messages.

Viewing Log Files

Log files are typically saved locally (see below). If they are saved on the CMS or any other machine running the CA DataMinder infrastructure, you can view these log files directly in the Administration console. You can also view certain non-infrastructure logs in the Administration console.

Where Log Files Are Saved

The log file location depends on the log type:

Which Machine?

> Log files are typically saved locally. For example, FSA log files are saved on the machine hosting the FSA; Account Import log files are saved on the CMS; and System log files are saved on the machine where an infrastructure error occurred. Note that log files stored on non-CMS machines are *not* replicated up to the CMS.

Which Folder?

> In the current CA DataMinder release, log files are typically saved in CA's \data\log subfolder. On 32-bit machines, find this subfolder in the Windows All Users profile. On 64-bit machines, find this subfolder below the \ProgramData folder.
>
> The relevant sections in this Deployment guide include log file details for individual utilities or agents.

Examples:

Windows 2003, Windows XP

> C:\Documents and Settings\All Users
> \Application Data\CA\CA DataMinder\data\log

Windows Vista, Windows 2008, Windows 7 and higher

> C:\ProgramData\CA\CA DataMinder\data\log

View Log Files in the Administration Console

The Administration console enables you to view logs on any machine in your enterprise running the CA DataMinder infrastructure, plus certain other logs on machines not running the infrastructure (these include Event Import, iConsole and FSA logs).

Note: To view log files in the Administration console, you must have the Machine: View Logfile privilege.

To view log files on the CMS

1. In the Administration console, connect to a CMS.
2. Choose Manage, Logfiles.
3. Browse the available log files and choose the one you want to view:

   🗎 indicates the current log file.

   ▤ indicates a closed log file.

To view log files on other machines

1. In the Administration console, connect to a CMS.
2. Choose Manage, Machine Administration.
3. Expand the machine hierarchy and select the machine you want.
4. Right-click and choose View Logfile.

Configure Log Files

Log files generated by the CA DataMinder infrastructure can be configured in the local machine policy. For some logs, you can also choose which types of event are logged. Log files that are not created by the infrastructure (such as iConsole and Event Import logs) are configured using an alternative mechanism.

Configuration for Non-infrastructure Logs

CA DataMinder log files that are not created by the infrastructure can be configured by, for example, editing the registry or an .ini configuration file.

More information:

PE Hub Log Files (see page 383)
Configure iConsole Log Files (see page 150)
Quarantine Manager Log Files (see page 237)

General Log Configuration in Machine Policy

Log files controlled or generated by the CA DataMinder infrastructure are configured in the relevant machine policy. For example, to modify logging settings for all gateway servers, you must edit the Common Gateway Policy.

Policy settings specify the maximum number and size of log files and whether log entries are copied to external logs. You need to edit the following settings in the Infrastructure, Logging policy folder:

Maximum Number of Log Files

Defaults to 10. The maximum number applies separately to each type of log file (that is, for each type of log file you can generate files up to the maximum number). When this maximum number is exceeded, the oldest log file is deleted.

Maximum Size of log Files

Defaults to 1,000. This specifies the maximum size (in KB) for each log file. When the current log file reaches its maximum size, CA DataMinder creates a new log file.

This size limit does not apply to Account Import log files.

Write to Windows Event Log

This specifies the default handling for copying CA DataMinder log entries to the Windows Application log. For example, you can choose to copy all messages or to only copy errors and warnings.

You can override the default handling and set custom log levels for different types of CA DataMinder log file.

Write to Syslog Server

This specifies the default handling for copying CA DataMinder log entries to a Syslog server such as ArcSight. As with the Windows Application log, you can choose which messages to copy. You can also override the default handling and set custom log levels for different types of CA DataMinder log file.

More information:

Write to the Windows Event Log (see page 516)
Write to Syslog Servers (see page 519)

Policy Configuration for Specific Log Types

Further configuration is possible for the following logs.

Activity Logs

In the Infrastructure, Logging, Activity policy folder, you can optionally edit policy settings to record:

Machine and User Logins

For user logins, you can also record summary information (login and logout times) or detailed information (logins, logouts and failed account creation attempts).

Machine and User Administration Changes

These changes include accounts being created, modified or deleted. You can also record this information in the User Administration logs.

Policy Changes

These include any saved changes to user or machine policies.

Cache Actions

These include cache changing and purge activity involving the temporary object store (an event cache), plus cache preload and clear down operations associated with the internal user and email address caches (used to optimize event processing).

Storage Connector Events

These refer to events associated with third party object stores. CA DataMinder uses storage connectors to integrate with third party object storage solutions.

The range of logged events includes object store connections and disconnections, storage queue activity and data retrieval.

More information:

User Administration Logs (see page 515)

Policy Incident Logs

Only available in the CMS machine policy.

These logs are created on the CMS only. In the Infrastructure, Logging, Policy Incidents policy folder, you can optionally edit settings to record the outcome each time a user policy trigger fires. The log includes both event-level and trigger-level entries. To configure these logs, you need to edit these settings:

Log Policy Incidents Locally To File

If enabled, CA DataMinder writes entries to the policy incident log on the CMS.

iConsole Address

Set this to the fully qualified domain name or IP address of a machine hosting the iConsole application server. For example:

```
CMS-HARDY.unipraxis.com
```

This address is incorporated into the event URL in each log entry (see below). Administrators can browse to this URL to view the associated incident in an iConsole.

Format of Event URL

You do not normally need to edit this policy setting. It specifies the format of the event URL included in the log message. Users can browse to this URL to view the event in the iConsole. The default URL includes variables that are replaced with actual values when a message is written to the log. For example, variable {0} is substituted with the machine specified by the Address of iConsole policy setting—see above.

More information:

About Policy Incident Logs (see page 508)

Replication Logs

In the Infrastructure, Logging, Replication policy folder, you can optionally edit settings to record when a CA DataMinder machine *receives* the following replicated data:

Captured Data

This can include captured or imported e-mails, Web pages, IM events, file events, and so on, plus any associated trigger details, event attributes, and policy actions.

A single captured event typically comprises several objects, for example, an event object, a trigger object, a 'capture action' object, and several 'event attribute' objects.

Infrastructure Data

This includes changes to user and machine accounts or policy updates.

Critical Data

This typically includes data recorded for diagnostic purposes. For example, when a machine becomes suspended a notification is replicated to the CMS.

In all cases, you can choose to record summary information (when replication starts and completes, and how many objects were successfully replicated) or you can record detailed information (log entries for each successfully replicated object).

System Logs

In the Infrastructure, Logging, System policy folder, you can optionally edit a setting to record any infrastructure errors while the CA DataMinder service is running. Under normal conditions, this log is empty.

Note: Any errors detected when the CA DataMinder service starts up are written to wgninfra.out. Find this file in the same folder as the System log files.

User Administration Logs

In the Infrastructure, Logging, User Administration policy folder, you can optionally edit a setting to record changes made to user accounts or groups. You can also record these changes in the Activity logs.

More information:

Activity Logs (see page 513)

Write to the Windows Event Log

CA DataMinder can copy log entries to the local Windows Application log (accessible through the Windows Event Viewer). This enables you to use third party monitoring and alerting software such as Microsoft Operations Manager (MOM) to, for example, notify your administrators when a CA DataMinder error occurs by forwarding the error message to pagers or sending an e-mail alert. To set this up, CA DataMinder provides two mechanisms:

- For log files maintained by the CA DataMinder infrastructure, you need to edit the local machine policy. See below for details.
- For ‘non-infrastructure’ log files, you need to edit the local registry.

More information:

Registry Configuration (see page 517)

Machine Policy Configuration

For infrastructure-maintained log files, you can edit the machine policy to ensure that CA DataMinder log entries are copied to the local Windows Application log. You must set a default log level, and you can also specify a custom log level for individual logs.

Default Log Level

The Write to Windows Event Log setting determines which log messages are copied to the Windows Application log. Find this setting in the Infrastructure > Logging folder.

You can set this to ‘None’ so that no messages are copied to the Windows log, or you can choose ‘Errors Only’, ‘Errors and Warnings’, or ‘All Messages’. But note that you can also specify a custom log level for each individual log; this overrides the default log level—see below.

More information:

General Log Configuration in Machine Policy (see page 512)

Custom Log Level for Individual Logs

For each log maintained by the CA DataMinder infrastructure, a Log Detail setting lets you set a custom log level. Find these settings in the Infrastructure, Logging, External Logging, Windows Event Log policy folder.

Each Log Detail setting specifies which messages are copied to the Windows log and overrides the default log level. In each case, you can choose ‘None’, ‘Errors Only’, or ‘Errors and Warnings'. You can also choose:

- ‘Use Default’ to use the default log level defined by the Write to Windows Event Log setting.
- ‘All Messages’. Any message written to the CA DataMinder log is also copied to the Windows log.

Registry Configuration

For CA DataMinder log files that are *not* maintained by the infrastructure (for example, the iConsole, Event Import and Universal Adapter logs) you must edit the registry if you want to copy CA DataMinder log entries to the local Windows Application log.

Specifically, you need modify the following registry key:

```
HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA DataMinder
    \CurrentVersion\Logging
```

Within this registry key, edit the following value:

EventLogLevel

Type: REG_DWORD

Data: Defaults to 0. This determines whether or not CA DataMinder log entries are copied to the local Windows Application log. The default value of zero disables Windows logging. That is, log entries are *not* copied to the Windows log. If you specify a non-zero value, entries *are* copied to the Windows log. The supported logging levels are:

0

Disable Windows logging. But see the iConsole 'lost connection' note below.

1

Errors only

2

Errors and warnings

3

As 2, plus informational and status messages

Lost connection between iConsole servers

If a connection failure occurs between the iConsole application server and front-end Web server, CA DataMinder log entries are automatically written to the local Windows log file on the front-end Web server, regardless of how EventLogLevel is configured. This precaution is vital for diagnosing such connection failures.

More information:

Log Level Restriction (see page 518)

Log Level Restriction

Be aware that the Windows log level is limited by the level already defined for the CA DataMinder log file. You cannot write more information to the Windows log than is written to the CA DataMinder log.

For example, if the iConsole LogLevel registry value is set to 1, only errors are written to the CA DataMinder log file. This means the iConsole can also only write errors to the Windows log (that is, EventLogLevel is effectively limited to 1).

Write to Syslog Servers

If required, CA DataMinder can send log messages to Syslog servers such as CA Enterprise Log Manager or ArcSight. To enable this, you need to edit the local machine policy to specify the message format, log level, and your Syslog server.

You can specify up to three different Syslog servers in the local machine policy. That is, you can simultaneously copy CA DataMinder log messages to three different Syslog servers.

Note: This guide uses the term 'Syslog server' to refer to the Syslog receiver, also commonly known as a syslogd or syslog daemon.

Machine Policy Configuration

For infrastructure-generated log files, you can edit the machine policy to ensure that CA DataMinder log entries are copied to a Syslog server. You must set a default log level, and you can also specify a custom log level for individual logs (see below). You must also configure the Syslog connection and, if required, edit the CEF message format.

More information:

Syslog Configuration (see page 521)

Default Log Level

The Write to Syslog Server setting determines which log messages are copied to Syslog servers. Find this setting in the Infrastructure, Logging folder.

You can set this to 'None', so that no messages are copied, or you can choose 'Errors Only', 'Errors and Warnings', or 'All Messages'. But note that you can also specify a custom log level for each individual log; this overrides the default log level—see below.

More information:

General Log Configuration in Machine Policy (see page 512)

Custom Log Level for Individual Logs

For each log maintained by the CA DataMinder infrastructure, a Log Detail setting lets you set a custom log level. This specifies which messages are copied to the Syslog server and overrides the default log level. In each case, you can choose ‘None’, ‘Errors Only’, or ‘Errors and Warnings'. You can also choose:

- ‘Use Default’ to use the default log level defined by the Write to Syslog Server setting.
- ‘All Messages’. Any message written to the CA DataMinder log is also copied to the Syslog server.

Find these settings in the Infrastructure, Logging, External Logging, Syslog *n* policy folder.

Syslog Configuration

For each Syslog server, you must specify the following settings. Find these in the Infrastructure, Logging, External Logging, Syslog n policy folders.

Server Name

Enter the IP address or fully qualified domain name of the Syslog server.

Server Port

Specify the port number that the Syslog listens on. By default, Syslog servers use port 514.

Maximum Message Length

Specifies the maximum length (in characters) for log messages copied to a Syslog server. The Syslog protocol defines a maximum length of 1024 characters, but many Syslog servers can accept longer messages.

Client Port

Specifies the port(s) that CA DataMinder uses to send log messages to Syslog server. If required, you can specify a range of consecutive port numbers (such as 510—515) or a comma-separated list of port numbers and ranges (such as 501,505,510—515).

Syslog Protocol

Specifies the format for data transfers to the Syslog server. Choose either:

IETF RFC 3164

All Syslog servers support this protocol.

IETF Syslog Internet Draft Document

Specifies an extension to the RFC 3164 protocol.

We recommend that you choose the RFC 3164 protocol unless you are certain that your Syslog server supports the extension published in the Internet Draft Document.

Message Format: Choose either:

Common Event Format

Choose this option if your Syslog server supports CEF. For example, ArcSight uses CEF. If you do choose CEF, some further policy configuration is needed; see the next section.

Unformatted Data

If your Syslog server does not support CEF, choose this option.

Common Event Format Configuration

CEF messages include an event severity value, between 0 and 10. If you specify 'Common Event Format' as the Message Format (see previous section), you can optionally change the default severity values assigned by CA DataMinder. To do this, edit the following settings. Find these in the Infrastructure > Logging > External Logging > Syslog *n* policy folders:

Error Messages Severity Value

> Defaults to 8. This severity value is assigned to error messages and high severity events when sent to Syslog servers as CEF messages.

Warning Messages Severity Value

> Defaults to 5. This severity value is assigned to warning messages and medium severity events when sent to Syslog servers as CEF messages.

Information Messages Severity Value

> Defaults to 1. This severity value is assigned to Information messages and low severity events when sent to Syslog servers as CEF messages.

Note: Policy incident log messages classify events as Low, Medium, or High severity.

More information:

About Policy Incident Logs (see page 508)

Chapter 24: Technical Information

This section contains technical information that can affect CA DataMinder deployment. This includes supplementary information for various deployment and uninstallation methods (including administrative installations); TCP port configuration instructions, and a list of known issues. The full list is summarized below.

This section contains the following topics:

Performing an Administrative Installation

To deploy CA DataMinder to client machines using command-line methods or managed methods such as Group Policy, we recommend that you first perform an administrative installation to your network. This enables client machines to install CA DataMinder directly from the network without generating excessive network traffic or requiring excessive free disk space on the client.

The administrative installation installs a source image of CA DataMinder onto the network in a target folder that you specify. The source image is called Client.msi. Store this source image in a network folder that all the target client machines can access.

To perform an administrative installation

Use the /a command-line option for Msiexec.exe.

- The syntax to create an administrative image for 32-bit client machines is:

  ```
  msiexec /a <Path>\client.msi
  ```

- The syntax to create an administrative image for 64-bit client machines is:

  ```
  msiexec /a <Path>\client_x64.msi
  ```

Note: See your Microsoft documentation for details about administrative installations.

More information:

Command Line Parameters for Msiexec.exe (see page 531)

Installation Transforms

Transform files let you apply customized configuration changes to a Windows Installer package. CA DataMinder typically provides scripts to create several useful transforms. Find the scripts to generate transforms in the \Support folder of your CA DataMinder distribution media.

For script-generated transforms, you must copy the script into the same folder as your administrative installation source image, Client.msi or Client_x64.msi. You need write access to this folder. You run the script by double-clicking it in Windows Explorer. The script generates the transform for the administrative installation source image.

Note: If no script is provided, you can obtain the transform from CA Support at http://ca.com/support.

.

You can reference transforms as an option for msiexec.exe, as part of the program for an SMS distribution package, or as a modification to a Group Policy installation package.

The scripts generate appropriate transforms depending on your environment: For example, the EnableAppmon.vbs script creates the transforms EnableAppmon_Client.mst for 32-bit client installations, and EnableAppmon_Client_x64.mst for 64-bit client installations.

The references in this section list the transforms that are currently supported.

Example

You want to capture a typical deployment and then roll it out to all clients. You know that CA DataMinder starts automatically after installation and want to disable AutoStart during the model deployment. You run the DisableAutoStart.vbs script to generate the DisableAutoStart_Client.mst (see page 527) transform for the installer.

Follow these steps:

1. Perform an administrative installation of the client package. You can do this in a *path* of your choice.

   ```
   msiexec /a path_to_distrib\windows\client.msi
   ```

2. Copy the DisableAutoStart.vbs script into the *path* that contains the administrative installation of the client package.

   ```
   copy path_to_distrib\support\DisableAutoStart.vbs path
   ```

3. Run the VBS script.

   ```
   cd path
   DisableAutoStart.vbs
   ```

   The script produces the transform DisableAutoStart_Client.mst.

4. Install the client and deploy the DisableAutoStart_Client.mst transform.

   ```
   msiexec /i path\client.msi TRANSFORMS=path\DisableAutoStart_Client.mst
   ```

More information:

Identify the CMS: SetParentName_Client.mst (see page 526)
Prevent Unauthorized Uninstallations: ClientLockDown_Client.mst (see page 526)
Enable Silent Uninstallations for SMS: SMSQuietUninstall_Client.mst (see page 527)
Prevent Automatic Startup: DisableAutoStart_Client.mst (see page 527)
Configure Outlook Endpoint Agent: EmailClientOptions.mst (see page 528)
Prevent Consoles being Installed: HideConsole_Client.mst (see page 530)
Install Application Integration: EnableAppmon_Client.mst (see page 531)
VBS Scripts for Generating Installation Transforms (see page 612)

Identify the CMS: SetParentName_Client.mst

For a command line or Group Policy installation, you can use a transform to identify the name or IP address of the CMS or gateway that the client machines connect to.

