Change and Configuration Management Database

IBM Tivoli Change and Configuration Management Database 1.1.1 Fix Pack 003 Readme (1.1.1.0-TIV-CCMDB-FP0003)

Version 1 1.1

Change and Configuration Management Database

IBM Tivoli Change and Configuration Management Database 1.1.1 Fix Pack 003 Readme (1.1.1.0-TIV-CCMDB-FP0003)

Version 1 1.1

Note
Before using this information and the product it supports, read the information in "Notices," on page 85.

This edition applies to version 1, release 1, modification 1 Fix Pack 003 of IBM Tivoli Change and Configuration Management Database and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication

This publication describes how to install the fix pack for the Configuration Discovery and Tracking and Process Management and Integration Platform features of the Tivoli Change and Configuration Management Database (IBM Tivoli CCMDB).

In addition, there is information about changes to the Configuration Discovery and Tracking and Process Management and Integration Platform features.

Intended audience

This publication is for administrators who want to install and use the latest fix pack.

Publications

Publications for the IBM Tivoli CCMDB library are available at the following Tivoli software library Web site:

http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.ccmdb.doc/ccmdb_welcome.htm

Accessing terminology online

The *Tivoli Software Glossary* includes definitions for many of the technical terms related to Tivoli® software. The *Tivoli Software Glossary* is available at the following Tivoli software library Web site:

http://publib.boulder.ibm.com/tividd/glossary/tivoliglossarymst.htm

The IBM® Terminology Web site consolidates the terminology from IBM product libraries in one convenient location. You can access the Terminology Web site at the following Web address:

http://www.ibm.com/software/globalization/terminology

Tivoli technical training

For Tivoli technical training information, refer to the following IBM Tivoli Education Web site at http://www-306.ibm.com/software/tivoli/education/.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

Online
Go to the IBM Software Support site at http://www.ibm.com/software/support/probsub.html and follow the instructions.

IBM Support Assistant
The IBM Support Assistant (ISA) is a free local software serviceability

workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to http://www.ibm.com/software/support/isa.

Conventions used in this publication

This publication uses several conventions for special terms and actions, operating system-dependent commands and paths, and margin graphics.

Typeface conventions

This publication uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multi-column lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip:**, and **Operating system considerations**:)
- Keywords and parameters in text

Italic

- Citations (examples: titles of books, diskette, and CDs)
- Words defined in text (example: a nonswitched line is called a *point-to-point line*)
- Emphasis of words and letters (examples: "Use the word *that* to introduce a restrictive clause." and "The LUN address must start with the letter *L*.")
- New terms in text (except in a definition list): a *view* is a frame in a workspace that contains data.
- Variables and values you must provide: ... where *myname* represents...

Monospace

- Example and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

Operating system-dependent variables and paths

This publication uses the UNIX® convention for specifying environment variables and for directory notation.

When using the Windows® command line, replace **$***variable* with **%***variable***%** for environment variables and replace each forward slash (/) with a backslash (\) in directory paths. The names of environment variables are not always the same in the Windows and UNIX environments. For example, %TEMP% in Windows environments is equivalent to $TMPDIR in UNIX environments.

Note: If you are using the bash shell on a Windows system, you can use the UNIX conventions.

Chapter 1. Prerequisites for Fix Pack 003

There are two features of IBM Tivoli Change and Configuration Management Database, Version 1.1.1: Configuration Discovery and Tracking and Process Management and Integration Platform. One or both of these features must be installed on your machine.

The Configuration Discovery and Tracking feature is sometimes referred to as IBM Tivoli Application Dependency Discovery Manager, Version 5.1.

It is not required that you install the Configuration Discovery and Tracking server, version 1.1.1.1 (Fix Pack 001) or Configuration Discovery and Tracking server, version 1.1.1.2 (Fix Pack 002).

If, for the Configuration Discovery and Tracking feature, you are using a supported Windows or Linux® for System z™ operating system, you must install 5.1.1.2-TIV-ITADDM-LA0001 before installing Fix Pack 003.

For information on upgrading to version 1.1.1, visit the IBM Tivoli CCMDB information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/index.jsp.

For the latest information regarding Fix Pack 003, go to the IBM Tivoli CCMDB Product Support Web site at http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliChangeandConfigurationManagementDatabase.html. To download Fix Pack 003, click **Fixes by version**. This link is under the Download heading.

Chapter 2. APARs and defects in Fix Pack 003

The following APARs and defects have been addressed in Fix Pack 003.

Configuration Discovery and Tracking feature APARs for Fix Pack 003

The following APARs are addressed by Fix Pack 003 for the Configuration Discovery and Tracking feature:

APAR number	Abstract
IY92027	Unable to launch in context from the Process Management and Integration Platform ComputerSystem to a Configuration Discovery and Tracking Java client.
IY93898	The storage sensor running on the anchor is timing out. This timeout causes the threads on the primary server to go into a state where they cannot be stopped or garbage collected.
IY94312	During a discovery of a HP-UX 11.23 system, a message is written to `stderr - WARN collation.stderr - adb: info`. Option *-k* is deprecated. Running the command natively on HP 11.23 also yields same result.
IY95270	Unable to differentiate between objects of the same name.
IY95474	Configuration items with many children, in this case, the BridgeSNMPSensor, can take a very long time to store their results if it is not implemented using the appropriate approach.
IY95675	Scopes loaded with loadscope cannot be discovered with api.sh.
IY95769	When the synchronizer tool loses the connection to the database, you have to restart the tool to regain the database connection.
IY95842	When you delete a configuration item and then refresh the topology, the topology does not complete the rebuild. When you restart the Configuration Discovery and Tracking server, the configuration items still exist.
IY95961	If you try to discover a system that hosts JBoss 4.1, the JBoss3xSensor is deployed and completes with success.
IY96098	STATEMANAGER and TOPOLOGYBUILDER consume excessive memory inside of the topology service.
IY96109	Discovery waits on sensors that do not finish or are complete.
IY96388	For devices categorized as routers, Show L2 topology is added to the pop-up menu.
IY96401	When comparing WebSphere nodes using a component comparison report, the system hangs. No results are reported.
IY96547	An API import followed by a sync does not refresh Details and JDO cache.
IY96549	There are two problems fixed by this APAR: PrimaryMacAddress set by discovery should be the lowest lexical one. SNMP discovery does not set the PrimaryMacAddress.
IY96560	If extended attributes are gathered inside transactions, the action fails and a memory leak occurs.

APAR number	Abstract
IY96565	A prerequisite for the UNIX credentials is that the user must be able to run lsof and must be in sys group, or lsof must be setgid sys or ″sudo″ access for lsof. However, Configuration Discovery and Tracking does not report the inability to run or find lsof.. The discovery is reported as successful.
IY96727	Multiple user interface performance issues are addressed: 1) When running a discovery, there is a delay of 50 seconds before seeing the scope dialog. 2) Clicking on the Topology tab, there is a delay of 5 minutes before the graph showing the business applications is displayed. 3) Showing physical topology of a business application takes 15 to 20 minutes and sometimes times out completely that causes the user interface to display a ″Client Session has timed out″ message.
IY96766	The WebLogic sensor does not detect JDBC dependencies because of issues with name resolution.
IY96877	The new or deleted configuration items in access collections are not copied to enterprise after the initial sync of the collection.
IY96948	Performance of moving custom server templates in the list is not optimal.
IY96972	The Configuration Discovery and Tracking J2EE domain naming rule requires the name to be unique.
IY97119	The Fix Pack 002 installation program fails on the prerequisite checks.
IY97163	Unable to obtain access with the sapccms sensor.
IY97224	For Fix Pack 002, disabling a sensor using the discover-sensor/ SensorName.xml file causing the sensor not to load and no further sensors to load.
IY97296	WebLogic discovery fails if the WebLogic is setup using IP address rather then name for the server and node.
IY97322	Discovery of Windows services should be done by the Windows computer system sensor, not the generic server sensor.
IY97459	When you run a discovery involving many sensors, including the WebLogic sensor, the storage time exceeds 30 minutes.
IY97511	The z/OS DLA is missing the Dependencies tab for IMS and CICS. The z/OS objects are not displayed in the Business Services wizard. The subsystem name is blank for IMS. The z/OS object are displayed in the Business Applications when using application descriptors that do not contain z/OS objects.
IY97512	An error occurs when running the WebSphere MQ sensor on a Windows machine, having the product installed in the C:/Program Files/MQSeries directory. It seems that the space in the directory name causes problems when parsing the queue name.
IY97516	Fix Pack 002 database template upgrade script reports an error with the ″Failed to upgrade the database templates″ message if bash is not in the /bin/bash directory.
IY97552	For Fix Pack 002, data is corrupted after session reuse.
IY97566	Application descriptors do not work on Oracle SIDs with uppercase characters.
IY97588	If *–DEPTH=-1* is used, XML generation by api.sh consumes excessive memory.
IY97598	According to the Web site, an L2Interface has an ″accessedVia″ relationship with a ComputerSystem on the parent attribute. According to the source code, an L2Interface has a ″contains″ relationship with a ComputerSystem on the parent attribute.

APAR number	Abstract
IY97635	Properties that are manually added to the collation.properties file are not migrated when you upgrade to Fix Pack 002.
IY97636	When a full synchronization occurs, there is an option to delete everything. The default is 'do not delete implicit relationships.' Prior to this APAR, all objects were always deleted.
IY97747	A scheduled synchronization of the enterprise configuration management database to the domain configuration management databases can be deleted and is removed from the user interface. However, the synchronization is run again at the next scheduled time and is displayed on the user interface.
IY97753	An error occurs when an SQL query is constrained to only return objects in the access collections a user is allowed to access.
IY97815	When you load the z/OS DLA books with the loadidml command, the Database tab of the component DB2 SubSystem contains the same row, repeated several times. In addition, an IMS subsystem has many transactions in a z/OS DLA book, only the last transaction is displayed in the user interface.
IY97828	The Configuration Discovery and Tracking software does not allow the WeblogicJDBCTxDataSource and WeblogicJDBCConnectionPool objects to have the same name yet Weblogic allows this configuration. Fix by moving the naming rules from SoftwareResource/J2EEResources to the subclasses.
IY97886	An exported PDF file about Physical Infrastructure is 0 byte.
IY97899	For Configuration Discovery and Tracking on the supported Windows operating systems, discovery with the supported SAP sensors does not work.
IY97900	When running a synchronization with many business applications, an out of memory condition can occur.
IY98062	The Configuration Discovery and Tracking discoveries are bring back the same IP addresses multiple times.
IY98100	When creating a component comparison report, save the report with CSV file type and the size is 0 byte.
IY98121	After installing Fix Pack 002, MQL reports an error.
IY98194	The IIS discovery is incomplete.
IY98200	Only Microsoft application on Windows are discovered. In the software components of a discovered Windows server, not all software is listed.
IY98302	If there are multiple instances of the DB2 software running on a server, only one instance is discovered.
IY98317	In a z/OS component, the DB2 software version is blank for the DB2 subsystem.
IY98563	When creating a new sensor, for example, the StackScan sensor, the ports in the GUI appear to be truncated in the 700 range.
IY98732	Configuration items in each of the access collections on the domain Configuration Discovery and Tracking servers are not merging up to the enterprise Configuration Discovery and Tracking server.
IY98793	The Configuration Discovery and Tracking installation process does not grant the appropriate access to database objects for versions to be created.