Note: If you specify the CMS or gateway by name, you must ensure that the client machines can resolve its name.

1. Find the CreateParentNameTransform.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Copy the script into the folder containing your administrative installation source image.
3. Run the script and enter the name or IP address of the CMS or gateway when it prompts you for it.

   The script creates the SetParentName_Client.mst (or SetParentName_Client_x64.mst) transform.

Prevent Unauthorized Uninstallations: ClientLockDown_Client.mst

For command line, Group Policy or SMS installations, you can use a transform to prevent users from uninstalling CA DataMinder with the Add/Remove Programs utility. The ClientLockDown.mst transform disables the Change and Remove buttons when a user selects CA DataMinder in the Add/Remove Programs dialog.

Follow these steps:

1. Find the ClientLockDown.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Run the script.

   It creates the ClientLockDown_Client.mst (or ClientLockDown_Client_x64.mst) transform.
3. Copy the transform into the folder containing your administrative installation source image.

Enable Silent Uninstallations for SMS: SMSQuietUninstall_Client.mst

For SMS installations, you can include a transform that enables silent uninstallations. That is, the transform removes the need for user cooperation when uninstalling.

1. Find the SMSQuietUninstall.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Run the script.

   It creates the SMSQuietUninstall_Client.mst (or SMSQuietUninstall_Client_x64.mst) transform.
3. Copy the transform into the folder containing your administrative installation source image.

Prevent Automatic Startup: DisableAutoStart_Client.mst

By default, CA DataMinder starts automatically immediately after installation on servers and client machines. But for command line, Group Policy or SMS installations, you can use the DisableAutostart_*Package*.mst transform to prevent this.

1. Find the DisableAutostart.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Run the DisableAutostart.vbs script. The script automatically detects which administrative installations are present in the local folder and generates a transform for each package:

   The script creates one or more DisableAutostart_*Package*.mst transforms.

   - DisableAutoStart_Client.mst
   - DisableAutoStart_Client_x64.mst
   - DisableAutoStart_Integration.mst
   - DisableAutoStart_Integration_x64.mst
   - DisableAutoStart_Server.mst
   - DisableAutoStart_Server_x64.mst
3. Copy the transforms into the folder containing your administrative installation source image.
4. Reference the transform as an option for Msiexec.exe, as part of the program for an SMS distribution package, or as a modification to a Group Policy installation package.

Configure Outlook Endpoint Agent: EmailClientOptions.mst

For command line, Group Policy or SMS installations, you can create registry values to control the behavior of the Outlook client agent. Typically you use these settings in order to alleviate email processing delays.

Follow these steps:

1. Find the EmailClientOptions.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Run the script.

   Note: For command syntax details, see the Support Tools (see page 611) section. You can configure any or all of the registry values described above.

   It creates the EmailClientOptions_Client.mst (or EmailClientOptions_Client_x64.mst) transform.
3. Copy the transform into the folder containing your administrative installation source image.

The transform creates new values in the following registry key:

```
HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA DataMinder\CurrentVersion\EMail
```

Or for 32-bit Outlook client running on a 64-bit OS:

```
HKEY_LOCAL_MACHINE\Software\WOW6432Node\ComputerAssociates\CA
DataMinder\CurrentVersion\EMail
```

Note: This script requires Windows Script Host 5.6 or later. You can download version 5.6 from the Microsoft Web site: http://msdn.microsoft.com/scripting.

You configure the script to create and specify the following registry values:

NeverPutPDLsInRecipientsList

> **Type:** REG_DWORD
>
> **Data:** Specifies that CA DataMinder should prevent unexpanded (or only partially expanded) PDLs from being saved as email recipients. This prevents control triggers from activating inadvertently.
>
> **Default:** zero

MaxNumExpandedRecipients

> **Type:** REG_DWORD
>
> **Data:** Specifies the maximum number of recipients extracted from distribution lists in a single email. This setting can alleviate email processing delays.
>
> **Default:** zero, CA DataMinder can extract an unlimited number of recipients.

DontProcessReceived

Note: Applies to incoming emails only.

Type: REG_DWORD

Data: Specifies that CA DataMinder does not process incoming emails until the recipient attempts to read them.

Default: zero, CA DataMinder processes incoming emails as soon as they arrive in the recipient's inbox.

Important! If DontProcessReceived is set to a non-zero value, this can affect the operation of control actions for incoming emails. First, a control action cannot generate Notify events. Second, if a control action specifies 'Block Quietly', the blocked email arrives and stays in the recipient's inbox until the recipient tries to read the email. At this point, the email is deleted.

Note: A non-zero DontProcessReceived only yields improved performance if Microsoft Outlook is running on the recipient machine at the time when the email arrives.

NoSenderExtendedInfo

Type: REG_DWORD

Data: Specifies that CA DataMinder only extracts basic information for each sender from the Exchange server, including the sender's display name, email address and the address format. Applies to incoming and manually captured emails only.

Default: zero, dlp> extracts basic information plus each sender's 'true' display name and email address aliases.

Important! If NoSenderExtendedInfo is set to a non-zero value, incoming email policy triggers configured to detect a sender's email address will not activate if the targeted address is an alias.

Note: If NoSenderExtendedInfo and DontProcessReceived (see above) are both set to non-zero values, NoSenderExtendedInfo yields improved performance only when incoming emails are read; if DontProcessreceived is zero or does not exist, then a non-zero NoSenderExtendedInfo yields improved performance both when incoming emails arrive in the recipient's inbox and when they are read.

OutlookMonitorIntervalInSeconds

Type: REG_DWORD

Data: Specifies how often CA DataMinder verifies whether the Outlook client agent is disabled in the current session. (A security feature in Outlook can automatically disable certain add-ins). Set this option to 0 to disable monitoring. Specifically, it checks the registry for the wgnemol.dll name and path.

Default: 5 seconds

OutlookRepairDisabledExtension

Type: REG_DWORD

Data: Specifies whether CA DataMinder should re-enable the Outlook client agent if it is found to be disabled. If this registry value is set to a non-zero value and the Outlook client agent is found to be disabled, CA DataMinder re-enables the client agent and writes a Windows application log entry to that effect.

Default: 0 (do not re-enable)

The transform creates the DWORD registry values in the following registry key on the machine hosting the Outlook client agent:

```
HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA DataMinder\CurrentVersion\EMail
```

Or for 32-bit Outlook client running on a 64-bit OS:

```
HKEY_LOCAL_MACHINE\Software\WOW6432Node\ComputerAssociates\CA
DataMinder\CurrentVersion\EMail
```

Prevent Consoles being Installed: HideConsole_Client.mst

For command line, Group Policy or SMS installations, you can prevent users from installing any CA DataMinder console. HideConsole.mst deselects and removes the Management Console feature from the Custom Setup screen in the installation wizard.

Without this transform, if you deliberately excluded consoles from your CA DataMinder client machines, a user could use Add/Remove Programs to subsequently install the Administration console and, potentially, access sensitive information. Using this transform eliminates any such risk.

1. Find the HideConsole.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Run the script.

   It creates the HideConsole_Client.mst (or HideConsole_Client_x64.mst) transform.
3. Copy the transform into the folder containing your administrative installation source image.

Install Application Integration: EnableAppmon_Client.mst

A default CA DataMinder installation does not include application integration. To install application integration as part of a command line, Group Policy or SMS installation, you must specify the EnableAppmon.mst transform.

Note: Application Integration enables CA DataMinder to monitor usage of desktop applications (excluding email and browser applications) and capture application usage metrics.

1. Find the EnableAppmon.vbs script in the \Support folder of your CA DataMinder distribution media.
2. Run the script.

   It creates the EnableAppmon_Client.mst (or EnableAppmon_Client_x64.mst) transform.
3. Copy the transform into the folder containing your administrative installation source image.

Stopping and Restarting the Infrastructure

You can manually stop and restart the infrastructure (wgninfra.exe) on any CA DataMinder machine. Run the following commands.

- To stop the infrastructure, run:

  ```
  net stop wgninfra
  ```
- To restart the infrastructure, run:

  ```
  net start wgninfra
  ```

Command Line Parameters for Msiexec.exe

Command line deployment operations for CA DataMinder can use the standard Msiexec.exe options. For install operations, you can also use various proprietary variables to configure the installation.

More information:

Operations (see page 532)
General Variables (see page 533)
Database Variables (see page 540)
Logon Variables (see page 547)

Operations

See your Microsoft documentation for full details about command line options for Msiexec.exe. The options you are most likely to need are listed here.

/a <Path>

Creates an administrative installation source image. <Path> is defined below. The syntax is:

```
msiexec /a <Path>\client.msi
```

/i <Path>

Specifies an install operation. <Path> is defined below. The basic syntax is as follows, but see the note below:

```
msiexec /i <Path>\client.msi
```

Note: For CA DataMinder installations, if you use the /i option, you must also specify a /qn or /qb option. See the next column for details.

/x <Path>

Specifies an uninstall operation. <Path> is defined below. The basic syntax is as follows, but see the note below:

```
msiexec /x <Path>\client.msi
```

Note: For CA DataMinder uninstall operations, if you use the /x option, you must also specify a /qn or /qb option. See the next column for details.

<Path>

For administrative installations, this specifies the path to the original shipped Client.msi. file. Msiexec.exe uses this file to generate a source image. Find this file on your CA DataMinder distribution media.

When installing or uninstalling client machines, this specifies the path to the Client.msi file in the CA DataMinder source image on your network. This is the file created by the administrative installation.

If the path includes folder names with spaces, you must enclose the path in quotes.

/qn or /qb

Various /q options let you specify silent or near-silent installations or uninstallations. For example:

- /qn specifies no user interface. The operation is completely silent and requires no user interaction. To install silently, the syntax is:

  ```
  msiexec /i <Path>\client.msi /qn
  ```

- /qb specifies a progress bar and a simple confirmation dialog when the operation is complete. To uninstall, the syntax is:

  ```
  msiexec /x <Path>\client.msi /qb
  ```

More information:

Performing an Administrative Installation (see page 523)

General Variables

You can use the following variables to configure command-line deployment operations based on Msiexec.exe. For key variables, supply a new value; for others, only include the variable if you want to use a nondefault value. Examples are shown at the end of this section.

INSTALLDIR=<path>

Specifies the installation folder in installation operations. <path> is the full path to the installation folder for CA DataMinder. INSTALLDIR must be in uppercase.

If omitted, the path defaults to:

- (on 32-bit systems) %ProgramFiles%\CA\CA DataMinder
- (on 64-bit systems) %ProgramFiles(x86)%\CA\CA DataMinder

INSTALLDIR64=<path>

Specifies the installation folder for 64-bit files in installation operations. <path> is the full path to the 64-bit installation folder for CA DataMinder. INSTALLDIR64 must be in uppercase.

If omitted, the path defaults to %ProgramFiles%\CA\CA DataMinder .

WGNADMINUSERNAME=<username>

Specifies the primary administrator in CMS installations. <username> is the account name for the primary administrator. Also, WGNADMINUSERNAME must be in uppercase. If omitted, the account name defaults to 'Administrator'.

WGNADMINPASSWORD=<password>

Specifies the primary administrator password in CMS installations. <password> is the password for the primary administrator account. Also, WGNADMINPASSWORD must be in uppercase.

Important! Do not omit this variable. By default, the password is null: set a new password for the primary administrator.

WGNSERVERTYPE=<type>

Specifies whether to install a CMS or gateway for server installations. <type> is either CMS or GATEWAY. Also, both WGNSERVERTYPE and its value must be in uppercase. If omitted, the server type defaults to CMS.

WGNPARENTSERVERNAME=<parent>

Identifies the parent server when installing a gateway or client machine, where <parent> is the name or IP address of the parent server. Also, WGNPARENTSERVERNAME must be in uppercase. For example:

```
WGNPARENTSERVERNAME=255.255.255.0
```

If you specify the CMS computer name, enclose the name in double quotes, for example:

```
WGNPARENTSERVERNAME="CMS-HARDY"
```

Important! Do not omit this variable. Specify a parent server or the installation fails.

WGNDELETEDATABASE=0

Specifies that the local database is not deleted when uninstalling CA DataMinder. Also, WGNDELETEDATABASE must be in uppercase.

TRANSFORMS=<transform>

Applies a transform when performing an administrative installation. <transform> is the name of the .MST transform file.

The transform file must be in the same folder as your administrative installation source image. Use a semicolon to separate a list of transforms. TRANSFORMS must be in uppercase.

For example, in an SMS installation you can apply a transform to enable silent uninstallations:

```
TRANSFORMS=SMSQuietUninstall.mst
```

WGNNOEXPLORER=1

When the Internet Explorer Integration feature is installed, CA DataMinder automatically integrates with the **Windows Explorer** browser.

This variable *disables* integration with the Windows Explorer browser. That is, if a user browses the Web using Windows Explorer as a browser, CA DataMinder does not monitor this activity. WGNNOEXPLORER must be in uppercase.

Note: Integration with **Internet Explorer** is not affected by this variable and web activity in Internet Explorer is detected as normal.

WGNNOOUTLOOKBROWSER=1

When the Internet Explorer and Outlook integration features are installed on the same machine, CA DataMinder automatically integrates with the **Microsoft Outlook** browser.

This variable *disables* integration with the Outlook browser. That is, if a user surfs the Web using Outlook as a browser, CA DataMinder does not monitor this activity. WGNNOOUTLOOKBROWSER must be in upper case.

Note: Email integration with **Microsoft Outlook** is not affected by this variable activity in Outlook. Email continues to be monitored as normal.

WGNDEFAULTUSERGROUPPATH=<group path>

Specifies the parent group for new users created on the local machine when installing a client machine. You can use this variable in multiple source images so that when different teams or departments install CA DataMinder, their respective users are added automatically to separate groups.

<group path> is the full path to the parent group; the path must start with the top-level group user (by default, ‘Users’). WGNDEFAULTUSERGROUPPATH must be in uppercase. Use forward slashes ‘/’ as path separators. If the path includes group names with spaces, enclose the path in quotes. Verify that the CMS machine policy setting ‘Allow Client Machines to Specify Default User Group’ is set to True.

If omitted, new users are added to the prevailing default user group.

ADDLOCAL=<componentname>

Specifies the name of a component you want to install.

REMOVE=<componentname>

Specifies the name of subcomponent you want to remove.

More information:

List of Components (see page 537)

Examples

Example 1

This example installs a CMS. The installation folder is C:\CA DataMinder. The primary administrator is frankschaeffer (password: dsf8534mnfg).

```
msiexec /i <Path>\server.msi
  INSTALLDIR="C:\CA DataMinder"
  WGNADMINUSERNAME=frankschaeffer
  WGNADMINPASSWORD=dsf8534mnfg
  WGNSERVERTYPE=CMS
```

Example 2

This example silently installs a client machine to C:\Program Files\CA DataMinder. The parent server is CMS-HARDY. The installation also includes two transforms.

```
msiexec /i <Path>\client.msi /qn
  INSTALLDIR="C:\Program Files\CA DataMinder"
  WGNPARENTSERVERNAME="CMS-HARDY"
  TRANSFORMS=ClientLockDown.mst;
  EnableAppmon.mst
```

Example 3

This example silently uninstalls a client machine but retains the local database:

```
msiexec /x <Path>\client.msi /qn WGNDELETEDATABASE=0
```

Example 4

This example ensures that new users are always added to the ‘Sales’ group:

```
msiexec /i <Path>\client.msi /qn
  WGNDEFAULTUSERGROUPPATH=Users/Sales
```

List of Components

This section lists the component names for each CA DataMinder msi image. These component names can be used with the ADDLOCAL and REMOVE variables. The list takes this format:

```
Component: <componentname>
```

Where <componentname> is the name that you assign to ADDLOCAL or REMOVE.