APAR number	Abstract
IY99037	When large quantities of data are gathered for domain reports, performance is slow. Pagination is added. In addition, new function is added to improve the usability of the domain reports in an enterprise environment: 1) The domain name is added to the output of the query results. 2) Export functionality has been added to the system and software inventory reports. To view this functionality the following collation.properties setting must be set to *true*: `com.collation.comain.pdfreport.enabled=true` When set to *true*, you can save the reports in the domain manager in PDF format.
IY99081	There are problems when discovering VMware ESX 2.5 using Fix Pack 002. Changes were made to support the ESX 3.0 version that seemed to break the discovery in the environment when discovering version ESX 2.5.
IY99151	It takes a long time to create and remove extended attributes.
IY99288	The naming rule for the CitrixServer does not allow for more than one CitrixServer per CitrixZone.
IY99290	The component server is incorrectly displayed in as a part of the business application.
IY99336	Configuration Discovery and Tracking cannot store results of WebLogic discovery because of a domain conflict.
IY99379	If you have two active directory users in the credentials list with the same scope, Windows discovery does not know which discovery to use.
IY99381	If you load the XML files generated with the z/OS DLA and the loadidml process completes successfully, the MQ results file can contain errors. The attributes required for the naming rules of this object were not set.
IY99385	It takes a long time open an application summary report (more than 30 minutes).
IY99401	In a custom server template, using *$COLL_PROGPATH* does not work. The log indicates that the environment variable is truncated.
IY99420	Using the bulk load program, you cannot create FunctionalGroup object and relationships.
IY99427	The change history is not created or viewable in the user interface for extended attributes.
IY99428	When running the OpenVMS sensor, the sensor reports the following error: ″A storage error occurred.″
IY99597	The Change History report does not catch the new table in Oracle databases.
IY99648	If you unzip the sdk/doc/model/CDMWebsite.zip file and open WebSitefiles/files/index.html, there are broken links to portions of the documentation.
IY99680	The Configuration Discovery and Tracking feature does not receive application descriptors for the SQL server on Windows machines. The following error message is displayed in the log file: ″Error: CreateFile: The network path was not found.″
IY99764	Problems occur when trying to create a federates relationship between a Business System and a Computer System.
IY99803	The TaddmTool.exe fails when `GetSystemInfo` WMI call returns null for product information.
IY99806	The sensor for the SQL sever is not started although the server processes match the template.

APAR number	Abstract
IY99808	If a PortScan sensor or StackScan sensor returns an empty or incomplete PortListResult, that sensor is rescheduled to an anchor and returns a complete list. The complete list is ignored and no SessionSensor is seeded.
IZ00015	Extraneous connections between Solaris, and switches and VLANs

Configuration Discovery and Tracking feature defects for Fix Pack 003

The following defects are addressed by Fix Pack 003 for the Configuration Discovery and Tracking feature:

Defect number	Abstract
8424	Relns missing in CDM for TPC.
9997	Change History in Domain Manager should have an option to view Relative and Absolute changes for a component.
10278	There is no page layout size setting in GUI when printing or exporting details.
11105	While Configuration Discovery and Tracking server is starting, clicking refresh button causes many Java errors in logs.
11398	When discovering large scopes, the UI becomes unusable if the Overview is sorted by description or status.
11545	DB2/Windows sensor fails with uncaught exception.
11582	/bin/csh is hard coded in unix detect locale function
11829	From the Create Query dialog, if the attribute field for a specific datasource component is removed from the datasource list, it no longer makes any sense to keep it in the first combo box of the attribute section (as it cannot be used).
12150	After creating a new Application template, there is no way to modify or view the contact-specific details.
12643	API unable to delete within transaction, hangs topomgr.
12666	Performance tuning of IDD sensor.
12688	Citrix users and groups are not accurate enough.
12719	WMI Provider deployment fails if ADMIN$ share is removed by MS hotfix
12720	JBoss dependency tab is empty though JCA tab has DB2 details
12759	IP address rather than hostname shown in Discovery Overview screen
12784	New WAS sensor fails to collect some data on standalone WAS servers
12802	This error, IndexOutOfBoundsException in api.getDomain(), is logged but is benign, as it is caught and processed.
12815	The contains operator used to be able to operate like java.String.indexOf() -- returning true if the presence of a substring was found. Now, substring detection is not working.
12865	Scoped properties in collation.properties do not work
12870	API: findChanges(.....,DELETED) not working.
12906	Race condition during Change Manager initialization causes discovery to hang
12922	StackScan sensor duplicates open port list on every discovery run

Defect number	Abstract
12934	Performance tuning of Configuration Discovery and Tracking IDD Sensor.
12952	Apache server missing from the seedconfiguration map.
13060	Sensors starting full anchor instead of local anchor.
13061	If a reference to a host is not fully qualified and the Configuration Discovery and Tracking server does not have DNS configured properly, any calls made to the method OsFactory.newOS() on that short name fail.
13083	Log in to Domain Manager, pick any discovery run, and click **Scope Details**. No output.
13095	Bulkload performance.
13184	When migrating from Fix Pack 001 to Fix Pack 002, collation.properties is not being updated correctly.
13360	Null pointer exception in presentation api getDetailsPanel.
13438	DB2 sensor requires bash to work.
13547	DiscoveryProfiles - Prerequisites for ApacheSensor are not correct, only AnchorSensor shows up as the required prerequisite.
13561	Level1 discovery after a Level3 discovery fails.
13615	Configuration Discovery and Tracking IDD Sensor not reporting UDP ports without NMAP.
13675	Deletion of a scope throws ObjectNotFoundException.
13806	Argus policy file writes take too long.
13838	DB2 do not discover multi-instance on multi-installation.
13887	System p discovery problem.
13936	Cannot delete a scope from scope set.
13942	Corrected excessive use of remote reverse DNS lookups.
13966	New database in DB2 does not show up in change history.
13977	getChangeHistory leaks 3000 threads/minute, gets OutOfMemory.
13986	Configuration Discovery and Tracking does not start unless COLLATION_HOME is set.
14012	Details panel information truncated when exported to PDF file.
14135	Change history shows the create and delete events for objects.
14168	AIX sensor does not set primaryMACAddress.
14185	MQ Sensor - detect application descriptors.
14237	Interfaces panel for switches is not sorted.
14243	In Product Console (UI) Websphere Component Comparision Run Repo.
14275	Change history does not work for Extended Attributes.
14284	Comparing Business applications across versions does not work.
14339	control stop command executes with any password in third attempt.
14366	DP - Newly created profile does not remain selected.
14508	WAS sensor - failure discovering Web services.
14527	L2 discovery of switches with duplicateMAC ports is wrong.

Process Management and Integration Platform feature APARs for Fix Pack 003

The following APARs are addressed by Fix Pack 003 for the Process Management and Integration Platform feature:

APAR number	Abstract
IY99646	Launch in context does not work with new data APIs
IY99026	Standard text value with an attribute name
IY98295	Data standardization function does not work if there is old incorrect data
IY98004	Recoverability if person role table gets corrupt
IY97947	CREATE CI, needs to use an array for extended attributes
IY97844	Exception when creating the same account twice
IY95658	Reporting application does not work if it starts before Request Manager
IY95707	Slow performance in ″REMOVEPERSONFROMACCOUNT″
IY96099	CONFIGITEMOWNER can not launch config item properties for owned CIs
IY99880	Performance: Long run slow down when using Work with CIs

Process Management and Integration Platform feature defects for Fix Pack 003

No defects were included Fix Pack 003 for the Process Management and Integration Platform feature.

Chapter 3. APARs and defects in Fix Pack 002

The following APARs and defects have been addressed in Fix Pack 002.

Configuration Discovery and Tracking feature APARs for Fix Pack 002

The following APARs were addressed in IBM Tivoli CCMDB Fix Pack 002 and are also included in this fix pack for the Configuration Discovery and Tracking feature:

APAR number	Abstract
IY86993	Topology build for small scope discoveries takes an excessively long time in large environments.
IY90384	The contactinfo, siteinfo, and admininfo classes are missing relationships in the ITSM CI portlet view.
IY90625	Unable to sort columns by name.
IY90759	DB2 sensor is consuming the DB2 client processes and not returning them to be further processed when it ultimately fails.
IY90985	Validation.jar does not work.
IY91604	Access collections do not have an icon.
IY92149	SAP sensor does not work after upgrading from 5.1 to 5.1.1
IY92250	Copy of custom server template is not working.
IY92296	No possibility to create custom application template based on oracleappserver.
IY92297	Oracle Application Servers with common domain are getting merged.
IY92418	WebSphere Application Server-> DB2 dependencies are not discovered if the WebSphere Application Server resource is created to a virtual interface on the DB2 server.
IY92728	Single IP scopes create and update CIs even if there is no device on the IP
IY92800	Windows systems are stored twice.
IY93028	Authorization file becomes corrupted if the access list is updated from two different users at the same time. This change addresses the auth file corruption, but multiple administrators **should not** update access lists at the same time. If multiple administrators update access lists at the same time, results are unpredictable. In the database, the entire access list that the administrator sees at the moment of clicking **Yes** is saved. If there are two administrators, they should refresh after each administrator makes a change. If the refresh does not occur, an administrator can undo the other administrator's changes. Customers should back up auth and userdata XML files on a regular basis to ensure recovery from such scenarios.
IY93164	Pop message typo: This operation may take a long time. Do you 'with' to continue? change with to wish.
IY93239	There is a problem with the business application finding an unexpected switch in the destination directory which can cause a blank line in the log file.
IY93365	Performance to retrieve CI details can take 45+ seconds.
IY93383	Full synchronization takes up to 30+ minutes to synchronize the eCMDB and the domain CMDB.

APAR number	Abstract
IY93505	Installing 5.1.1 fix pack 1 fails when running with TADDM userid, FP1 install requires ROOT.
IY93527	Authorization process: Sequential processing of authorization requests by eCMDB Authorization Manager with a significantly large number of users and permissions, access collections, roles associated with users results in a performance hit.
IY93581	After restarting the Configuration Tracking and Discover (CDT) server, you cannot log in to either CDT or the eCMDB using USERS created in eCMDB.
IY93656	GUI installer does not work on AIX.
IY93717	DB2 discovery on Windows.session layer not allowing DB2 authorizations to be used to establish WMI sessions.
IY93839	Cannot create relationships on ITSYSTEM classes using the bulk load program.
IY93877	The eCMDB CI retrieval performance is slow due to going to domain for detail. Before implementing this change, back up your database. In order to turn on full sync edit the *%ECDT_HOME%*/etc/domainquery on the eCMDB and add this to the first list `SYNC_ALL_ATTRS` Complete a 'full synchronization' so all attributes can be pulled in. A full synchronization operation can take a long time with this change. To backout this change, you need to restore the database backup taken above.
IY93911	Chained anchor server does not deploy.
IY93968	Need the index on cause_id for incremental synchronizations.
IY93969	Improve Storage sensor performance.
IY94097	Cannot delete unwanted custom server templates.
IY94118	Not able to discover Weblogic if the host has multiple IP's.
IY94282	No configuration files are not collected if a single one of the configuration files fails due to permissions.
IY94603	The items listed in the packages field need to display in the software installed report. This APAR also added some new database views. **Note:** See Technote#: 1249657 on the support Web site for the list of added database views. You can also go to this link: http://www-1.ibm.com/support/docview.wss?rs=3036&context=SSPLFC&q1=1249657&uid=swg21249657&loc=en_US&cs=utf-8&lang=en
IY94908	DB2FIND.SH script fails on the grep command when there are two or more users with a common leading subset of characters.
IY94932	Configuration Tracking and Discovery is unable to discover WAS using the WAS 6.0 client.
IY94996	WMI discovery fails across Domain environment.
IY95227	During SQL server discovery, the SQL sensor stops if it finds a null in SID or VERSION field.
IY95509	RMID stops processes that take too long to respond.
IY95510	z/OS DLA inserts the wrong tables.
IY95579	The loadscopy.jy script does not check duplicated entries. The loadscope.jy script is not parsing IP address correctly.
IY95770	Deploying Windows gateway install fails.

APAR number	Abstract
IY95947	When a new gateway with a scope limitation is added, if Configuration Discovery and Tracking is not restarted, the first discovery using that gateway fails, as the scope is not recognized.
IY95954	It take a long time (several minutes) to open the Product Console when there are many collections.
IY96078	The discovery does not complete. Change propagation in ChangeHistory causes excessive memory consumption in ChangeManager (inside of EventsCore).