Server.msi

Policy Engine: PolicyEngine

Socket API: SocketAPI

Event Import - DataImport

Remote Policy Engine Connector: ActiveImportConnector

Templates: INITemplates

Content Indexer Console: ContentIndexerUI

Content Indexer Server: ContentIndexerService

Content Proxy Server: ContentProxyServer

Remote Data Manager: RDMServer

Quarantine Manager: QuarantineManager

EMC Centera Connector: Centera

IBM Content Manager Connector: IBMCM

Administration Console: ServerConsole

Data Management Console: DataExtensions

Web.msi

- Front-end Web Server: WebServer
- Application Server: WebService

Reports.msi

- Dashboard: Dashboard
- Compliance Reports
  - Compliance Audit Report: ComplianceAuditReport
  - Employees Not Reviewed Report: EmployeesNotReviewedReport
  - Proof of Supervision Report: ProofOfSupervisionReport
  - Repeat Offender Report: RepeatOffenderReport
  - Review Latency Report: ReviewLatencyReport
  - Reviewer Activity Report: ReviewerActivityReport
- Incident Reports
  - Incident Cause Frequency Report: IncidentCauseFrequency
  - Incident Rate by Policy Report: IncidentRateByPolicy
  - Incident Summary Report: IncidentSummaryReport
  - Incidents by Location Report: IncidentsByLocation
  - Incidents by Policy and Action Report: IncidentsByPolicyAndAction
  - Incidents by Policy and Channel Report: IncidentsByPolicyAndChannel
  - Incidents by Policy and Time Period Report: IncidentsByPolicyAndTimePeriod
- Issue Reports
  - Detailed Issue Report: DetailedIssueReport
  - Issues by Status or Resolution Report: IssuesByStatusOrResolution
- Review Queue
  - Administrative Reports: RQAdminReports
  - Reviewer Search: RQReviewerSearch

Integration.msi

- Email server agents
  - Domino Server Agent: DSA
  - Exchange Server Agent (32-bit package only): ESA
  - IIS SMTP Agent: IIS_SMTP (32-bit package only)
- Archive Agents
  - EMC SourceOne: EMCS1
  - Symantec Enterprise Vault: SEV
  - ZANTAZ Digital Safe Adapter: ZDS
- **File Scanning Agent: FSA**
  - Administration Console: ServerConsole
  - NIST Database Connector: NIST
  - Remote Policy Engine Connector: FSAHub
  - FSA Remote Connector: FSARemoteConnector
- External Agent API: RDI
- Remote Policy Engine Connector: RDIHub
- ICAP Agent: ICAP
- Bloomberg and IM Support: IMImport

Integration_x64.msi

- Exchange Server Agent: ESA_x64

UA.msi

- De-duplication Support: SingleInstancing

Database Variables

When installing a CMS or gateway from a command line, you can use these variables to configure the database engine and the database account used by CA DataMinder. For key variables, you *must* supply a new value; for others, you only need to include the variable if you want to use a non-default value. Examples are shown at the end of this section.

Note: Note that these database details cannot not be validated during the installation; you must ensure that you have entered them correctly, otherwise the installation will fail.

Important! All database variables and their values must be in uppercase.

- **General variables**

  ```
  WGNDATABASETYPE=<type>
  WGNDATABASESERVER=<server>
  WGNDATABASEIPPORT=CA Portal
  WGNDATABASESERVICENAME=<service>
  WGNDATABASENAME=<dbname>
  WGNDATABASEUSERNAME=<username>
  WGNDATA=<path>
  WGNDATABASEPASSWORD=<password>
  ```

- **Primary user variables**

  ```
  WGNDBPRIMARYUSERNAME=<username>
  WGNDBPRIMARYPASSWORD=<password>
  WGNDBPRIMARYCREATEACCOUNT=1
  WGNDBPRIMARYTABLESPACENAME=<name>
  ```

- **Search user variables**

  ```
  WGNDBSEARCHUSERNAME=<username>
  WGNDBSEARCHPASSWORD=<password>
  WGNDBSEARCHCREATEACCOUNT=1
  WGNDBSEARCHTABLESPACENAME=<name>
  ```

- **Owner user variables**

  ```
  WGNDBOWNERUSERNAME=<username>
  WGNDBOWNERUSERPASSWORD=<password>
  WGNDBOWNERCREATEACCOUNT=1
  WGNDBOWNERTABLESPACENAME=<name>
  ```

- **Admin user variables**

  ```
  WGNDBADMINUSERNAME=<username>
  WGNDBADMINPASSWORD=<password>
  ```

- **Data warehouse user variables**

  ```
  WGNDBREPORTINGUSERNAME=<username>
  WGNDBREPORTINGPASSWORD=<password>
  WGNDBREPORTINGCREATEACCOUNT=1
  WGNDBREPORTINGTABLESPACENAME=<name>
  ```

- **Unrestricted search user variables**
  ```
  WGNDBUNRESTRICTEDUSERNAME=<username>
  WGNDBUNRESTRICTEDPASSWORD=<password>
  WGNDBUNRESTRICTEDCREATEACCOUNT=1
  WGNDBUNRESTRICTEDTABLESPACENAME=<name>
  ```
- **Data warehouse variables**
  ```
  WGNDWENABLE=1
  WGNDWEVENTPARTICIPANTS=1
  ```

These variables are described below.

WGNDATABASETYPE=<type>

Identifies the database engine that you will use on the CMS. <type> is either MSSQL or ORACLE. If you omit this variable, the database engine defaults to MSSQL.

WGNDATABASESERVER=<server>

Identifies the host server for your database, where <server> is the name or IP address of the database server. If you specify a computer name, you must enclose the name in double quotes.

If you omit this variable, the database server defaults to localhost. This specifies a database on the local machine.

WGNDATABASEIPPORT=CA Portal

Specifies the TCP/IP port number used by the database host server. If you omit this variable, the default depends on the type of database engine (as specified by WGNDATABASETYPE). Note that you do not normally need to change these defaults, meaning you can normally omit this variable. The default port numbers are:

1433 for SQL Server.

1521 for Oracle.

WGNDATABASESERVICENAME=<service>

For Oracle databases, this specifies a service name to identify the correct database tables, where <service> is the name of the Oracle service.

Important! Do not omit this variable; you *must* supply a service name.

Note: This is not required for SQL Server databases.

WGNDATABASENAME=<dbname>

For SQL Server databases, this specifies the name of the database, where <dbname> is the database you want to create or (if already created) use on the database server.

If you omit this variable, the database name defaults to WGN_<machine name> where <machine name> is the name of the server on which you are installing the CMS or gateway.

Note: This is not required for Oracle databases.

WGNDATABASEUSERNAME=<username>

Specifies the user name for a valid database account, where <username> is the user name. WGNDATABASEUSERNAME must be in upper case. CA DataMinder uses this account to access the CMS database. You must ensure that the account has appropriate roles and privileges.

Note: This value is still supported, but has been superseded by WGNDBPRIMARYUSERNAME.

WGNDATA=<path>

Specifies the name and network location of the data folder. This folder contains all the configuration data and captured data used by the CA DataMinder installation. <path> is the full path to this folder.

To specify a remote location, you cannot specify a mapped network drive. Instead, you must enter a network file share, using the universal naming convention, for example:

```
\\MyMachine\share_name\target_folder
```

If you omit this variable, the folder defaults to CA's \data\log subfolder of the Windows All Users profile.

WGNDATABASEPASSWORD=<password>

Specifies the password for the database account used by CA DataMinder, that is, the account specified by WGNDATABASEUSERNAME. If you omit this variable, the password defaults to null (that is, no password is needed to access the database).

Note: This value is still supported, but has been superseded by WGNDBPRIMARYPASSWORD.

WGNDBPRIMARYUSERNAME=<username>

Specifies the user name for a valid CA DataMinder database user account, where <username> is the user name. CA DataMinder uses this account to access the CMS database. You must ensure that the account has appropriate roles and privileges, unless you intend to specify WGNDBPRIMARYCREATEACCOUNT which assigns the appropriate roles and privileges automatically.

WGNDBPRIMARYPASSWORD=<password>

Specifies the password for the database account used by CA DataMinder, that is, the account specified by WGNDBPRIMARYUSERNAME. If you omit this variable, the password defaults to null (that is, no password is needed to access the database).

WGNDBPRIMARYCREATEACCOUNT=1

Set this optional variable to 1 to specify that if no valid database account exists on the server, then CA DataMinder will automatically create one with the appropriate roles and privileges.

WGNDBPRIMARYTABLESPACENAME=<name>

(Oracle only) This specifies a valid name name to identify the correct tablespace for the primary database account, where <name> is the name of the Oracle database tablespace.

Important! Do not omit this variable; you *must* supply a tablespace name.

WGNDBSEARCHUSERNAME=<username>

Specifies the user name for the database search user account, where <username> is the user name. CA DataMinder uses this account when searching

the CMS database for events. You must ensure that the account has appropriate roles and privileges, unless you specify WGNDBSEARCHCREATEACCOUNT which assigns the appropriate roles and privileges automatically.

Note: This is a secure account that is subject to ‘row level security’ when searching the database for events. This ensures that reviewers cannot see events associated with users outside of their management groups.

WGNDBSEARCHPASSWORD=<password>

Specifies the password for the database search user account used by CA DataMinder, that is, the account specified by WGNDBSEARCHUSERNAME. If you omit this variable, the password defaults to null (that is, no password is needed to search the database).

WGNDBSEARCHCREATEACCOUNT=1

Set this optional variable to 1 to specify that if no valid search user account exists on the server, then CA DataMinder will automatically create one with the appropriate roles and privileges.

WGNDBSEARCHTABLESPACENAME=<name>

(Oracle only) This specifies a valid name to identify the correct tablespace for the search user account, where <name> is the name of the Oracle database tablespace.

Important! Do not omit this variable; you **must** supply a tablespace name.

WGNDBOWNERUSERNAME=<username>

(Oracle only) This optional variable specifies the user name for the database owner user account, where <username> is the user name. This account owns the database schema in preference to the primary user. You must ensure that the account has appropriate roles and privileges.

WGNDBOWNERPASSWORD=<password>

This optional variable specifies the password for the database owner account used by CA DataMinder, that is, the account specified by WGNDBOWNERUSERNAME. If you specify WGNDBOWNERUSERNAME but omit this variable, the password defaults to null (that is, no password is needed to access the database).

WGNDBOWNERCREATEACCOUNT=1

(Oracle only) Set this optional variable to 1 to specify that if no valid database owner account exists on the server, then CA DataMinder will automatically create one with the appropriate roles and privileges.

WGNDBOWNERTABLESPACENAME=<name>

(Oracle only) This specifies a valid name to identify the correct tablespace for the database owner account, where <name> is the name of the Oracle database tablespace.

Important! Do not omit this variable; you **must** supply a tablespace name.

WGNDBADMINUSERNAME=<username>

Specifies the user name for database administrator account, where <username> is the user name. CA DataMinder uses this account to create users. You must ensure that the account has appropriate roles and privileges.

This variable is only required if any of the following are set to 1:

- WGNDBPRIMARYCREATEACCOUNT
- WGNDBSEARCHCREATEACCOUNT
- WGNDBOWNERCREATEACCOUNT

WGNDBADMINPASSWORD=<password>

Specifies the password for the database administrator account used by CA DataMinder. That is, the account specified by WGNDBADMINUSERNAME.

Important! Do not omit this variable; you **must** supply a password for this account.

WGNDBREPORTINGUSERNAME=<username>

This optional variable specifies the user name for the data warehouse user account, where <username> is the user name. External reporting applications (such as BusinessObjects Enterprise) use this database account to connect to the Data Warehouse and CMS database.

WGNDBREPORTINGPASSWORD=<password>

This optional variable specifies the password for the data warehouse user account used by CA DataMinder, that is, the account specified by WGNDBREPORTINGUSERNAME. If you specify WGNDBREPORTINGUSERNAME but omit this variable, the password defaults to null (that is, no password is needed to access the database).

WGNDBREPORTINGCREATEACCOUNT=1

Set this optional variable to 1 to specify that if no valid data warehouse user account exists on the server, then CA DataMinder will automatically create one with the appropriate roles and privileges.

WGNDBREPORTINGTABLESPACENAME=<name>

This specifies a valid name to identify the correct tablespace for the data warehouse user account, where <name> is the name of the Oracle database tablespace.

Important! Do not omit this variable; you **must** supply a tablespace name.

WGNDBUNRESTRICTEDUSERNAME=<username>

This optional variable specifies the user name for the unrestricted search user account, where <username> is the user name.

This database account corresponds to the 'Unrestricted' security model. CA DataMinder consoles and external reporting tools can use this database account when searching the CA DataMinder data warehouse and CMS database for events. Unlike normal Search User database accounts, the Unrestricted Search User is *not* subject to row level security (RLS) when searching the database.

WGNDBUNRESTRICTEDPASSWORD=<password>

This optional variable specifies the password for the unrestricted search user account used by CA DataMinder, that is, the account specified by WGNDBUNRESTRICTEDUSERNAME. If you specify WGNDBUNRESTRICTEDUSERNAME but omit this variable, the password defaults to null (that is, no password is needed to access the database).

WGNDBUNRESTRICTEDCREATEACCOUNT=1

Set this optional variable to 1 to specify that if no valid unrestricted search user account exists on the server, then CA DataMinder will automatically create one with the appropriate roles and privileges.

WGNDBUNRESTRICTEDTABLESPACENAME=<name>

This specifies a valid name to identify the correct tablespace for the unrestricted search user account, where <name> is the name of the Oracle database tablespace.

Important! Do not omit this variable; you **must** supply a tablespace name.

WGNDWENABLE=1

Set this optional variable to 1 to install a data warehouse.

WGNDWEVENTPARTICIPANTS=1

Set this optional variable to 1 to collect event participant data in the data warehouse.

This option is essential if you want to run reports or dashboards that show results broken down by user group, or if you want to apply management group row level security.

Important! If you collect event participant data, the associated table in the data warehouse can grow very large. Our testing indicates that the table can increase the size of the database by 30-40%. The reason is because an individual event can have many participants. For further details, see the *Reports Integration Guide*.

Examples

Example 1

This example installs a CMS and specifies a SQL Server database engine. The host server is MyDBServer, which uses IP port 1480; the database account is fred (password: xfz3105labp); the \Data folder is located on Server02 at \CA DataMinder\CA DataMinder_data.

```
msiexec /i <Path>\server.msi
  //
  // General variables go here
  //
  WGNDATABASETYPE=MSSQL
  WGNDATABASESERVER=MyDBServer
  WGNDATABASEIPPORT=1480
  WGNDATABASENAME=My_CA DataMinder_Database
  WGNDATABASEUSERNAME=fred
  WGNDATABASEPASSWORD=xfz3105labp
  WGNDATA=\\Server02\CA DataMinder\CA DataMinder_data
```

Example 2

This example upgrades a SQL Server database engine. The command create a Reporting User and an Unrestricted Search User. It also enables the data warehouse and turns on collection of event participant data.

```
msiexec /i %camden_msi_path% /qb
  //
  // General variables and
  // database variables go here
  //
  WGNDBREPORTINGUSERNAME=wgnreporting
  WGNDBREPORTINGPASSWORD=%PASSWORD%
  WGNDBREPORTINGCREATEACCOUNT=1
  WGNDBUNRESTRICTEDUSERNAME=wgnrestricted
  WGNDBUNRESTRICTEDPASSWORD=udz8343lgs
  WGNDBUNRESTRICTEDCREATEACCOUNT=1
  WGNDWENABLE=1
  WGNDWEVENTPARTICIPANTS=1
  WGNDBADMINUSERNAME=sa
  WGNDBADMINPASSWORD=hyd8734pcs
```

Logon Variables

Use the following variables to specify logon accounts for CA DataMinder services. Examples are shown at the end of this section.

WGNINFRAFULLNAME=<username>

If the CA DataMinder infrastructure service (wgninfra.exe) needs to log on as a named user, use this variable to specify the domain account.

WGNINFRAFULLNAME must be in uppercase. Names must be in domain\username format.

WGNINFRAFPASSWORD=<password>

Specifies the password for the logon account specified by WGNINFRAFULLNAME.

WGNINFRAPASSWORD must be in uppercase.

WGNIMPSVFULLNAME=<username>

When Event Import runs as a service, it needs to log on to the CMS host machine. Use this variable to specify the domain account that the Event Import service uses.

WGNIMPSVFULLNAME must be in uppercase. Names must be in domain\username format.

WGNIMPSVPASSWORD=<password>

Specifies the password for the logon account specified by WGNIMPSVFULLNAME.

WGNIMPSVPASSWORD must be in uppercase.

WGNPESVFULLNAME=<username>

Use this variable to specify the PE domain user for CA DataMinder policy engines.

WGNPESVFULLNAME must be in uppercase. Names must be in domain\username format.

WGNPESVPASSWORD=<password>

Specifies the password for the logon account specified by WGNPESVFULLNAME.

WGNPESVPASSWORD must be in uppercase.

WGNAUTSPFULLNAME=<username>

Use this variable to specify the domain account that the Content Indexer Service uss.e

WGNAUTSPFULLNAME must be in uppercase. Names must be in domain\username format.

WGNAUTSPPASSWORD=<password>

Specifies the password for the logon account specified by WGNAUTSPFULLNAME.

WGNAUTSPPASSWORD must be in uppercase.

Environment Variables

The CA DataMinder installers create the following Windows environment variables:

WGNDATADIR

This variable is set to the CA DataMinder data folder. You specify the data folder when you install a CMS or gateway.

- Example for Windows Server 2008:
  `C:\ProgramData\CA\CA DataMinder\data`

WGNINSTALLDIR

This variable is set to the 32-bit CA DataMinder installation folder. You specify the installation folder when you install a CA DataMinder component.

- Example for 32-bit Windows:
  `C:\Program files\CA\CA DataMinder\`
- Example for 64-bit Windows:
  `C:\Program files (x86)\CA\CA DataMinder\`

WGNINSTALLDIR64

This variable is set to the 64-bit CA DataMinder installation folder. This variable is only set on 64-bit operating systems by CA DataMinder 64-bit installers.

You specify the installation folder when you install a CA DataMinder component.

- Example for 64-bit Windows:
  `C:\Program files\CA\CA DataMinder\`

Standalone Installations

A standalone installation is a single machine operating simultaneously as a CMS and client machine. The main reasons for standalone installations are as demonstration machines.

For a standalone to offer the full range of CA DataMinder features, it must have a locally installed Oracle or Microsoft SQL Server database.

Installing a Full-featured Standalone

To install a standalone that can demonstrate all key CA DataMinder features

1. Run the server installation wizard.
    a. In the Custom Setup screen, install all the Enterprise Server and Management Console features that you will need.
    b. In the Server Type screen, choose to install a Central Management Server.
    c. You can now step through the usual CMS installation screens, specifying the data location, database engine, and primary administrator credentials in the usual way.
2. Run the client installation wizard. Step through the client installation screens in the usual way, specifying which client integration features to install and the connectivity settings.