Configuration Discovery and Tracking feature defects for Fix Pack 002

The following defects were addressed in IBM Tivoli CCMDB Fix Pack 002 and are also included in this fix pack for the Configuration Discovery and Tracking feature:

Defect number	Abstract
5382	When deleting a business application containing many components, getting the "changes detected, please reload view" message at the bottom of the UI more than once. sometimes twice, sometimes three times.
6692.2	Go to the application topology graph, go to Topology -> Options and click on the "Hide components with no dependencies" checkbox, then click on OK. The resulting graph still has the components with no dependencies and, in fact, many of the components that had dependencies are hidden.
8116	Change history appears to not be working because of default date choice.
9111	Error msg requiring scope despite assigned scope to schedule.
9266	When you try to create a schedule, and you don't choose a scope to run against, it brings up an error message which says, "no scopes are defined. define a scope".
9267	For a scheduled discovery run, multi-selection of scope sets is not allowed. If you try to select multiple scope sets, only the first scope set is displayed as part of the schedule. You can select one scope set for a scheduled discovery run.
9268	Re-created schedules with same name have unexpected scopes.
9374	The wrong error coming from OracleAppSensor.
9763	Inventory report for databases does not identify type of database.
9798	StorageErrors on AppConfigJdo.
9950	Can not set extended attributes from CustomServer extensions.
9951	GenericComputerSystemSensor should launch a CustomComputerSystemSensor if it is unable to determine OS type.
10143	Scope restrictions on gateways do not work reliably.
10807	Problem with PATH due to globalization changes.
11083	Delete component failing from dormant component report.
11084	Discover runs are never deleted from the database.
11092	Suppress "ignore" custom server templates in app templates.
11098	Custom Server Template for OracleInstance does not work.
11222	WAS sensor fails if WAS 6 client jars are used.

Defect number	Abstract
11396	System Inventory report needs sort capabilities.

Process Management and Integration Platform feature APARs for Fix Pack 002

The following APARs were addressed in IBM Tivoli CCMDB Fix Pack 002 and are also included in this fix pack for the Process Management and Integration Platform feature:

APAR number	Abstract
IY92712	Deleting of CI in protected lifecycle state is possible
IY93905	Missing computer system name rule in PMIP
IY94029	Performance: >15 minutes to view a CI
IY94703	NAME attribute not set
IY94713	FINDRELATIONSHIP API is used and for each relationship getting
IY94714	ASSIGNCITOACCOUNT performance issue
IY94717	CCMDB performance issue
IY94846	Auditing doesn't pull all data causing unknown to be displayed
IY95267	Remove CI from account failed
IY95294	Report ″INSTALLED SOFTWARE″ returns no data
IY95334	Get organizations that are used by any account
IY95188	Assign person to organization performance issue
IY95433	Audit report returns negative number in results report

Process Management and Integration Platform feature defects for Fix Pack 002

The following defects were addressed in IBM Tivoli CCMDB Fix Pack 002 and are also included in this fix pack for the Process Management and Integration Platform feature:

Defect number	Abstract
201642	Change getPermissions to use find instead of findCollections
201779	RAT: Fixes for Report Definition Editor
201817	Enforce aggregation operations run in query->sort->group order
201952	CI: Create Computer System CI, erroneous char in IP interface

Chapter 4. APARs and defects for Fix Pack 001

The following APARs and defects were addressed in Fix Pack 001:

Configuration Discovery and Tracking feature APARs for Fix Pack 001

The following APARs were addressed in IBM Tivoli CCMDB Fix Pack 001 and are also included in this fix pack for the Configuration Discovery and Tracking feature:

APAR number	Abstract
IY85344	Access errors in DB2 sensor due to not using CSH.
IY89029	An issue with CLOB tables and the 9i Oracle driver. This is corrected by an updated Oracle driver, which has been included in this Fix Pack.
IY89316	Discovery results are corrupted because of lack of server filesystem space.
IY89859	Extended attributes not showing up in eCMDB.
IY90792	The change history does not list current version change.
IY90926	If a software server is listening only on the loopback interface (127.0.0.1), its name does not match the computer system it was actually running on.
IY91098	A WebSphere custom server template does not identify the application running on a WebSphere server.
IY91226	Cannot access the User tab in the Domain Manager.
IY91509	Dependencies disappear after running a discovery.
IY91570	ComputerSystem objects stored excessively.
IY91596	Generic server sensor takes too long to store if there are NFS dependencies.
IY91688	When running a discovery, the system hangs after the TopologyBuilder progress shows 100% complete.
IY91791	Discovery of network devices yields excessive amount of storage errors.
IY92072	SSH fails on OpenVMS due to use of ″echo″ command.
IY92164	WebSphere discovery does not work through an anchor server.
IY92499	JAVA UI unknown server report fails if any attributes are missing.
IY92555	Application templates are not stored after creation.
IY85393	The change report is not available until 6 hours after changes are made.

Configuration Discovery and Tracking feature defects for Fix Pack 001

The following defects were addressed in IBM Tivoli CCMDB Fix Pack 001 and are also included in this fix pack for the Configuration Discovery and Tracking feature:

Defect number	Abstract
5399	Delete any component in the topology or tree, and it gives you an error that says, ″Error: Component is in use by other components″.
5451	When you are discovering a Windows system, without a gateway, extra files are left on the Windows system.

Defect number	Abstract
7618	Reading the memory size from the /dev/mem directory on HPUX systems with more than 64 GB of memory causes CPU spin on the HPUX system.
8041	The confirmation window for the Dormant Component report needs a scrollbar.
8787	The Topobuilder intermittently hangs in cache clearing.
9110	Externalized the adb command so it can be invoked with sudo access by modifying the collation.properties file.
9114	A WebLogic discovery shows following error on the console: sensor failed in remote server:null.
9117	Multiple JMS servers running on a machine, collecting a configuration file from each server, all JMS servers getting the same configuration file back.
9119	Compare two HP computer systems and click on the link to show the difference between installed patches. The link seems to be ineffective and the Product Console shows an NPE error.
9121	Create a version, then try to compare against it, and it does not show up in the list. Only the current version is listed. If you restart the GUI, then the new version appears.
9122	Run a comparison between the configuration files of threeJMS servers, click on the link that should show the difference between the JSM/main_config file. You get an NPE error.
9123	When you compare the configuration files from two custom servers, the two columns at the top of the report are both named JMS/main_config.
9130	When doing a specific query, get NPE error.
9251	The custom server template for WebSphere will not collect configuration files.
9260	Launch in context URLs are not working.
9358.1	The configuration file changes are not detected correctly.
9370	Trying to edit a configuration file that has been captured causes a NPE error.
9380	Running bin/dbupdate.sh drops all the CCMDB database views if CDT is using a DB2 CMDB database.
9383	The Runtime tab for Custom Servers lists the process ID and program name in the "Ports" section. The actual ports are not visible.
9401	Application descriptors do not work on a Windows system.
9407	CDT/TADDM do not discover an application on Windows if the application's listen socket is the first one on the process list.
9614	The Windows WMI provider was using aports.dll to get network socket information on Windows 2003 when it should be using the built-in netstat command.