Installing a Standalone to Demonstrate Client Agents

To install a standalone to demonstrate CA DataMinder client agents

1. You must configure and install Oracle or SQL Server on the target standalone machine.
2. Run the client installation wizard.
3. In the Custom Setup screen, install all the Client Integration and Management Console features that you will need.
4. In the Connectivity screen, choose the Standalone Mode option. The screen sequence now changes from being client-focused to CMS-focused. For example, the next screen is the Data Location screen.
5. You can now step through the usual CMS installation screens, specifying the data location, database engine, and primary administrator credentials in the usual way.

Set Account Credentials with WgnCred.exe

CA DataMinder provides a command line utility (wgncred.exe) to set account credentials for various components. For example, you can use it to securely cache database logon credentials for the FSA, so avoiding the need to include these credentials in the scanning job definition file.

After installing the relevant CA DataMinder component, you must run wgncred.exe on the host machine of that component to set the required credentials. For example, to set the Notes password for Event Import, you must run wgncred.exe on the Domino server hosting Event Import.

More information:

Supported CA DataMinder Components (see page 551)
Account Credentials Operations (see page 552)

Supported CA DataMinder Components

You can use wgncred.exe to set the credentials for components. Each has its own component ID.

Event Import for Notes Import Client

When importing Lotus Notes e-mails, Event Import accesses the Notes databases as Notes user. Instead of specifying the password for this user in the import configuration file, we recommend that you cache this password using wgncreed.exe.

Component ID: **NotesImport**

File Scaning Agent (FSA)

When scanning database records, the FSA connects to the target database using the authentication method specified in the scanning job definition. As a secure alternative to storing a database user's password in plain text in the XML job definition file, you can use wgncred.exe to cache the password. (This is especially important when scanning Oracle databases because we recommend that the FSA does *not* use 'Windows authentication').

You indicate that you are using cached database credentials when you create a database scanning job in the Administration console; see the line help for details; search for 'connection string'.

Component ID: **FSADBSCAN**

Policy Engines connecting to Voltage SecureEmail Servers

When decrypting emails encrypted with Voltage SecureMail, policy engines use the cached account credentials to connect to a Voltage SecureMail server.

Component ID: **Voltage**

Quarantine Manager connecting to Lotus Notes mailboxes

To specify which user account the Quarantine Manager (QM) uses to release quarantined Notes e-mails, you must specify a Domino mailbox in the registry. You must also cache the password for the Notes 'Current User' on the QM host server.

Component ID: **QMGRNOTES**

Quarantine Manager connecting to SMTP servers

To specify which user account the Quarantine Manager (QM) uses to release quarantined SMTP e-mails, you must specify a Domino mailbox in the registry. You must also cache the password for the Notes 'Current User' on the QM host server.

Component ID: **QMGRSMTP**

RDM connecting to IBM DB2 CommonStore

The RDM uses the cached account credentials to retrieve emails from the IBM DB2 Content Manager archive during an event search. For CommonStore for Lotus Domino, the RDM must also connect to the Notes Client as part of the retrieval process.

(Exchange) Content Manager Component ID: **RDMCSX**

(Domino) Content Manager component ID:**RDMCSLD**

Notes Client component ID:**RDMCSLDClient**

RDM connecting to Symantec Enterprise Vault for Domino

The RDM uses the default Notes user on your Enterprise Vault server to retrieve archived Domino emails. Use wgncred.exe to cache the password for this account.

Component ID: **RDMSEVLDClient**

Socket API

To log onto an SMTP server in order to send user notifications, the Socket API uses the account specified by the UserID registry value. Use wgncred.exe to cache the password for this account.

Component ID: **EANotification**

More information:

Notes Servers: Set Credentials (see page 235)
SMTP Servers: Set Credentials (see page 235)

Account Credentials Operations

From a command prompt in the \system subfolder in the CA DataMinder installation folder, you can use wgncred.exe to:

- Set full account credentials
- Clear a specific account password
- Manually set an account password
- Manually set an account password
- Show password status for supported components

More information:

Set an Account Password (see page 553)
Manually Clear an Account Password (see page 553)
Show Password Status for Supported Components (see page 554)
Set Full Account Credentials (see page 554)

Set an Account Password

WgnCred.exe can list all the supported components that are currently installed on the host machine. You simply choose the component you want and enter a password.

Follow these steps:

1. From a command prompt in the \System subfolder in the CA DataMinder installation folder, run:

   ```
   wgncred -set
   ```

   A list of components is displayed with their corresponding ID numbers and component identifiers.

2. When prompted, enter the ID number of the component you want to set a password for.
3. Type a password for the selected component.

Manually Clear an Account Password

WgnCred.exe can list all the supported components that are currently installed on the host machine. You simply choose the component you want and enter a password.

Follow these steps:

1. From a command prompt in the \System subfolder in the CA DataMinder installation folder, run:

   ```
   wgncred.exe -clear
   ```

   A list of components is displayed with their corresponding ID numbers and component identifiers.

2. When prompted, type the ID number of the component you want to clear the password for.

Show Password Status for Supported Components

You can view all supported components and their password status (that is, whether a component currently has a password set).

From a command prompt in the \System subfolder in the CA DataMinder installation folder, run:

```
wgncred -query
```

Set Full Account Credentials

From a command prompt in the \System subfolder in the CA DataMinder installation folder, run:

```
wgncred -set <component identifier> <domain> <username> <password>
```

where:

<Component identifier>

Is a unique ID for the component.

<domain>

Is the user's domain.

<username>

Is the user's name.

<password>

Is the password for the user account used by the component.

For example, to cache database logon credentials for the FSA where the domain, name and password for the database user are respectively unipraxis, srimmel, and MyDatabasePW, run this command:

```
wgncred -set FSADBScan unipraxis srimmel MyDatabasePW
```

Importing and Exporting CMS Profiles

You can export CMS profiles to an XML file and then import them to other CMSs in your CA DataMinder enterprise. CMS profiles consist of certain files in the CA DataMinder file system, plus the CMS installation properties located in the CA DataMinder WgnWellKnownString, WgnFile, Wgn3Role and Wgn3ResourceRole database tables. For example, a profile includes:

- **User attributes:** The name and index of all user attributes.
- **Role privileges:** The name of each role and the privileges assigned to that role.
- **Audit settings:** All audit settings.
- **Audit templates:** The names and contents of the audit templates.
- **Severity thresholds:** The low, medium and high severity threshold values for triggers.
- **Policy Roles**: The names and associated resources of all policy roles.

Export a CMS Profile

To export a CMS profile

1. On the CMS machine you want to export from, use a command line to browse to the \System subfolder of the CA DataMinder installation directory.
2. Run the following command line:

   ```
   wgninfra -exec wigan/infrastruct/machine/Profile Export <file name>
   ```

   Where <file name> specifies the path and name of the XML file containing the exported profile. The location of the XML file must be accessible from the CMS that you want the profile to. The command syntax is case-sensitive.

   This exports the CMS profile from the file system, and from the WgnWellKnownString, WgnFile, Wgn3Role and Wgn3ResourceRole database tables, to the specified XML file.

Note: You cannot export single properties. If any data is not required, you need to edit the XML file created by the export process.

Import a CMS profile

To import a CMS profile

1. On the CMS machine you want to import to, use a command line to browse to the \System subfolder of the CA DataMinder installation directory.
2. Now run the following command line:

   ```
   wgninfra -exec wigan/infrastruct/machine/Profile Import <file name>
   ```

   Where <file name> specifies the path and file name of the profile that you previously exported. The command syntax is case-sensitive.

   This imports the XML file and updates the file system and the WgnWellKnownString, WgnFile, Wgn3Role and Wgn3ResourceRole database tables with files and data from the CMS where the XML file was exported from.

Note: Unlike export operations, it *is* possible to import single properties. If any data is not required, you need to edit the XML file created by the export process.

Retain or Overwrite Profile Properties

When importing a CMS profile, you can specify how CA DataMinder reconciles mismatches between existing CMS properties and properties in the imported profile.

Property Exists on CMS But Not in Imported Profile

If a property exists on the CMS but not in the imported profile, you can retain or delete the existing property. To configure the import operation, edit the following attribute in the XML profile:

retainvalues

This optional attribute has the following values:

retainvalues="True"

This is the default value if retainvalues is not specified.

The existing property on the CMS is retained.

retainvalues="False"

The existing property on the CMS is deleted.

Property Exists an CMS and Imported Profile

If a property exists both on the CMS and in the imported profile, you can retain the existing property or overwrite it with the imported property. To configure the import operation, edit the following attribute in the XML profile:

overwritevalues

This optional attribute has the following values:

overwritevalues="True"

The existing property on the CMS is overwritten by the imported property.

overwritevalues="False"

This is the default value if overwritevalues is not specified.

The existing property on the CMS is retained.

Property Exists in Imported Profile But Not on CMS

If a property exists in the imported profile but not on the CMS, the imported property is always added to the CMS profile.

Apply Changes to Whole CMS profile or Individual Sections

You can set retainvalues and overwritevalues for the whole CMS profile or for an individual section, such as the user attributes section.

Apply Changes to Whole CMS Profile

Edit retainvalues or overwritevalue in the <profile> tag. For example:

```
<profile retainvalues="True" version="1.0" machinetype="CMS">
```

If a property exists in both the XML file and on the CMS, then the value of that property on the CMS is overwritten by the value in the XML file.

Apply Changes to Individual Section

Edit retainvalues or overwritevalue in the relevant section tag. For example, to retain existing user attributes, edit the <userattributes> tag:

```
<userattributes retainvalues="True">
```

The supported section tags are:

```
<userattributes>
<roles>
<audit_setttings>
<audit_templates>
<trigger_severities>
<resource_roles>
```

Note: If retainvalues or overwritevaluesattribute exist in the <profile> tag and in a section tag, the section tag takes precedence.

Reset CMS Profile Default Properties

You can reset the default settings of a CMS profile. You can do this for an entire CMS profile, or any properties listed within it.

Reset the Entire CMS Profile

To reset the entire CMS profile

1. In the <profile> tag, set the retainvalues attribute to False.

   ```
   <profile retainvalues="False" version="1.0" machinetype="CMS">
   ```

2. Remove all parameters from between each set of section markers.

3. In each <userattribute> tag, set the retainvalues attribute to False.

   ```
   <userattributes retainvalues="false">
   ```

   Note: If the retainvalue attribute exists in the <profile> tag and the section tag, then the value in the section tag takes precedence.

4. Save and import the XML file.

5. Restart the CMS.

Reset Specific Properties in a CMS Profile

To reset specific properties in a CMS profile

1. In the section containing the properties you want to reset, set the retainvalues attribute to False. For example:

   ```
   <userattributes retainvalues="False">
   ```

2. Remove the parameters you want to reset to their default settings. For example, to reset the entire list of user attributes, the syntax is:

   ```
   <userattributes retainvalues="False"> </userattributes>
   ```

   Note: If the retainvalue attribute exists in the <profile> tag and the section tag, then the value in the section tag takes precedence.

3. Save and import the XML file.

4. Restart the CMS.

Manage System Files

If a problem occurs with your CA DataMinder installation, you may be instructed to examine system files stored in the CMS database. Analysis of the system files can often help CA technical staff diagnose the problem.

CA DataMinder system files include such items as iConsole search definitions, content registration agents, definition files for social security numbers, dynamic address lists and audit mail templates. CA DataMinder stores these system files in the CMS database. Each system file has two parts: the file metadata and the actual file content. The metadata includes such details as the file name and creation date. The file content itself comprises XML or binary data.

The Administration console includes the System File Explorer for managing system files. The System File Explorer is very similar to Windows Explorer and displays the internal file system within the CMS database. It lets you browse the internal file system to view, edit, and replace system files.

More information:

View System Files (see page 560)
Edit System Files (see page 560)
Replace System Files (see page 561)

View System Files

The Administration console lets you view system files in read only mode. Viewing these files in read only mode prevents you from accidentally deleting or modifying crucial data.

Note: You must have the 'Admin: Manage System Files' administrative privilege in order to view system files.

To view system files

1. Click a CMS in the Administration console.
2. Click Tools, View System Files.

   The System File Explorer appears.
3. Browse the system folders to find the file you want to view.
4. Right-click the file and do one of the following:
   - Click Open or Open With to view the file contents in Read Only mode.
   - Click Properties to view the file metadata and other details.

Edit System Files

You cannot edit system files directly in the CMS database. Instead, the Edit System Files feature in the Administration console lets you edit a *copy* of the system file. You can then import the edited copy onto the CMS, replacing the original system file.

Note: You must have the 'Admin: Manage System Files' administrative privilege to edit copies of system files.

To edit system files

1. Click a CMS in the Administration console.
2. Click Tools, Edit System Files.

   The System File Explorer appears.
3. Browse the system folders to find the file you want to view.
4. Right-click the file and click Open or Open With.

   A copy of the file opens in your preferred editor.
5. Amend the file and save the changes.

   The copy of the system file is saved to the folder you specify on the local server or a remote server. If necessary, you can now import this copy onto the CMS to replace the original system file.

Replace System Files

If you want to replace a system file with a different version, you must import a new version of the file. When you import a system file, the content of the original system file is replaced by the content of imported new file. However, the metadata of the original system file is retained.

For backup purposes, we recommend that you first export the original file to a folder on the local server or a remote server.

To export system files

1. Click a CMS in the Administration console.
2. Click Tools.
3. Click View System Files or Edit System Files.

   The System File Explorer appears.
4. Browse the system folders to find the file you want to export.
5. Right-click the file and click Export to copy the file from the CMS database to a folder on the local server or a remote server.

To import a system file

1. Click a CMS in the Administration console.
2. Click Tools.
3. Click View System Files or Edit System Files.

   The System File Explorer appears.
4. Browse the system folders to find the file you want to *replace*.
5. Right-click the file and click Import.

   The Import File dialog appears.
6. Browse to the file whose content you want to *import* and click Open

   The file content is imported, replacing the content of the original system file. The metadata of the original system file is *not* replaced.

IPv6 Address and Port Formats

This section describes the syntax to specify IPv4 and IPv6 addresses, address ranges, ports, and port ranges.

Specifying IPv6 Adresses

Use the following formats:

- Dotted format for IPv4:
  `192.168.0.3`
- Colon-separated format for IPv6:
  `fe80::214:c2ff:fec8:c920`
- All addresses in the range including the lower and upper address:
  `10.0.1.53-10.0.1.80`
- All addresses from 10.0.0.0 to 10.0.255.255:
  `10.0`
- All addresses from 10.0.0.0 to 11.255.255.255:
  `10/7`
- All addresses from "2001:0db8:85a3::" to "2001:0db8:85a3:ffff:ffff:ffff:ffff:ffff":
  `2001:0db8:85a3/48`

Appending Port Numbers to IP Addresses

Enclose IPv6 addresses in brackets when you specify a port number or port range.

Use the following formats:

- IPv4 address and a port:
  `192.168.0.5:10`
- IPv6 address and a port:
  `[fe80::e828:209d:20e:c0ae]:375`
- All addresses where the port number ranges from 137 through 139:
  `:137-139`
- All addresses from 192.168.0.0 to 192.168.255.255 where the port number is from 1024 through 65535:
  `192.168:1024-65535`
- All addresses from fe80:: to fe81:: where the port number is 80:
  `[fe80::]-[fe81::]:80`

Allocating UDP and TCP Ports

Network communication between CA DataMinder machines depends on ports being correctly allocated. Each machine uses a combination of client and server ports. In order to transfer data, a machine creates a connection between a client port and a server port on a remote machine. For example, when the CMS receives a notification to replicate captured events up from a client machine, it creates a connection between one of its client ports and a server port on the client machine.

Communications in CA DataMinder is bidirectional. That is, a CMS can initiate a new connection with a client and a client is able to initiate a new connection to a CMS. This means that all machine types are sources and destinations of UDP packets, and all machine types act as TCP servers and clients.

Example Communication Sequence

As an example, consider the sequence of communications when replicating *captured* data from a gateway or client to a CMS.

Note: Replication of *policy* data from a CMS to a gateway or client uses a symmetric method.

1. Gateway notifies CMS that data is available

   The gateway sends the CMS a CA DataMinder machine identification service UDP packet containing the number of items to replicate and other replication details. The relevant startup.properties setting on both the gateway and the CMS is default.pingports.

2. CMS prepares to pull data up from gateway

   The CMS opens a TCP socket to the gateway and sends and receives packets requesting replication service information from the RMI registry. The relevant startup.properties setting on the gateway is default.rmiports.

   Note: This socket can stay open for up to 15 seconds after the data transfer completes before closing, to save re-negotiating the connection if it is required within that time.

3. CMS pulls data from gateway

   The CMS opens a TCP socket to the gateway and reads captured data from the gateway. The relevant startup.properties setting on the gateway is default.serverports.

   Note: This socket can stay open for up to 15 seconds after the data transfer completes before closing to save re-negotiating the connection if it is required within that time.

Controlling Port Allocation

Port allocation on each machine is controlled by the configuration file startup.properties. This file is in the \system folder below the installation folder on each CA DataMinder machine. There are two possible situations, both rare, when you may need to edit this file to control the allocation of TCP and UDP ports:

- **A port is already allocated to another application**

  Certain port numbers must be the same across all your CA DataMinder machines. For example, this applies to the server port assigned to the CA DataMinder machine identification service. But if an existing application on a CA DataMinder machine already uses all of the default ports specified in startup.properties, you will need to allocate a new port number for use by CA DataMinder.