Process Management and Integration Platform feature APARs for Fix Pack 001

The following APARs were addressed in IBM Tivoli CCMDB Fix Pack 001 and are also included in this fix pack for the Process Management and Integration Platform feature:

APAR number	Abstract
IY88838	Cancel from a CM navbar action returns to a blank window
IY91457	CI sort does not sort drop-down data
IY92024	Unable to Launch in Context from PMIP computer system to CDT Java client

Process Management and Integration Platform feature defects for Fix Pack 001

The following defects were addressed in IBM Tivoli CCMDB Fix Pack 001 and are also included in this fix pack for the Process Management and Integration Platform feature:

Defect number	Abstract
193987	Coordinate two-table udpates, if one fails both must roll back
196593	Beginning spaces as part of the MSS name
197570	Disable Save as Draft feature in the Define Targets and Refine SPB task portlets
198501	When selecting a RFC to transfer from the Work with RFC table, get the RFC from the back end database instead of the cached RFC from the Work with RFC table
198511	Add ChangeAdmin, ChangeManager, ReleaseManager, and ReleaseAdmin groups to the DynamicFlowAdmin role of the HWEProcessController EJB application
199166	Add help descriptions to the “Manage Areas” extended attribute
199288	Update MSS does not check dup to the existing MSS
199296	Prevent wrapping of Filter by label in the Forward Schedule of Change filter
199312	Increase window size of Close Window pop-up box to prevent wrapping of text descriptions in different languages
199328	Use resource bundle for TCMIM labels that missed translation
199467	Need to add Read permission to ITAPM roles
199868	Make sure that the correct table row is selected after a sort in the TCM software package selection task portlet
199977	Get Exception After Trying to Update DSC Relationship Panel
200495	Create sql scripts to query and update the version for component
200553	Application always exported, even when path isn’t supplied
200691	Uninstall is broken - subsequent re-installs will fail
200706	Read permission is not created for ConfigAuditor role
200737	Incorrect directory used for Config Management metadata
200754	Re-associate BPEL applications with libraries after upgrade
200758	Uninstall: Exception in CCMDB_uninstall.log
200924	Uninstall after fixpack - needs to support changed IDs
200934	ECMDB domain port needs to be set to correct port
201101	Upgrade of IMs failed ExternalCommandException

Defect number	Abstract
201113	Uninstall is failing to uninstall change/config and dynaflows
201214	Assign People to Organization selects first User in list
201215	ExtendAttribute: Attribute not defined with 2 accounts assigned
201216	ExtendAttribute: Remove from Single Acct leave Acct Label in CI
201232	Exception if cancel button is used before running install
201271	Search: with Match Case not selected no data returns
201296	Audit: When multiple accounts is not enabled Acct Column display

Chapter 5. Installation overview for Fix Pack 003

If you have installed only the Configuration Discovery and Tracking feature, follow the steps in Section A and skip Section B.

If you have installed both features, Configuration Discovery and Tracking and Process Management and Integration Platform, read both Sections A and B.

Important: If you have installed both features on different machines, you must first apply the fix pack to the Configuration Discovery and Tracking server.

Determining which installation wizard was used to install the base Configuration Discovery and Tracking feature

Note: If you attempt to install this fix pack using the IBM Tivoli CCMDB installation wizard, and you had previously used the Configuration Discovery and Tracking installation wizard, you get the following error message:

```
This fix pack cannot be applied to the current installation of
Configuration Discovery and Tracking using this installation
program. You must use the installation program that is
supplied with the Configuration Discovery and Tracking
product.
```

It makes a difference whether you used the Configuration Discovery and Tracking or the IBM Tivoli CCMDB installation wizard to initially install the baseConfiguration Discovery and Tracking feature.

To determine which installation wizard was used, complete the following steps:

1. Go to the base Configuration Discovery and Tracking installation directory. The default directory is located at /opt/IBM/CCMDB/.
2. If the Configuration Discovery and Tracking feature was installed with the IBM Tivoli CCMDB installation wizard, you see a **vpdexport** directory. If you see this directory, you used the IBM Tivoli CCMDB installation wizard.

Section A: Installing Fix Pack 003 for the Configuration Discovery and Tracking feature

You must install the Configuration Discovery and Tracking portion of this fix pack using the same method that you used when you installed the base Configuration Discovery and Tracking feature.

That is, if you installed the Configuration Discovery and Tracking feature for the first time using its own installer, you must use that same method to install this fix pack, as described in Section AA.

If you installed the Configuration Discovery and Tracking feature for the first time using the IBM Tivoli CCMDB installation wizard, you must use that method to install this fix pack, as described in Section AB. This requirement holds whether you initially installed Version 1.1 and then upgraded to Version 1.1.1, or whether you initially installed Version 1.1.1.

If you have installed the Enterprise Configuration Discovery and Tracking server and the Configuration Discovery and Tracking server, you must upgrade the Enterprise Configuration Discovery and Tracking server and the Configuration Discovery and Tracking server at the same time. If the Enterprise Configuration Discovery and Tracking server and the Configuration Discovery and Tracking server are not at the same code level, they do not work properly.

This fix pack is supported on a limited number of operating systems. Before beginning the installation, verify that your system is running a supported operating system and that you meet all hardware and software requirements.

Supported operating systems

The following operating systems are supported for the Configuration Discovery and Tracking portion of this fix pack:

Table 1. Support operating systems for the Configuration Discovery and Tracking portion of Fix Pack 003

Operating system and supported release	**Support details**
AIX® 5.2 (Release previous to current platform release)	• Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
AIX 5.3 (Current platform release)	• Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
RedHat Enterprise Linux 3.0 x86 32 bit (Two releases previous to current platform release)	• Client or endpoint supported • Server or anchor supported
RedHat Enterprise Linux 4.0 x86 32 bit (Release previous to current platform release)	• Client or endpoint supported • Server or anchor supported
RedHat Enterprise Linux 4.0 for System z (Release previous to current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
Solaris 9 SPARC (Release previous to current platform release)	• Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
Solaris 10 SPARC (Current platform release)	• Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
SuSE Linux Enterprise Server 9.0 x86 32 bit (Release previous to current platform release)	• Client or endpoint supported • Server or anchor supported

Table 1. Support operating systems for the Configuration Discovery and Tracking portion of Fix Pack 003 (continued)

Operating system and supported release	Support details
SuSE Linux Enterprise Server 9.0 for System z (Release previous to current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
SuSE Linux Enterprise Server 10.0 x86 32 bit (Current platform release)	• Client or endpoint supported • Server or anchor supported
SuSE Linux Enterprise Server 10.0 for System z (Current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode • Server or anchor supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode
Windows 2000 Professional (Two releases previous to current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported
Windows 2000 Advanced Server (Two releases previous to current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported
Windows 2000 DataCenter Server (Two releases previous to current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported
Windows Server 2003 DataCenter (Current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported
Windows Server 2003 Standard Edition (Current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported • Server or gateway supported
Windows Server 2003 Enterprise Edition (Current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported • Server or gateway supported
Windows Server 2003 Enterprise x64 Edition AMD64 and EM64T (Current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode, limited manual installation is required
Windows Server 2003 Standard x64 Edition AMD64 and EM64T (Current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported: 32b app tested on 32b kernel, 32b app tested on 64b kernel under 'toleration' mode, limited manual installation is required
Windows XP Professional (Release previous to current platform release)	• Requires 5.1.1.2-TIV-ITADDM-LA0001 • Client or endpoint supported

IBM Tivoli CCMDB does not support Configuration Discovery and Tracking's Windows operating systems.

If you have an environment with the Process Management and Integration Platform feature and an Enterprise Configuration Discovery and Tracking server with the previous supported platforms (the Enterprise Configuration Discovery and Tracking server cannot be running on a Windows or Linux on System z operating system), you can add a Domain Configuration Discovery and Tracking server running on a Linux on System z operating system to this environment setup.

For more information about operating systems supported by IBM Tivoli CCMDB, visit the IBM Tivoli CCMDB information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/index.jsp.

Refer to the online documentation for the Configuration Discovery and Tracking feature for more information about anchors and gateways. Publications for the Configuration Discovery and Tracking library are available at the following Tivoli software library Web site: http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.ccmdb.doc/ccmdb_welcome.htm

All Windows gateways must be running the Bitvise WinSSHD 4.x version of SSH. The Configuration Discovery and Tracking server communicates with the gateway using SSH, regardless of the platform that the server is using.

If the Configuration Discovery and Tracking server is running on a Windows or Linux for System z system operating system, anchors on Windows systems are supported.

If the Configuration Discovery and Tracking server is running on a Windows or Linux for System z system operating system, an anchor and a gateway running on the same Windows system is supported.

Hardware requirements

The following list describes the processor, memory, and disk space requirements for a Configuration Discovery and Tracking server. The requirements are the same whether the machine is a single Configuration Discovery and Tracking server, an enterprise Configuration Discovery and Tracking server, or a domain Configuration Discovery and Tracking server.

Each Configuration Discovery and Tracking server requires a machine with:

- 100 GB of available disk space
- 2-4 CPUs with a minimum process speed of 2 GHz
- 2-8 GB of memory

  For a Linux for System z operating system, you should have at least 4 GB of memory to run all of the Java™ services.

You must install the database on a different machine. That machine and the Configuration Discovery and Tracking server must each have 2 GB of memory. For medium to large environments, use more memory.

Software requirements

Before you install this fix pack, you must have one of the following software components installed:

- The Configuration Discovery and Tracking feature for Tivoli Change and Configuration Management Database Version 1.1.1

- The Configuration Discovery and Tracking feature for Tivoli Change and Configuration Management Database Version 1.1.1 Fix Pack 001
- The Configuration Discovery and Tracking feature for Tivoli Change and Configuration Management Database Version 1.1.1 Fix Pack 002
- The Configuration Discovery and Tracking feature for Tivoli Change and Configuration Management Database Version 1.1.1 Limited Availability Fix 001

Firewall considerations

If a firewall exists between IBM Tivoli CCMDB and Configuration Discovery and Tracking resources, you cannot use the silent install and silent uninstall options.

Section AA. Using the Configuration Discovery and Tracking installation wizard

If you originally installed the Configuration Discovery and Tracking feature using the Configuration Discovery and Tracking installation wizard, complete the following tasks to install Fix Pack 003 using the Configuration Discovery and Tracking installation wizard.

Before you install Fix Pack 003, complete the following tasks:

1. If running in an enterprise environment, perform an incremental sync of all domains before installing Fix Pack 003.
2. Create a backup of the database. To create a backup of a DB2® database, complete the following steps:
   a. Stop the Configuration Discovery and Tracking server.
   b. Use one of the following procedures:
      - For Linux, Solaris, AIX, and Linux on System z operating systems, log in as the DB2 database instance owner. For example, you could use the **db2inst1** ID.
      - For Windows operating systems, open the DB2CMD command prompt.
   c. Run the following command:

      ```
      db2 backup database cmdb
      ```

      Replace *cmdb* with the name of your database.

   If you use an Oracle database, refer to the Oracle documentation for instructions about how to back up the database.
3. If you created any custom database views or triggers, drop them. If you do not drop the custom database views and triggers, the installation process can not complete successfully.
4. Create a complete backup of the Configuration Discovery and Tracking server files. Depending on the operating system, use the zip or tar command to zip the entire directory. For Linux, Solaris, AIX, and Linux on System z operating systems, you want to tar the /opt/IBM/cmdb directory. For Windows operating systems, you want to zip the C:\ibm\cmdb directory.
5. If you have made changes to the cmdb-context.xml file, save a copy of the cmdb-context.xml file. For Linux, Solaris, AIX, and Linux on System z operating systems, the file is usually located in the /opt/IBM/cmdb/dist/deploy-tomcat/ROOT/WEB-INF/cmdb-context.xml directory. For Windows operating systems, the file is usually located in the c:\ibm\cmdb\dist\deploy-tomcat\ROOT\WEB-INF\cmdb-context.xml directory.