- **CA DataMinder machines connect through a firewall or router**

  If laptop users operating outside your firewall need to communicate with your CMS inside the firewall, you will need to configure the firewall to allow communication between the ports defined in the startup.properties file of the CMS and laptop client machine.

More information:

Allocating Port Numbers to default.pingports and default.rmiports (see page 565)
Configuration Changes to Support an Internet Firewall or Router (see page 566)

Reallocating Default Port Numbers

Normally, CA DataMinder uses fixed port numbers but in rare cases you may need to reallocate these. There are three settings in each machine’s startup.properties file that you may need to change or add:

default.pingports

This setting controls the server port numbers reserved for the machine identification service. These are UDP ports.

Note: All your CA DataMinder installation need to have identical values in this setting.

default.rmiports

This setting controls the server port numbers assigned to the RMI registry. These are TCP ports.

default.serverports

This setting controls the TCP server port numbers allocated for CA services such as database management and data replication. To specify a set of numbers or a range, add this setting to startup.properties.

Allocating Port Numbers to default.pingports and default.rmiports

These two settings are handled in a similar way. In both cases, the default port numbers have been carefully chosen to minimize the risk that the default port numbers are already allocated to another service.

- The default.pingports setting controls the server port numbers reserved for the machine identification service. These are UDP ports. The default value for this setting is:

  ```
  default.pingports=56098 57098
  ```

  If the default.pingports setting is removed from the startup.properties file, CA DataMinder uses port number 56098. On some machines, this port number may already be used by another service so a CA DataMinder installation will fail unless you specify a new number.

  CA DataMinder will use the first available UDP port listed in the setting.

- The default.rmiports setting controls the server port numbers assigned to the RMI registry. These are TCP ports. The default value for this setting is:

  ```
  default.rmiports=56099 57099
  ```

  If the default.rmiports setting is removed from the startup.properties file, CA DataMinder will use port number 56099. On some machines, this port number may already be used by another service so a CA DataMinder installation will fail unless you specify a fallback number.

  CA DataMinder will use the first available TCP port listed in the setting.

For example, to reallocate port numbers 5000 to the machine identification service and 5001 to the RMI registry, copy these lines to the startup.properties file on each CA DataMinder machine:

```
default.pingports=5000

default.rmiports=5001
```

Reallocating default.serverports

This setting allows you to control the TCP server port numbers allocated for CA DataMinder services such as database management and data replication. It has a default setting of default.serverports=56100. This means port numbers are chosen sequentially from 56100 onwards.

Note: The ports specified by default.serverports must be configured on your firewall or router to allow **TCP**.

Use spaces to separate port numbers; use a hyphen suffix to specify the first number in a range. For example:

```
default.serverports=8000 8020 8030-
```

This allocates ports in the following order:

8000, 8020, 8030, 8031, 8032 and so on

How many server ports must you allocate?

This depends on the level of encryption used to protect network communications (this level is defined in the machine policy). But if you allocate a minimum of five server ports, this will be enough for all likely circumstances.

Configuration Changes to Support an Internet Firewall or Router

For details about CA DataMinder and the Windows XP SP2 firewall, see ‘Known Issues’ on Firewall configuration on Windows XP SP2 and 2003 SP1.

Allow Incoming and Outgoing Connections

To allow infrastructure and policy changes to be replicated from parent servers to child machines, and captured or imported events to be replicated from child machines to parent servers, your firewall or router must be configured to allow connections to be made on the specified ports in *both incoming and outgoing directions!*

Mobile Client Machines

To allow mobile client machines (that is, laptop users) to be part of your CA DataMinder installation, you may need to reconfigure your firewall or router to allow communication to each of the server ports specified in the startup.properties file. That is, the firewall or router configuration must allow use of the port numbers specified by these settings:

```
default.pingports
default.rmiports
default.serverports
```

See the previous sections for details about these settings.

Export, Import and Copy Policies

You can use the CA DataMinder Policy Editor to export and import policies to and from files, copy a policy from one account to another, and check policy versions. These operations are equally applicable to user policies, group policies and machine policies. For details, see the Administration console online help; search the index for ‘policy, utilities’.

This functionality is also available via a command line using wgnpol.exe, which is installed automatically using the server installation wizard and, for standalone client installations, the client installation wizard.

More information:

Installing a Standalone to Demonstrate Client Agents (see page 550)

Wgnpol.exe Command Line Syntax

You can use wgnpol.exe as a standalone command line utility. This can be useful to enable you to incorporate policy operations into scripts or batch files. In addition, you can use wgnpol.exe to import items from a CSV file to a setting in a policy.

Wgnpol.exe commands take the following format:

```
wgnpol <operation> <policy> [options]
```

Note: For quick usage instructions from a command prompt, run wgnpol -?.

More information:

wgnpol.exe - <Operation> (see page 568)
wgnpol.exe - <Policy> (see page 570)
wgnpol.exe - [Options] (see page 572)

wgnpol.exe - <Operation>

The following commands are supported in wgnpol.exe command line operations:

export

Exports a policy to a file. Specify the source policy plus the name and path of the target file.

import

Imports a policy from a file. Specify the target policy plus the name and path of the target file.

copy

Copies a policy from one account to another. Specify the source and target policies. The syntax is:

```
wgnpol copy <Source policy> <Target policy> [Options]
```

implist

Imports an external list to a specific policy list setting. The syntax is:

```
wgnpol implist
<Target policy>
<List setting>
<CSV file>
[Options]
```

where:

<Source policy> and <Target policy> are described on <Policy>.

<List setting> defines the target list setting. You can identify this setting either by its full XML path in the policy file, its full display path as shown in the Policy Editor, or, for certain outgoing e-mail control triggers, by a reserved keyword. See <List setting> below for more details.

<CSV file> is a comma delimited file containing the list items you want to import.

version

Displays the assigned version of the specified policy.

impsetting

Sets the value or attribute for a policy setting, or set the attribute for a policy folder. The syntax is:

```
wgnpol impsetting <Target policy> <Policynode> <CSV file>
```

where:

<Target policy> is described on <Policy>

<Policynode> defines a policy setting or policy folder (typically a trigger folder).

<CSV file> is a comma separated file containing the instructions for the policy change.

For a policy setting, the first item in the CSV file is the new value of the setting (for example, 'True', '10' or 'Always'). For a folder, this value can only be "enable" or "disable".

The second and subsequent items can be:
"enforce" or "unenforce"
"hide" or "show"
"reset"

Note: You cannot use impsetting to update policy list settings.

More information:

wgnpol.exe - <Policy> (see page 570)
wgnpol.exe - [Options] (see page 572)

wgnpol.exe - <Policy>

This is the account name for the user, group, or machine whose policy you want to operate on. If you have duplicate group names, you may also need to specify the policy path. (For example, your CA DataMinder installation may include two 'Sales' subgroups, one in the 'Asia' group and the other in the 'Europe' group.)

Policy Names and Paths

When specifying a policy name or path:

- Policy names and paths are not case-sensitive.
- If a machine, group or user name includes spaces, you must enclose the entire path in double quotes.
- Use / separators to specify a policy path, for example, "Sales/North America/Canada".
- Wgnpol.exe does not recognize the backslash \ as a policy path separator. But backslashes are allowed in user names, such as unipraxis\lyndasteel.

All exported policy files contain a comment tag:

```
<!--Name: YourPolicyName-->
```

You can configure wgnpol.exe to read the original policy name from this comment in an exported policy file. To do this, replace the policy name with an asterisk (*) in the command line. For example, to import the Sales.xml policy file back into its original group, run:

```
wgnpol import * -f Sales.xml
```

Users and Groups

When specifying the account name for a user or group:

- If the top level user group has *not* been renamed (by default, 'Users'), type usermaster to specify its policy. If it *has* been renamed (for example, to 'All Unipraxis users'), you specify its policy by typing its new name in the normal way.
- The policy path root is the management group of the CMS logon account. When you run wgnpol.exe, you must provide a logon account for the CMS. Wgnpol.exe takes the management group defined for this account as the root of the policy path.

Machines

To specify an account name for:

- Common client, use machinecommonclient
- Common gateway, use machinecommongateway
- CMS, use <CMS machine name>
- Utility machine, <utility machine name>

<List Setting> or <Policy node>

For implist or impsetting operations, you must specify the target policy setting that will receive the imported items. There are three alternative methods for specifying the target setting:

- **Reserved keywords:** For outgoing emails, the reserved keywords searchtext1, searchtext2 and searchtext3 refer respectively to the Included Search Text list setting in the Search Text 1, 2 and 3 control triggers.
- **Policy path based on folder and setting display names:** In the CA DataMinder Policy Editor screen, the status bar shows the policy path of the folder or setting currently selected, based on the folder and setting display names (the names shown in the policy tree, not the underlying XML node names). You can use these display names to specify the target setting for an import operation.
- **Policy path based on XML node names:** You can use the XML node names within policy files to specify a setting.

More information:

Wgnpol.exe Examples (see page 573)

wgnpol.exe - [Options]

The options supported in wgnpol.exe command line operations are listed in the table below.

-a

Specifies automatic CMS logon. That is, wgnpol.exe uses the cached credentials for a CA DataMinder user account (if available) to log on automatically to the CMS.

-c

For export operations only. This specifies that the complete policy is exported to a file, including all settings and folders inherited from a parent policy.

If this parameter is omitted, wgnpol.exe only exports a sparse policy. That is, it only exports those settings and folders that have been directly modified in the specified policy. Any settings and folders that were inherited unchanged from a parent policy are not exported to the target file.

-o

For import operations only. This specifies that complete policies (that is, those previously exported using -c) are included in the import operation.

-e

For implist operations only. Specifies that any duplicated list items are omitted from the import operation.

Only 'exact matches' are omitted. Duplicate matching is case-sensitive.

-f <file>

Specifies the path and file name to import or export to. Not valid for copy operations.

-m

For export operations only. This specifies that multiple policies for the selected group and all its subgroups are exported to a file. This includes the master policy.

File names for each subgroup policy include a policy ID in them to ensure that filenames remain constant across multiple exports.

-b

For export operations using -m only. Specifies that unmodified (blank) policies are not exported.

-s <server>

Specifies the name or IP address of the machine on which to operate. If you omit this option, wgnpol.exe defaults to the local machine.

-u <username>

Specifies the user name for the CA DataMinder account that you want to use to log on to the CMS. If you omit this option, wgnpol.exe prompts for a user name.

-p <password>

Specifies the password for the CA DataMinder account that you want to use to log on to the CMS. If you omit this option, wgnpol.exe prompts for a password.

-v

Specifies 'verbose' (fully detailed) output. This may be useful if, for example, you want to output operation details to a logfile.

-w <columns>

For implist operations only. Specifies how many data columns (the data 'width') you want to import from the source spreadsheet or CSV file.

For example, -w 3 specifies that the first three fields in each line of a CSV file (or the first three columns in a spreadsheet) are imported to the target list setting. In the CSV extract below, the three email addresses but not the "Equities Research" value will be imported into the policy list:

gary@xyz.com,barry@xyz.com,larry@xyz.com,"EquitiesResearch"

-? or -h

Displays usage instructions.

More information:

Wgnpol.exe Examples (see page 573)

Wgnpol.exe Examples

More information:

Export a Policy (see page 574)
Import a Policy (see page 575)
Import a Policy List (see page 575)
Copy a Policy (see page 575)
Check Version Details (see page 576)
Change a Policy Setting (see page 576)
Enable a Trigger (see page 576)

Export a Policy

Unless stated otherwise, these command line operation examples all export policies to a file called sales.xml.

- To export the policy for user unipraxis\lyndasteel from the machine CL-TAYLOR to a file called lyndasteel.xml, run:

  ```
  wgnpol export unipraxis\lyndasteel -s CL-TAYLOR -f lyndasteel.xml
  ```

- To export the policy for user unipraxis\lyndasteel, logging on to the CMS as CA DataMinder user unipraxis\spencerrimmel (whose password is 19apm77), to a file called lyndasteel.xml, run:

  ```
  wgnpol export unipraxis\lyndasteel
    -u "unipraxis\spencerrimmel"
    -p 19apm77 -f lyndasteel.xml
  ```

- If the top level user group has not been renamed, run this command to export the policy for the North American Sales group:

  ```
  wgnpol export "usermaster/north america/sales" -f sales.xml
  ```

- If the top level user group has been renamed (to 'All Unipraxis Users'), run this command to export the policy for the North American Sales group:

  ```
  wgnpol export "all unipraxis users
    /north america/sales" -f sales.xml
  ```

- If you log on to the CMS using an account that has /North America as its management group, run this command to export the policy for the North American Sales group:

  ```
  wgnpol export "north america/sales" -f sales.xml
  ```

  Or even just:

  ```
  wgnpol export "sales" -f sales.xml
  ```

- To export the policy for the gateway GW-CHICO to a file called new_gateway.xml, run:

  ```
  wgnpol export GW-CHICO -f new_gateway.xml
  ```

- To export all policies for user group 'All Unipraxis Users' and each of its subgroups to new files based on the name hierarchy.xml, run:

  ```
  wgnpol export "All Unipraxis Users" -f hierarchy.xml -m
  ```

- To export only the policies for user group 'All Unipraxis Users' and its subgroups, that have been edited, run:

  ```
  wgnpol export "All Unipraxis Users" -f hierarchy.xml -m -b
  ```

Import a Policy

To import the policy for user unipraxis\lyndasteel from a file called lyndasteel.xml, run:

```
wgnpol import unipraxis\lsteel -f lsteel.xml
```

- If the top level user group has not been renamed, run this command to import the policy for the North American Direct Marketing group from a file called new_dirmarketing.xml, run:

  ```
  wgnpol import "usermaster/north america
  /direct marketing" -f new_dirmarketing.xml
  ```

- To import a machine policy to the common gateway policy from a file called new_gateway.xml, run:

  ```
  wgnpol import machinecommongateway -f new_client.xml
  ```

Import a Policy List

Reserved keywords: To import the first two fields in each record of the source file keynames.csv to the Included Search Text setting in the Search Text 2 control trigger for the policy of user unipraxis\lyndasteel, run this command:

```
wgnpol implist unipraxis\lsteel searchtext2 "keynames.csv"
/w 2 /e /v /s CMS-HARDY /a
```

Policy path based on folder and setting display names: To import the first four fields in each record of the source file keywords.csv to the Included Search Text setting in the Search Text 1 control trigger for the policy of user unipraxis\lyndasteel, run this command:

```
wgnpol implist unipraxis\lsteel
  "Control/outgoing e-mails/control triggers
  /search text 1/included search text"
  "keywords.csv" /w 4 /e /v /s CMS-HARDY /a
```

Copy a Policy

If you log on to the CMS using an account that has
/North America as its management group, run this command to copy a group policy from North America/Sales to North America/Direct Marketing:

```
wgnpol copy "sales" "direct marketing"
```

- To copy a machine policy from UNI-TAYLOR to the common client policy, run:

  ```
  wgnpol copy UNI-TAYLOR machinecommonclient
  ```

- To copy a machine policy from GW-GROUCHO to GW-CHICO, logging on to the CMS as CA DataMinder user unipraxis\srimmel (whose password is 19apm77), run:

  ```
  wgnpol copy GW-GROUCHO GW-CHICO -u "unipraxis\srimmel" -p 19apm77
  ```

Check Version Details

To check the assigned version of the common client policy, run this command:

```
wgnpol version machinecommonclient
```

- To check the assigned version of policy for user unipraxis\lyndasteel, run this command:

  ```
  wgnpol version unipraxis\lyndasteel
  ```

- If the top level user group has been renamed (to 'All Unipraxis Users'), run this command to check the policy version for the North American Sales group:

  ```
  wgnpol version "all unipraxis users/north america/sales"
  ```

Change a Policy Setting

To change the Intervention setting in a control action to 'Warn' and enforce the change, run this command:

```
wgnpol impsetting unipraxis\lsteel
  "Control/Outgoing E-mails/Control Actions
  /Control Action 1/Intervention" intervention.csv
```

where intervention.csv has the following content:

```
warn,enforce
```

Enable a Trigger

To enable and hide a Control trigger, run this command:

```
wgnpol impsetting unipraxis\lsteel "Control/Outgoing E-mails
  /Control Triggers/Search Text 1" enable.csv
```

where enable.csv has the following content:

```
enable,hide
```

Replication Holding Cache

A replication holding cache is used to store captured or imported events that failed to replicate successfully to a parent server. If a parent server is unable to store a replicated event for any reason, it reports the failure back to the child machine which writes an entry for the 'failed' event to the replication holding cache.

By default, the child machine will attempt to resend events in the cache three times. After that, cached events remain in the cache until you manually delete them or reset the cache (that is, restore them to the main replication queue). If required, you can reset the retry limit—see below.

Likewise, as a safeguard against serious or persistent replication failures, the holding cache has a maximum event threshold which, if exceeded, causes the local infrastructure to be suspended. This means, for example, that an import operation is suspended if the CMS is unable to store events arriving from an Event Import server. By default, this threshold is set to 100 events. To reset this threshold, see below.

Note: When you reset the cache on suspended machines, the machines are automatically resumed.

More information:

Automatically Resume Suspended Machines on Scheduled Cache Reset (see page 580)

Cache Configuration

To override the default number of replication retries for failed events defaults, or to change the cache's maximum event threshold:

1. Add the following parameters to startup.properties on each child machine—find this file in the \system subfolder of the CA DataMinder installation folder. For example, to specify five retries and a 200 event cache limit, add these parameters:

   ```
   rep.retryThreshold=5
   rep.cacheSuspendThreshold=200
   ```

2. After editing startup.properties, restart the local CA DataMinder infrastructure.

More information:

Stopping and Restarting the Infrastructure (see page 531)

Managing the Cache

If replication failures occur repeatedly, you need to examine the events in the holding cache to determine and resolve the cause of the failures. You can then either reset or clear the cache.

To dump the cache contents

To dump the contents of the cache to a text file, run the following command. Cached events are output to CacheData.txt in the \System subfolder of the CA DataMinder installation folder.

```
wgninfra -exec wigan/infrastruct/replication/NetworkMonitor DumpHoldingCacheData
CacheData.txt
```

To clear the holding cache

To clear the replication cache (that is, delete the cached events), run the following command:

```
wgninfra -exec wigan/infrastruct/replication/NetworkMonitor ManageHoldingCacheData
delete <minID> <maxID>
```

where <minID> and <maxID> are event identifiers in the cache dump file. All events within the specified range of identifiers will be deleted. To delete all events in the cache, specify the identifiers for the top and bottom rows in the dump file (that is, the first and last events in the cache).

To reset the holding cache

After diagnosing and resolving a replication problem, you can reset the cache. That is, you can move cached events back to the main replication queue for re-sending to the parent server. CA DataMinder provides methods to do this manually and automatically.

More information:

Reset the Holding Cache (see page 579)

Reset the Holding Cache

When you reset the holding cache, cached events are moved back to the main replication queue for re-sending to the parent server. You typically reset the holding cache after diagnosing and resolving a replication problem.

You can manually reset the holding cache. You can also set up scheduled operations to automatically reset the cache at regular intervals. Note that scheduled cache resets on suspended machines will automatically resume the machine (if the suspension was cache-related).

Manually Reset the Holding Cache

To do this, run the following command:

```
wgninfra -exec wigan/infrastruct/replication/NetworkMonitor ManageHoldingCacheData
reset <minID> <maxID>
```

Where <minID> and <maxID> are event identifiers in the cache dump file. All events within the specified range of identifiers will be reset. To reset all events in the cache, specify the identifiers for the top and bottom rows in the dump file (that is, the first and last events in the cache).

You can also configure the local machine policy to automatically reset the holding cache at scheduled intervals—see the next section.

Automatically Reset the Holding Cache

In some situations, it is not practical to manually reset the holding cache on every machine that requires it. For this reason, CA DataMinder allows you to schedule operations to automatically reset the holding cache.

For example, a blob (Binary Large Object) file is quarantined by antivirus software on a client machine. Here, the maximum number of retries (fresh attempts to replicate the quarantined event) will be quickly used up. After diagnosing the problem by checking the log files on the parent server, the only intervention required from the CA DataMinder administrator is to remove the blob file from quarantine; the scheduled cache reset will then automatically return the blob file to the main replication queue.

To schedule automatic cache resets, you edit the following settings in the \Replication policy folder of the local machine policy:

- **Replication Holding Cache Reset Frequency**

  This setting specifies how often (in days) the replication holding cache is reset. Defaults to 3.

  Setting this value to zero enables the holding cache to be reset more than once per day with the Reset Time setting controlling the time (in minutes) between resets.

  Note: Setting both the Reset Time and Reset Frequency values to zero prevents holding cache entries from being reset.

- **Replication Holding Cache Reset Time**

  This setting is dependent on the value of the Reset Frequency. That is, if the Reset Frequency setting is:

  - zero, then this setting is the number of minutes between cache resets. Defaults to 180 (minutes).
  - non-zero, then this setting represents the number of minutes from midnight at which the replication holding cache is reset. For example, to schedule a reset for 9:00pm, enter 1260 (that is, 21 x 60). The reset will run at this time on each day that a reset is scheduled. Defaults to 180 (3:00am).

  Note: Setting both the Reset Time and Reset Frequency values to zero prevents holding cache entries from being reset.

Automatically Resume Suspended Machines on Scheduled Cache Reset

To streamline machine administration, if you reset the holding cache on a suspended child machine, the infrastructure on that machine is automatically resumed *if it was suspended because its maximum event threshold had been exceeded*.

This ensures, for example, that scheduled automatic cache resets are indeed fully automatic and require no manual intervention.

Note: If you manually reset the cache on a suspended machine, you must also manually resume the infrastructure on that machine.

CA DataMinder Installations on 64-bit Machines

CA DataMinder components are typically 32-bit applications. Currently, the exceptions are the Exchange 2007 and 2010 server agents, which are 64-bit applications (because Exchange Server 2007 and 2010 require a 64-bit operating system and hardware).

Exchange 2007 and 2010 Server Agent and PE Hub

You can only install the Exchange 2007 and 2010 server agents on 64-bit system. However, when you install the server agent, a 32-bit PE hub is also installed automatically. As a result, the Exchange server agent and the PE hub have different installation folders and registry keys.

However, there are no differences between the Exchange server agent and the PE hub in terms of log file location or performance counters.

Installation Folder

On 64-bit machines, 32-bit applications are installed to their own '32-bit' installation folder.

The 64-bit Exchange 2007 and 2010 server agents are installed to a subfolder below the \Program Files folder, while 32-bit components such as the PE hub are installed below the \Program Files (x86) folder.

Registry Keys

When 32-bit applications such as the PE hub are installed on 64-bit systems, any associated registry keys are created in their own '32-bit' subkey.

Registry values for the Exchange 2007 and 2010 server agents are created in the following subkey:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA DataMinder
    \CurrentVersion\Exchange
```

Registry values for the PE hub are created in the following subkey:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6332Node\ComputerAssociates\CA DataMinder
    \CurrentVersion\Policy Engine Hub
```

Log Files

On a 64-bit system, log files for the Exchange 2007 and 2010 server agents and the PE hub are all in the same location. Find these CA DataMinder logs in the \data\log subfolder of the Windows All Users profile.

Performance Counters

On a 64-bit system, all CA DataMinder performance counters, including counters for a 32-bit PE hub, are supported in the default 64-bit version of the Performance applet.

Integration with CA Identity Manager

CA DataMinder integrates directly with CA Identity Manager r12.5 SP3 or later.

This allows you to use CA Identity Manager to maintain your CA DataMinder user accounts. This ensures that your CA DataMinder accounts are consistent with identity-based policies used by other applications that rely on CA Identity Manager for provisioning user access rights and entitlements.

After CA Identity Manager has been set up to connect to CA DataMinder, any changes to user accounts in CA Identity Manager can be imported to the corresponding CA DataMinder accounts on your CMS.

For instructions on integrating CA DataMinder with CA Identity Manager, see the *CA Identity Manager Connectors Guide*, available from CA Support: http://ca.com/support

Chapter 25: Applying Hotfixes

This chapter describes how to apply CA DataMinder hotfix patch packages directly and how to update an administrative installation source image to include the fixes contained within a patch package

This section contains the following topics:

Patch a Local Installation

To patch the local CA DataMinder machine, do either of the following:

- Double-click the .msp file in Windows Explorer. This launches the CA DataMinder installation wizard. Click Update to apply the .msp file.
- Run this command:

```
msiexec /p <Full path to MSP> REINSTALL=ALL REINSTALLMODE=omus
```

Where <Full path to MSP> specifies the full path and name of the .msp patch package. The /p option indicates a patch operation.

Create a New, Patched Administrative Installation Source Image

To create a new administrative installation source image that incorporates the fixes in the hotfix patch package:

1. Create a new source image by running the following command:

   ```
   msiexec /a <Full path to original MSI>
   ```

   Where the /a option creates an administrative installation source image and <Full path to original MSI> specifies the full path and name of the source image .msi file on the original CA DataMinder distribution media. This command launches the CA DataMinder installation wizard, which prompts you to specify the network location of the new source image.

2. To patch the new source image, run the command described in Patch an Existing Administrative Installation Source Image.

Patch an Existing Administrative Installation Source Image

To update an existing administrative installation source image to incorporate the fixes in the hotfix patch package, run this command:

```
msiexec /a <Full path to MSI> /p <Full path to HF MSP>
```

Where <Full path to MSI> and <Full path to HF MSP> specify respectively the full path and name of the existing administrative installation source image .msi file and the Hotfix .msp patch package. The /a option updates an administrative installation source image; the /p option indicates a patch operation.

Force a Local Repair from a Patched Source Image

To update the local CA DataMinder machine using a patched source image, run this command:

```
msiexec /i <Full path to MSI> REINSTALL=ALL REINSTALLMODE=vomus
```

Where <Full path to MSI> specifies the full path and name of the .msi file, that is, the patched source image.

Uninstall a Hotfix Patch Package

You can uninstall hotfix patch packages from a command line or using an applet such as 'Add or Remove Programs' and 'Programs and Features'.

Note: Hotfix patch packages can only be uninstalled using Windows Installer 3.x or later. You cannot uninstall hotfix patch packages using Windows Installer 2.x.

- **Add or Remove Programs:** Only supported for Windows 2003 and Windows XP. In the Add or Remove Programs applet, select the Show updates check box. This lists all the hotfix patch packages installed on the local system. Select the relevant CA DataMinder hotfix patch package and click the Remove button.
- **Command line:** To remove an uninstallable patch package, run the following command:

  ```
  msiexec /i <GUID> MSIPATCHREMOVE=<Full path to MSP> /qb
  ```

  Where:

  <GUID>

  Specifies the CA DataMinder product code (also called the GUID) for the *product package* that was patched with the hotfix.

  For example, if you are uninstalling a *reports hotfix patch package*, you must specify the GUID for the *reports installation package*.

  You must enclose the GUID in braces { } when running the uninstall command.

  GUIDs vary by installation package and product version. See the following section for lists of supported GUIDs.

  <Full path to MSP>

  Specifies the full path and name of the hotfix patch package; the /qb option specifies a progress bar and simple confirmation dialog when the uninstall is complete.

  Note: To remove multiple hotfix patch packages, set MSIPATCHREMOVE to a semicolon-separated list of <Full path to MSP> values.

CA DataMinder Product Codes

The CA DataMinder product codes vary by product version. They are listed below:

CA DataMinder r14.5 Product Codes

Package Name	Product Code
CCS_x64.msi	{6552A17F-4093-48BC-8F63-9A271B406B78}
CCS_Prescan.msi	{5E813054-0140-4360-B527-AB6CDFC2DAE3}
CCS_Prescan_x64.msi	{B095ABD3-7D39-4C3D-93D6-9AC9E87D567E}
Client.msi	{214D2AE5-8A79-4431-85DF-34AC26F38217}
Client_x64.msi	{2DF6C2E8-7B43-41CD-A4A3-A3C7B52F4C21}
Content.msi	{1E17ED37-EE7A-4DEA-A1A6-5E577178D378}
Integration.msi	{CEBD3470-0EDD-437D-A605-F057B28187AE}
Integration_x64.msi	{729BEF07-C515-470B-9EC4-BDE410F0CF11}
Policies.msi	{D1E5557B-ECF4-4B69-9ED0-2635FAF15822}
Reports.msi	{66CA9B0D-D2CE-41A6-8CEE-702E303C4C17}
Server.msi	{020FA7F3-0B55-4F22-8052-DAD0B01FCE2B}
Server_x64.msi	{FD03FB1F-7746-4F0B-BD7F-94EC9D4652D6}
Sevddup.msi	{619DCC5F-26B0-4B20-9245-553696BFFF0D}
UA.msi	{69C5F15A-1E8E-452D-8AFF-E5D18009F13E}
Web.msi	{19BE9257-9AE8-48FD-944E-B6B2F77AF7DC}

CA DataMinder r14.1 Product Codes

Package Name	Product Code
CCS_x64.msi	{EC6F90E7-721B-4D3B-8017-457AC4848BAD}
CCS_Prescan.msi	{DF9E2563-6E97-4C72-9DB1-D35ED5A531C4}
CCS_Prescan_x64.msi	{A3821FC0-6422-4D56-862C-B225F794DF90}
Client.msi	{D7DA635F-8611-482B-B92F-FC7C6ACBE8D6}
Client_x64.msi	{ACEE3D21-3376-4D85-9429-7A981898656D}
Content.msi	{FD6FEC74-5D3D-4217-9475-1EFED99E4E7C}
Integration.msi	{93E9C36B-E656-4AD5-A1AC-ACDFBBE1EDB3}
Integration_x64.msi	{00843DD6-8992-4DB5-9581-E4AA5EFE9117}
Policies.msi	{D938FE2A-2783-4EFA-B219-54A1AFDB9701}

Package Name	Product Code
Reports.msi	{211925EB-92C3-4FC6-B1CB-1ABE10585CFC}
Server.msi	{020FA7F3-0B55-4F22-8052-DAD0B01FCE2B}
Server_x64.msi	{E788813A-2F42-45E6-AFC5-D13A178F45A7}
Sevddup.msi	{68D647D8-517F-4F98-B762-4C50304D6328}
UA.msi	{C77751CD-6D96-4EAA-B1EB-8A285CBEBF4E}
Web.msi	{DB63A972-57B6-4724-9F36-30B3ED1A5DB5}

CA DLP r14.0 Product Codes

Package Name	Product Code
CAMM_Config.msi	{AC353FC0-8091-476D-879B-48201D171D2D}
CCS_x64.msi	{C429EE98-B7E4-402A-9244-8E9927E9BF8C}
CCS_Prescan.msi	{880426D2-BB07-456B-B0DB-F25F690E902C}
CCS_Prescan_x64.msi	{C424CAF2-F53F-4F9D-B7AC-DD433B6FD245}
Client.msi	{0E16B697-35D8-43B4-A235-FEB5F26ECB67}
Client_x64.msi	{2FE9BC0A-43DB-40FB-B9A6-66D58AD069A9}
Content_IDOL.msi	{B63B0561-E137-4B80-9236-2181D1218CBF}
Integration.msi	{76871A9E-B150-4E9C-B629-9D09DCB76EE4}
Integration_x64.msi	{182915C0-6B3E-4D2F-B8F8-17FF9EE7C4CE}
Policies.msi	{FF309C58-6B30-4F08-A600-D7A17976D59D}
Reports.msi	{9C154C57-DF5C-42F7-9E90-1FAF0DA9016A}
Server.msi	{A72164E5-5F5E-44ED-AD1B-6C0758DBC55A}
Sevddup.msi	{7FE0DD74-5C32-4B17-8A75-A5CDB85A1322}
UA.msi	{AF4AE9B6-2719-4642-BFB6-9AB2ECD9EC52}
Web.msi	{22CC12E0-692B-4100-8ACD-F6D3727D2F11}

CA DLP r12.5 Product Codes

Package Name	Product Code
CAMM_Config.msi	{53CA3047-3F54-4D67-82C7-D92AE35E2024}

Package Name	Product Code
Client.msi	{994F6DA5-8603-4C22-BAF8-7AFA3E91B98B}
Content.msi	{F99857CC-6010-44DE-91CD-8F3C0CE86BE3}
FastStartConfig_DLP.msi	{F40DA5BC-9C26-423D-B37E-43A9EFB823D8}
Integration.msi	{3F690152-B52B-4214-A4EF-D2B58616438A}
Integration_x64.msi	{59CA1605-2F6D-4C17-8927-260C3AD8D136}
Reports.msi	{9E0B740F-41D3-4BE0-B3A1-595F5CEDABC3}
Server.msi	{F294914D-190D-42EC-9C32-EE1743CFE177}
Sevddup.msi	{469EDBF9-86E5-4330-A4D8-2AF68D87A968}
UA.msi	{8C081DC6-B6D9-4B88-A2B2-9403301D9B21}
Web.msi	{D4E4CA78-B2CE-420E-9CBE-4342AF50F76F}

CA DLP r12.0 Product Codes

Package Name	Product Code
CAMM_Config.msi	{0A14C0FA-8462-4988-845D-109DD8FB31D7}
Client.msi	{85FA4587-35C5-4C5C-AC40-2BAEEEFA51EE}
FastStartConfig_DLP.msi	{A1985217-29EC-4AB6-B305-EC3ACA0C242F}
Integration.msi	{2669CDA3-5479-4FE9-87C8-0608D5042A55}
Integration_x64.msi	{9796B92C-2927-41AC-94FE-3AE6DB502DC8}
Reports.msi	{A9F3BC1C-96A0-4CCD-A409-2D3DC642F51A}
Server.msi	{8973D0EF-34A2-45EE-BAB7-DA20A08765A7}
UA.msi	{50BBB163-7A17-4F15-BE87-E388907EFBEF}
Web.msi	{DF5DBFF5-9565-470B-9963-964D9069A3D7}

Orchestria APM 6.0

Package Name	Product Code
Client.msi	{A889228E-9DEE-4102-9A77-A4DF89A8689A}
Content.msi	{831799BB-93A8-4994-931D-CBA924989CBC}
FastStartConfig_DLP.msi	{C86B4C09-4469-48A2-8F5B-2BB3290FCA3B}

Package Name	Product Code
Integration.msi	{49DD405A-9D3D-40F9-BEF4-DEDDCCD2C5D2}
Integration_x64.msi	{94FD0328-5120-432B-AE46-F30A312AEA95}
Server.msi	{07535CA7-3475-498A-AB45-0A0A772D4645}
UA.msi	{6D4036E7-EF34-4AB2-A7CF-F228F618EC5E}
Web.msi	{7D5C9BDE-5629-41D2-B7AF-C3B7C5C9E63B}

Known Issues

Performance counters may be lost

After applying, removing or repairing of any Hotfix or Service Pack, performance counters may be lost. We therefore recommend that you reboot the system after these operations.

Policy reset may occur

After applying, removing or repairing of any Hotfix or Service Pack, a policy reset may occur. If this happens, you will need to re-install your license file and reapply all extended policies.

Infrastructure-dependent services may fail to restart

After applying, removing or repairing of any Hotfix or Service Pack, services dependent on the CA DataMinder infrastructure service may have been stopped and subsequently failed to restart. We therefore recommend that you reboot the system after these operations.

Restart Microsoft IIS before applying the hotfix

When patching a CA DataMinder component other than the iConsole when the iConsole is installed, we recommend that you restart Microsoft IIS before applying the hotfix to prevent reboot requests after you apply the hotfix.

Chapter 26: Known issues

This section describes various known deployment issues. These include deployment issues, troubleshooting email server agents, support for Far Eastern characters, and IM import issues.

If you need help, contact CA Technical Support (see page 21).

This section contains the following topics:

General Deployment

More information:

Firewall Configuration on Endpoints

The CA DataMinder installation wizard automatically registers the CA DataMinder infrastructure as a firewall exception. This enables data, including policy updates, to replicate unhindered through the firewall between CA DataMinder endpoints and servers.

However, this automatic configuration requires the firewall to allow the CA DataMinder exceptions:

- On Windows XP endpoints and Windows 2003 servers hosting CA DataMinder, you must turn **off** the ‘Don’t allow exceptions’ setting. Find this setting on the General tab of the Windows Firewall applet.
- On Windows Vista and 7 endpoints and Windows 2008 servers hosting CA DataMinder, you must turn **off** the 'Block all incoming connections' setting. Find this setting in the Domain Networks section of the Windows Firewall applet.

Important! If either of these settings is **on**, the Windows Firewall allows no firewall exceptions, including the CA DataMinder infrastructure. This means that the endpoint computer cannot contact its parent server. As a consequence, the endpoint is unable to receive any user or machine policies. This effectively paralyzes any CA DataMinder agents on the endpoint computer so they are unable to monitor, capture or control user activity.

Stopping or Starting the Infrastructure without Rebooting

If you need to stop or restart the CA DataMinder infrastructure, you can do so using the wgninfra service.

To stop the infrastructure

Run this command:

```
net stop wgninfra
```

To restart infrastructure

Run this command:

```
net start wgninfra
```

Do Not Install to Encrypted Folders

Do not install CA DataMinder to an encrypted folder or file system. This also applies to your CMS data folder (see step 6 of Installing a CMS, gateway or utility machine).

Note: For SQL Server users, you must also ensure that the \Data folder is not compressed. See the *Database Guide*; search the index for ‘data folder’.

CFSA Can Prevent BitLocker From Encrypting USB Devices

The Client File System Agent (CFSA) can affect the operation of the BitLocker To Go encryption feature on endpoint computers.

If the CFSA is installed on an endpoint computer and configured to apply policy to files being copied to removable devices (such as USB drives or SD cards), BitLocker cannot initialize removable devices for encryption. That is, it cannot give these devices the "lockdown treatment". This is because the BitLocker initialization process is denied write access to the device by the CFSA.

Note: This problem only occurs if the CFSA is explicitly configured to apply policy to removable devices. Also, if a removable device has been initialized by BitLocker running on a different computer, the device can used on any endpoint computer hosting the CFSA, even if the CFSA is configured to apply policy to removable devices.

RDM, EAS and Windows 2003

This is a test you can run to confirm that CA DataMinder console users can retrieve e-mails archived in the Zantaz Exchange Archive Solution (EAS).

1. On the Remote Data Manager (RDM) server, ensure that you are logged on to Windows with the same logon account as that used by the CA DataMinder infrastructure service (wgninfra.exe).
2. Open Internet Explorer and use the following URL to browse to EAS:

   ```
   http://<EASWebServer>/EAS_APP
     /easweb.dll?ServerGetMsg&msgid=
     <MSGID>&serverID=<ServerID>&compressed=0
   ```

   where:

   <EASWebServer>

   Is the name or IP address of the EAS IIS server.

   <MSGID>

   Is the numerical EAS message ID to be retrieved.

   <ServerID>

   Is the numerical ID of the EAS server containing that message.

   If Internet Explorer can retrieve a .MSG file from EAS in this way without being prompted for Windows credentials, then it will also be possible for RDM to do so.

Laptop Users and Dial-up Connections

Laptop users who normally connect to the CMS using a dial-up connection may be prompted for their dial-up connection details if they subsequently connect to the CMS over a LAN. To prevent the Dial-up Connection dialog from appearing, laptop users must edit the dial-up settings in their Internet Explorer properties. To do this, they must:

1. Open Internet Options in the Control Panel.
2. Go to the Connections tab.
3. In the Dial-up Settings list, choose ‘Dial whenever a network connection is present’.

Email Server Agents

Failure to Generate Email Events

If the Exchange server agent fails to generate email events, the following questions can help you to diagnose the problem.

Has a Policy Engine Been Activated?

Use Process Explorer (from www.sysinternals.com) to ensure that a wgnpesv.exe service is running on the policy engine host machine and using the correct user account. If it is not:

- Has the policy engine service been set to run as the correct named user?
- Have the same credentials been assigned to the policy engine hub?
- Have the policy engines been defined and configured correctly in the hub registry?
- Does the NT System or Application event log on either machine contain relevant information?

Check the policy engine hub log file.

Has a Policy Engine Stopped Working?

To force the policy engine hub to disconnect and reconnect to a policy engine, remove the policy engine from the ActivePolicyEngines registry value then add it again. This is a policy engine hub registry value.

Note: For the Exchange server agent, if you need to restart the policy engine hub, you must first stop the Internet Services.

Is Email Address Mapping Set Up Correctly?

- Check the machine policy settings on the policy engine host machine.
- Check that the UserSpecificAddrPattern registry value is correctly configured with a pattern match for the email addresses you want to detect. This is a policy engine hub registry value.

Is wgnemno.dll Registered as an Add-in?

(Domino sever agents only) Confirm that the following line has been added to notes.ini on the Domino host server. Find this file in the same folder as the Domino executables, typically \Lotus\Domino.

```
EXTMGR_ADDINS=wgnemno.dll
```

If this line is not present, manually add this line and restart Domino.

Unable to Expand Distribution Lists with Hidden Membership

(Exchange sever agents only) The fix for enabling a policy engine to expand hidden distribution lists depends on which version of Exchange Server you are using.

Fix for Exchange Server 2003 or Later

If a distribution list in Exchange's Global Address list has been configured to hide list members, policy engines will be unable to expand these distribution lists to identify and apply policy triggers to individual recipients.

If you want policy engines to expand these distribution lists, you must ensure that the policy engine service account belongs to a group whose members have the necessary permissions to expand 'hidden membership' distribution lists.

More information:

Specify a PE Domain User (see page 340)

Multiple Notifications in Response to a Single Email

When an email activates a control trigger, the sender of the email can sometimes receive multiple warning or notification messages. This happens when CA DataMinder applies policy to multiple copies of the original email.

This can occur if an email is sent, via an email server with no server agent installed. In this situation, the following scenarios can result in multiple notification messages being sent:

- The recipients' mailboxes are hosted on multiple email servers. Each of these email servers has a server agent installed. Policy is applied to the email each time it reaches the server agent on one of these subsequent email servers. This example is shown below.
- Two or more recipients have mailboxes hosted on the same email server. The 'sending' email server sends copies of the email to each recipient's mailbox. The server agent on the email server hosting these mailboxes applies policy to each recipient's copy of the email.

The two scenarios above are the most likely cause of multiple notification messages. But other configurations can also cause multiple notifications. For example, multiple copies of the original email may simply be routed through an email server hosting the server agent. Likewise, a non-standard journaling setup may result in the server agent failing to recognize a journaled email and applying policy to that email.

To avoid any of the above scenarios, we recommend that you install server agents on all email servers in the CA DataMinder enterprise.

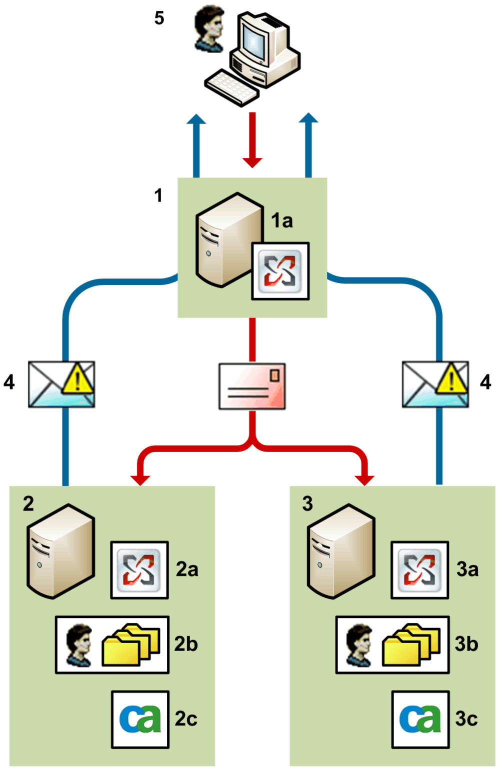

Multiple notification messages example

1 A user sends an email and it transits through the email server (**1a**), undetected. The recipients' mailboxes (**2b** and **3b**) are on two separate email servers (**2a** and **3a**).

2 and **3** An instance of the email arrives at each email server. Each instance is detected by an email server agent (**2c** and **3c**) and triggers a warning notification.

4 Each email server agent then sends a warning message to the sender.

5 The user receives multiple notification messages in response to a single email.

Event Import

Importing from Exchange Requires MAPI Client and CDO 1.2.1 Component

When you import emails directly from mailboxes on an Exchange server, the Event Import host server requires the 'Messaging API and Collaboration Data Objects 1.2.1' component.

You can download this component from the Microsoft Web site. Before you install it on the Event Import host machine, verify that Microsoft Outlook is *not* also installed; if Outlook is installed, you must uninstall it.

Note: Event Import does not require the 'Messaging API and Collaboration Data Objects 1.2.1' component when importing .pst or .msg files. To import these files, Event Import must have Microsoft Outlook installed (32-bit only).

Special Requirement for Windows Server 2003 x64 Machines

When running an Event Import job on a Windows Server 2003 x64 Edition machine, the host machine must be using the Exchange Server 2007 "Messaging API and Collaboration Data Objects 1.2.1" (MAPI client and CDO). If the host machine is using the Outlook 2003 MAPI client, the import job will fail to run.

Note: This problem does not affect Exchange import jobs running on 32-bit Windows Server 2003 machines. However, even on these machines we recommend that the host machine uses the Exchange Server 2007 "MAPI client and CDO" (unless you are importing Unicode .pst files, in which case it is acceptable to use the Outlook 2003 MAPI client).

Cannot Access an Exchange Mailbox

When importing from an Exchange mailbox, the import operation may occasionally fail and report a ‘no user identifier available’ error. This is because the user account that the CA DataMinder Event Import service is logging on as does not have access to the mailbox. To check whether this user has the necessary permissions to open the mailbox:

1. Log on to Microsoft Outlook using the same user account as currently used by the Event Import service.
2. In Outlook, try to open the mailbox you want to import.
3. If this is unsuccessful, you may need to either:
   - Assign the necessary permissions to the current user.
   - Change the logon user that the service is using to be the same user account as the mailbox.
4. If step 2 is successful, then the access rights of the current user are not the problem in this case.

Imported email Timestamps Are Truncated

When importing e-mails using Outlook 2003 in Cached Exchange mode, Microsoft truncates the timestamps for e-mail events. That is, the timestamps are rounded down to the nearest minute. CA DataMinder then uses this timestamp to record when the event was captured, as shown in the table below:

Actual time sent	**Capture date**
2005-05-27 16:29:09.201	2005-05-27 16:29:00
2005-05-27 16:29:30.413	2005-05-27 16:29:00
2005-05-27 16:29:51.952	2005-05-27 16:29:00

These examples show that each e-mail was actually sent at a different time, but the capture date for each event is identical. The inaccuracy of this data could have an impact on other CA DataMinder features, such as filtering by capture date.

NSF Import and ‘End MIME to CD Conversion’ Messages

Before using wgnimp.exe from a command to import .NSF files, you need to adjust the logging level for the local Notes client. Failure to do so can cause multiple redundant ‘End MIME to CD Conversion’ messages to be output to screen when you run the import operation. To adjust the Notes logging level, add the following line to notes.ini on the Event Import host server:.

```
converter_log_level=10
```

After adding this line, you will need to restart any local Notes application and any local .NSF Event Import operations.

Cannot Import Unparented Emails from a Notes Database

When importing emails from Lotus Notes NSF files, be aware that the ‘source folder’ import parameters will fail to detect ‘unparented’ emails. These are emails which are **not** contained in a folder within a Notes database (that is, a NSF file). This normally only occurs if the emails are stored in a database that is not based on the standard email template.

The import parameters that require the source emails to be contained in folders are:

- NSF.FolderName
- NSF.FailedMessageFolder

In practice, this mainly affects emails that have been moved to the ‘failure’ folder following an unsuccessful import attempt (where the failure folder is defined by the NSF.FailedMessageFolder parameter). It could also affect emails in an NSF file with a flat folder structure. For example, this could apply to a subset of emails that have moved to a different NSF file for reporting or administrative purposes.

You can still import unparented emails, but the parameters specified above will not work. Instead, you can only use the NSF.DominoFileName parameter to locate the source emails.

Far Eastern Characters

Note the limitations on using Far Eastern characters in installation paths and computer names, and displaying these characters in CA DataMinder consoles.

Do Not Use Far Eastern Characters in Installation Paths

CA DataMinder cannot handle installation paths that contain Far Eastern characters. If you install Content Services components to a non-default location, *the target path must not include folders whose names contain Far Eastern characters!*

Computer Names with Far Eastern Characters

CA DataMinder does not support computers with names that contain Far Eastern characters. If you try to install CA DataMinder on these computers, the installation fails. However, you can install CA DataMinder on computers running Far Eastern versions of Windows if the computer name only contains Roman-based ('English') characters.

Far Eastern Characters in Event MetaData Do Not Display in iConsole

Symptom:

If an iConsole search returns an event containing Unicode event metadata (such as Chinese, Japanese, Korean and Russian characters), this metadata does not display correctly. Instead, these characters are replaced with question marks. For example, if the trigger name or search text was defined using Korean characters in the user policy, these items display as '????' strings in the iConsole search results page.

In addition, if a reviewer searches for events by Unicode trigger name, the search fails to return events where the Unicode metadata was not captured. For example, an email trigger has a Korean trigger name. When this trigger fires on a computer running an English version of Windows, the email is captured but the trigger name is lost. Consequently, when a reviewer searches for emails associated with this trigger name, the search fails to find the email captured on this computer.

Note: This problem is rare. It primarily affects events captured on endpoint computers running a Western version of Windows and where policy triggers include settings that use Unicode characters, such as Chinese, Japanese, Korean or Russian text.

Solution:

Implement Unicode support on all CA DataMinder endpoint computers running non-double byte character set versions of Windows (that is, non-DBCS Windows) and which are likely to capture events containing Unicode characters. For example, this applies to Japanese email trigger details captured on a computer running an English version of Windows.

To implement Unicode support:

1. Stop the 'CA DataMinder infrastructure' service. From a command prompt, run:
   ```
   net stop wgninfra
   ```
2. Edit the startup.properties file. Find this file in the \system subfolder of the CA DataMinder installation folder.
3. Open this file and add the following line to the [Database] section:
   ```
   db.charset=UTF-8
   ```
4. Restart the CA DataMinder infrastructure service. From a command prompt, run:
   ```
   net start wgninfra
   ```

Important! Do *not* make this change to startup.properties on CA DataMinder computers running Japanese, Korean or Chinese versions of Windows.

iConsole

More information:

HTTP 404 Error When Browsing to the iConsole URL (see page 602)
Unable to Download or Forward Original .msg File (see page 603)
Unable to Send Audit Emails (see page 604)
Problem with Multiple iConsoles on the Same Client Machine (see page 605)
Display Problems Due to IE Enhanced Security Configuration (see page 605)
Dashboard Event Totals Seem Wrong After Drilling into Report or Chart (see page 606)

HTTP 404 Error When Browsing to the iConsole URL

If users get an HTTP 404 error (‘page cannot be found’) when trying to browse to iConsole, you need to check that the required IIS Web Service extensions are registered and allowed:

- ASP.NET v1.1.4322
- CA iConsole

If either have been prohibited, you must set their status to ‘Allowed’ to enable users to browse to the iConsole.

Unable to Download or Forward Original .msg File

Applies to Microsoft IIS 5.x only.

If the iConsole application server uses IIS 5.x, attempts by reviewers to download emails can result in errors:

- When using the Download email feature to download an event’s original .msg file, the user gets a ‘Unable to retrieve mail message’ error.
- When trying to send an audit email with the original message attached, the user gets a ‘Unable to send mail’ error.

The reason for both error messages is the same. In order for IIS 5.x to use MAPI services on the application server, the local IWAM_<machinename> user must have local administrator rights.

To assign these rights

1. Open the Computer Management applet. This is available from the My Computer applet.
2. Browse to Local Users and Groups:
   a. Open the Groups folder, right-click Administrator and choose Properties.
   b. In the Properties dialog, add the user IWAM_<machinename> to the Administrator group.
3. Restart the IIS service.

Unable to Send Audit Emails

Applies to Microsoft IIS 5.x only.

If the iConsole application server uses IIS 5.x, users can encounter the following error messages when trying to send audit emails.

```
"SMTP server not running / Connection not configured correctly; The transport failed to connect to the server."
```

```
"Relay not configured correctly; The server rejected one or more recipient addresses. The server response was: 550 5.7.1. Unable to relay for lynda.steel@unipraxis.com."
```

These errors occur if the SMTP server is not running or the connection is not configured correctly. F

Typically, an audit email contains the original .msg file as an attachment. If this is the case, and the email failed to send, users will encounter the following error message on trying to resend the same email:

```
"System.Reflection.TargetInvocationException:"
```

The reason for this is because when the iConsole sends an email with an attachment, it creates a temporary file for the attachment which is then deleted when the email is sent. If the SMTP server is not running or the connection is not configured correctly, the temporary file is not deleted and the email cannot be sent.

To enable the iConsole to send audit emails with or without the original .msg file, you must first ensure that the SMTP server is correctly configured and running and then restart IIS.

More information:

Set Up SMTP Email (see page 143)

Problem with Multiple iConsoles on the Same Client Machine

A configuration setting in Internet Explorer can cause unexpected behavior if multiple iConsole instances are running simultaneously on the same client machine. Specifically, a local registry value can cause all iConsole instances to share the same browser session.

For example, if you have two iConsoles open in separate windows on the same machine, each connecting to a different CMS, this problem can inadvertently cause the second iConsole to reconnect to the CMS specified in the first iConsole.

If an iConsole user experiences this sort of problem, we recommend that you check the registry on the relevant host machine:

1. Locate the following registry key on the iConsole host machine (that is, the browser host machine:

   ```
   HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion
       \Explorer\BrowseNewProcess
   ```

2. Within this \BrowseNewProcess subkey, you need to edit the following value:

   ```
   BrowseNewProcess
   ```

   You must ensure this registry value is set to Yes for correct iConsole operation.

   Note: By default, this registry value is not present in the registry, but its operation defaults to Yes. Therefore you do not need to add this value if it is absent; you only need to set it to Yes if it is present and already set to No.

Display Problems Due to IE Enhanced Security Configuration

When running the iConsole from a machine that has Internet Explorer Enhanced Security Configuration (ESC) installed, you must add 'about:blank' to the list of trusted sites for the Local Intranet. Alternatively, we recommend that you disable ESC. Failure to implement these workarounds may result in some iConsole pop-up dialogs failing to appear.

Dashboard Event Totals Seem Wrong After Drilling into Report or Chart

In certain conditions, if a reviewer drills down into a report or dashboard chart, the number of results does not match the event total shown in the report or chart. This apparent disparity can occur if further events have been captured or reviewed in the intervening period since the snapshot totals were last calculated.

Snapshot totals are recalculated each time a data warehousing job runs. By default, these jobs run every hour. Consequently, snapshots (such as 'total unreviewed events') reflect the total number of events *at the time when the job ran*.

If the number of events (or incidents) in the CMS database rises or falls before the next data warehousing job runs (for example, because a manager reviews some previously unreviewed events), then the snapshot total shown in the report or dashboard will no longer tally with the actual number of underlying events in the CMS database. If a reviewer were to drill down into the report or dashboard at this point, they would see an apparent disparity in the number of events.

Example

Consider this timeline:

15.00 PM	A data warehousing jobs finds 100 unreviewed events and adds them to the data warehouse.
15.15 PM	A manager refreshes their dashboard. The snapshot total for Unreviewed Events is 100.
15.16 PM	The same manager drills down into the dashboard to see the underlying events. The Search Results screen does indeed find 100 unreviewed events.
15.30 PM	A reviewer audits 25 of the unreviewed events in the iConsole.
15.45 PM	The manager refreshes their dashboard; the snapshot total for Unreviewed Events is still 100. This is because there has been no new data warehousing job since 15.00 PM.
15.46 PM	The manager drills down into the dashboard again. But this time, the Search Results screen only finds 75 unreviewed events!
16.00 PM	The next data warehousing job runs and finds 75 unreviewed events. The snapshot total and number of underlying events are back in sync.

Note: Such potential disparities only affect snapshots of event counts based on audit status. They cannot occur with snapshots based on non-changing event attributes such as events counts by policy.

IM Import

More information:

Bloomberg IM Dump Files Are in US ASCII Format

Bloomberg IM dump files are in US ASCII format, but can contain data captured from terminals in other languages such as German and Japanese. In this cases, archived IM conversations cannot be restored to their original language because no code page information is available to IMFrontEnd.exe.

Identifying Conversation Participants

If IM Import detects inconsistent formatting in Bloomberg IM dump files, it attempts to handle these inconsistencies but inevitably some minor anomalous formatting can occur. For example, Event Import may be unable to identify the full list of participants and so fail to create conversation events for these users.

Timestamps in IB Unified Dump Files Are in EST

All events in the CMS database are stored in UTC time, but IB Unified dump files record events in EST. However, IMFrontEnd.exe is unable to reliably and universally convert EST times to UTC because the dump file data does not contain time zone information to enable it to calculate daylight saving adjustments. Instead, IMFrontEnd.exe assumes that it is running in the same time zone as the Bloomberg server that generated the IB Unified dump file and converts EST times to UTC on this basis.

Format Requirement for Imported Attachments

The specification for the IB Unified format does not explicitly refer to attachments, but our analysis has determined that attachments are referenced by an 'Attachment:' entry in the message header. IM Import uses this header entry to import attachments. However, because this header entry is undocumented, it could change in the future, possibly preventing attachment files from being imported into CA DataMinder.

Increase Size of .CNV Cache to Improve Import Performance

For IM certain dump files, IMFrontEnd.exe performance can be significantly improved by increasing the size of the .CNV file cache. To do this, you need to raise the value of the MaxFiles parameter.

As IMFrontEnd.exe processes a dump file, it needs to continually open and write to the relevant .CNV file. But for certain dump file formats, the internal structure of the dump file requires that individual .CNV files need to be continually reopened and updated, rather than being generated and closed in a single operation.

For these dump file formats, the comments that collectively make up an individual IM conversation, and which must be written to a single .CNV file, are not aggregated but occur throughout the dump file. This means that as IMFrontEnd.exe sequentially processes the dump file, it needs to continually open and close the relevant .CNV files. These file operations account for a high proportion of the total dump file processing time, so increasing the number of cached .CNV files can substantially reduce the processing overhead.

‘More recipients’ Entries Are Ignored

Some messages include the text 'Note: More recipients' in their header. These entries are undefined in the format specification and currently disregarded by IM Import. Therefore, they do not appear in imported IM conversations.

Anomalous Join and Leave Chat Room Actions

When reviewing an IM conversation in the Data Management console, participants may occasionally appear to join (or leave) chat rooms twice with no intervening 'leave' (or 'join') action. This is caused by inconsistent handling of these participant actions in Instant Bloomberg dump files.

Mismatch Between Participants Information

When display names in the 'Join' and 'Says' actions do not match exactly, the name variant that cannot be mapped to a CA DataMinder user is included in the participants list, but no IM event is created for that user.

Quarantine Manager

Quarantined Emails

Note the following issues for encrypted emails that are held in quarantine:

Some encrypted emails cannot be decrypted while in quarantine

If CA DataMinder quarantines an encrypted email, a copy of the email is held in the quarantine queue on the CMS. If possible, the copy is stored in clear text (that is, unencrypted).

- **Emails encrypted using Voltage SecureMail**

  If CA DataMinder quarantines an encrypted email, it can decrypt the email and apply policy based on the text content. A decrypted version of the email is available to reviewers in the iConsole while the email is held in quarantine.

- **Other encrypted emails**

  Encrypted emails detected by the Exchange or Domino server agents are not decrypted by CA DataMinder. Such emails cannot be read by a reviewer if they are quarantined and may not be readable by the recipients if released from quarantine.

Some encrypted emails are decrypted when they are released from quarantine

If possible, encrypted emails are decrypted before they are added to the quarantine queue. But when a reviewer releases an encrypted email from quarantine, it is not always possible to re-encrypt the email.

- **Emails encrypted using Voltage SecureMail**

  When the email is released from quarantine, it is re-encrypted before being forwarded to *external* recipients.

  But the email is not re-encrypted when it is forwarded to *internal* recipients. The email is forwarded unencrypted to internal recipients.

- **Other encrypted emails**

  Applies only to encrypted emails captured by Outlook or Notes endpoint agents or the Exchange 2010 server agent (where the TransportDecryptionSetting is set to Mandatory).

  When the email is released from quarantine, it is forwarded *unencrypted* to the intended recipients.

Appendix A: Support Tools

This section contains the following topics:

Introduction

This section describes the various transforms and utilities in the \Support folder on the CA DataMinder distribution media.

Oracle User Scripts and Associated SPs

When you install your CMS, the CA DataMinder server installation wizard lets you specify the database accounts that CA DataMinder needs to access the CMS database. However, if you intend to use an Oracle as your database engine, you can run the following scripts and SPs (stored procedures) to manually specify your Oracle primary user and schema owner before installing the CMS.

CreateOracleUser.sql

> Run this script to create the primary user and (if required) schema owner. The primary user is the main CA DataMinder database account; the infrastructure uses this account to access the CMS database.

WgnSetupOwner.sql

> This file contains a stored procedure that is installed and invoked by CreateOracleUser.sql.

For further details about manually configuring your CA DataMinder database accounts, see the *Database Guide*.

VBS Scripts for Generating Installation Transforms

For command line, Group Policy or SMS installations, you use Windows\ Installer transform files (.mst) to configure the installation. Visual Basic Script are provided to create these transforms.

Note: For details on using these transform in client machine installations, see the *Platform Deployment Guide's* Technical Information (see page 523) section.

CreateParentNameTransform.vbs

Creates a transform SetParentName.mst that sets the name or IP address of the CMS or gateway that the client must connect to.

ClientLockDown.vbs

Creates the ClientLockDown.mst transform. that disables the Change and Remove buttons when a user selects CA DataMinder in the Add/Remove Programs dialog.

SMSQuietUninstall.vbs

Creates the SMSQuietUninstall.mst transform that enables subsequent silent SMS uninstallations. That is, the transform removes the need for user cooperation when uninstalling.

DisableAutostart.vbs

Creates the DisableAutostart.mst transform that prevents CA DataMinder from starting automatically immediately after installation on servers and client machines.

EnableAppmon.vbs

Creates the EnableAppmon.mst transform that installs application integration. The default installation (that is, a Typical setup in the installation wizard) does not include application integration.

HideConsole.vbs

Creates the HideConsole.mst transform that deselects and removes the Management Console feature from the Custom Setup screen in the installation wizard, and prevents users from installing any console.

EmailClientOptions.vbs

Creates the EmailClientOptions.mst transform that creates registry values to control the behavior of the Outlook endpoint agent.

Deactivation Utility: Wgnbrb.exe

It is possible for a system administrator to deactivate the client component on specified machines, or on all client machines known to the CMS. This is a fail-safe for CA DataMinder to ensure that, after a reboot, it will be disabled on the client machine.

This utility is provided by Wgnbrb.exe. It can be run from any server or client machine with CA DataMinder installed, but it is safest to run it from the CMS itself (the CMS will never become disabled, but see the warning below).

For full usage instructions, run wgnbrb -? from a command prompt. Typical examples are shown below:

- View the status of CA DataMinder on all client machines:

  ```
  wgnbrb -all
  ```

- Deactivate CA DataMinder on all client machines:

  ```
  wgnbrb -all -deactivate
  ```

- Reactivate CA DataMinder on all client machines:

  ```
  wgnbrb -all -reactivate
  ```

- Deactivate CA DataMinder on a defined set of machines:

  ```
  wgnbrb -deactivate MACHINE1 MACHINE2 MACHINE3
  ```

  Important! You will deactivate the CMS if you forget to specify a client machine when you run Wgnbrb.exe on the CMS itself - see below.

- Deactivate CA DataMinder on the CMS by running this command on the CMS:

  ```
  wgnbrb -deactivate
  ```

Note: A user can run Wgnbrb.exe from on any machine providing that their Windows account on the host machine is able to write changes to the registry on the client machine(s) that they want to deactivate. Typically, this means they must be a domain administrator. Alternatively, if you do not want to grant administrator rights to a user, see MS Knowledge Base article Q186433 for an explanation of the WinReg registry key.

Data Management Utility: Wgnmgmt.exe

You use the Data Management utility to back up or restore your CMS and gateways servers.

The CMS and gateway servers use password-protected keys to provide highly secure data management. If you need to restore your CMS or a gateway, you will need to restore its data management key. The Data Management utility is provided to export and re-import such keys.

Exporting Password-protected Key Backups

After installing your CMS and gateways, you must run WgnMgmt.exe on each server to export a password-protected file containing the necessary key details. Note that you need only do this once on each server; you will not need to do it again. Keep this file in a secure location, for example, on a floppy disk in a fireproof safe. The command syntax is:

```
Wgnmgmt e [set the File Name variable] <password>
```

where [set the File Name variable] is the name of the target file (for example, data_management.dat) and <password> is the password you will supply to reimport the key. The password must be at least five characters long.

Re-importing Backed Up Registry Keys

To restore a CMS or gateway, stop the CA DataMinder infrastructure and re-import the data management key. The command syntax is:

```
Wgnmgmt i <my_filename> <password>
```

where <my_filename> and <password> are the backup file and password that you specified when exporting the key (see the previous section).

Note:

- You can find full instructions for backing up and restoring your CMS and gateways in the *Platform Deployment Guide*.
- To restore a client machine, you simply reinstall the client and reconnect to the CMS.

Version Check Utility: Wgncheck.exe

Wgncheck.exe gathers diagnostic data, including CA DataMinder log files, that can be forwarded to CA technical support staff. The output is sent to a zip file on your desktop.

To run the utility, double-click Wgncheck.exe or run this command:

```
wgncheck
```

To see the full usage, run this command:

```
wgncheck -?
```

Client Reparent Utility: wgnReparent.exe

WgnReparent.exe allows administrators to move machines between parents (gateway servers or the CMS, within the same or different hierarchies). If you move a machine to a new CMS infrastructure, you are prompted for confirmation. Typical examples are shown below:

- Display the parents for all machines known to the CMS:

  ```
  WgnReparent -all -show
  ```

- Display the parents for one or more client machines:

  ```
  WgnReparent -show MACHINE1 MACHINE2 MACHINE3
  ```

- Move one or more machines to a new parent:

  ```
  WgnReparent -parent NEW_PARENT MACHINE1 MACHINE2 MACHINE3
  ```

- Move the local machine to a new parent:

  ```
  WgnReparent -parent NEW_PARENT
  ```

- Move all child machines of a specified gateway to a new parent:

  ```
  WgnReparent -parent NEW_PARENT -children GATEWAY
  ```

- For full usage instructions, run:

  ```
  wgnReparent -?
  ```

Note: To use this utility to reparent a child machine to a parent in a different hierarchy, you must be a local administrator on the client machine or CMS.

De-duplication Filter for Symantec Enterprise Vault: wgnsevdd.dll

The \Support\Sevddup folder contains the De-Duplication Filter for Symantec Enterprise Vault (SEV). This filter prevents multiple copies of the same email being archived in SEV. It is implemented as a custom filter for SEV, wgnsevdd.dll.

For full details about the De-Duplication Filter for SEV, see the *De-Duplication Filter for Symantec Enterprise Vault Configuration Guide*.

Appendix B: Accessibility Features

CA Technologies is committed to ensuring that all customers, regardless of ability, can successfully use its products and supporting documentation to accomplish vital business tasks. This section outlines the accessibility features that are supported by CA DataMinder.

Display

To increase visibility on your computer display, you can adjust the following options:

Font style, color, and size of items

Defines font color, size, and other visual combinations.

The CA DataMinder iConsole also supports a High Visibility mode. This increases the size of text and images in the iConsole screens.

Screen resolution

Defines the pixel count to enlarge objects on the screen.

Cursor width and blink rate

Defines the cursor width or blink rate, which makes the cursor easier to find or minimize its blinking.

Icon size

Defines the size of icons. You can make icons larger for visibility or smaller for increased screen space.

High contrast schemes

Defines color combinations. You can select colors that are easier to see.

Sound

Use sound as a visual alternative or to make computer sounds easier to hear or distinguish by adjusting the following options:

Volume

Sets the computer sound up or down.

Text-to-Speech

Sets the computer's hear command options and text read aloud.

Warnings

Defines visual warnings.

Notices

Defines the aural or visual cues when accessibility features are turned on or off.

Schemes

Associates computer sounds with specific system events.

Captions

Displays captions for speech and sounds.

Keyboard

You can make the following keyboard adjustments:

Repeat Rate

Defines how quickly a character repeats when a key is struck.

Tones

Defines tones when pressing certain keys.

Sticky Keys

Defines the modifier key, such as Shift, Ctrl, Alt, or the Windows Logo key, for shortcut key combinations. Sticky keys remain active until another key is pressed.

Mouse

You can use the following options to make your mouse faster and easier to use:

Click Speed

Defines how fast to click the mouse button to make a selection.

Click Lock

Sets the mouse to highlight or drag without holding down the mouse button.

Reverse Action

Sets the reverse function controlled by the left and right mouse keys.

Blink Rate

Defines how fast the cursor blinks or if it blinks at all.

Pointer Options

Let you do the following:

- Hide the pointer while typing
- Show the location of the pointer
- Set the speed that the pointer moves on the screen
- Choose the pointer's size and color for increased visibility
- Move the pointer to a default location in a dialog box

Index