After the installation process is complete, you need to merge the contents of the cmdb-context.xml file with the contents of the cmdb-context.xml file installed with Fix Pack 003 before starting the Configuration Discovery and Tracking server.

In addition, contact support to determine if the changes to the cmdb-context.xml file can alternatively be defined in the collation.properties file to prevent this issue in the future.

6. Complete this step if you installed Fix Pack 002. If you did not install Fix Pack 002, skip this step. This step is for upgrading a Domain Configuration Discovery and Tracking server and a Enterprise Configuration Discovery and Tracking server. If the Configuration Discovery and Tracking server uses an Oracle backend database, complete the following steps:
   a. Change to the dbscripts directory.
   b. Edit the fp3versiongrants.sql file.
   c. Replace *taddmusr* with *primary oracle database user* setup for Configuration Discovery and Tracking.
   d. Run the fp3versiongrants.sql script using the *archive oracle user* as the login.

Installing Fix Pack 003 for the Configuration Discovery and Tracking feature

When installing this fix pack, for Linux, Solaris, AIX, and Linux on System z operating systems, if you installed the Configuration Discovery and Tracking server with a root user ID, use a root user ID to install the fix pack. Likewise, if you installed the Configuration Discovery and Tracking server with a non-root user ID, use a non-root user ID to install the fix pack.

For Windows operating systems, use a Windows logon ID with Administrator authority.

(For Linux, Solaris, AIX, and Linux on System z operating systems) If you have non-root user dot profile (for example, .profile , .bashrc , .cshrc) that starts another shell, you need to disable the shell call or the profile before running the install. Otherwise, the installation process hangs.

To install Fix Pack 003 for the Configuration Discovery and Tracking feature, complete the following steps:

1. Open a command prompt window and set the *JAVA_HOME* environment variable:

   For Linux operating systems
   ```
   export JAVA_HOME=$COLLATION_HOME/external/jdk-1.4.2-Linux-i686
   ```

   For Solaris operating systems
   ```
   export JAVA_HOME=$COLLATION_HOME/external/jdk-1.4.2-SunOS-sparc
   ```

   For AIX operating systems
   ```
   export JAVA_HOME=$COLLATION_HOME/external/jdk-1.4.2-AIX-powerpc
   ```

   For Windows operating systems
   ```
   set JAVA_HOME=$COLLATION_HOME\external\jdk-1.4.2-Windows-i386
   ```

   For Linux on System z operating systems
   ```
   export JAVA_HOME=$COLLATION_HOME/external/jdk-1.4.2-Linux-s390
   ```
2. Navigate to the directory where you downloaded the fix pack.
3. Unpack the fix pack file.

4. Change to the fix pack directory.
5. (Linux, Solaris, AIX, and Linux on System z operating systems only) Change the access profile for the installation file with the following command:

   ```
   chmod 755 installFixPack.sh
   ```

6. Run the installation process. Use one of the following commands:

   For Linux, Solaris, AIX, and Linux on System z operating systems:

   ```
   ./installFixPack.sh
   ```

   For Windows operating systems:

   ```
   installFixPack.bat
   ```

   You can run the installation process in silent mode. You do not need a response file for the silent installation mode. To run the silent installation mode, use one of the following commands:

   For Linux, Solaris, AIX, and Linux on System z operating systems:

   ```
   ./installFixPack.sh -silent
   ```

   For Windows operating systems:

   ```
   installFixPack.bat -silent
   ```

7. Click **Next**.
8. Select the **I have reviewed the information above.** checkbox and click **Next**.
9. After the installation process is complete, a page, indicating the status of the installation, is displayed. Click **Finish** to close the installation program.
10. If running in an enterprise environment, for all domains that were upgraded, run the purge recyclebin command using SQLPlus if the backend database is Oracle.
11. If running in an enterprise environment, check that dist/sync/deleteImplicitRelationsOnFullSync does not exist. If the file exists, delete it. Perform a full synchronization to synchronize the new database schema changes. If the file does not exist, a pop-up message is displayed when performing a full synchronization. Click **OK** to ensure implicit relationships are not deleted.
12. If you saved a backup copy of the cmdb-context.xml file, merge the contents of the backup copy with the contents of the cmdb-context.xml file installed as a part of Fix Pack 003. For Linux, Solaris, AIX, and Linux on System z operating systems, the file is usually located in the /opt/IBM/cmdb/dist/deploy-tomcat/ROOT/WEB-INF/cmdb-context.xml directory. For Windows operating systems, the file is usually located in the c:\ibm\cmdb\dist\deploy-tomcat\ROOT\WEB-INF\cmdb-context.xml directory. Complete this step before starting the Configuration Discovery and Tracking server.

If the installation fails, the log files contain information that you can use to troubleshoot the problem.

After you install the Configuration Discovery and Tracking server, you may have to clear the Java Webstart cache on each client machine where you log in to the Product Console. You must clear the Java Webstart cache on each client machine where you log in to the Product Console after an upgrade.

To clear the Java Webstart cache, complete the following steps:

1. Go to the directory where Java 1.5 is installed.
2. Change directories to the \jre\bin directory.
3. Double-click on **javaws.exe**.

a. In the Java Application Cache Viewer window, click **Edit → Preferences**.
b. In the Java Control panel, under the Temporary Internet Files section, click **Delete Files**, check all boxes, and click **OK**.
c. Click **OK**.

Locating log files

The software generates log files that contain message information. You can access this information for general use or when requested by IBM Software Support.

For Linux, Solaris, and AIX operating systems, message log files can be found in the following directory: <*install_dir*>/installLogs.

For Windows operating systems, message log files can be found in the following directory: <*install_dir*>\installLogs.

For Linux for System z operating systems, message log files can be found in the following directory: <*install_dir*>/installLogs.

Upgrading the database manually

If the database failed to upgrade during the fix pack installation process, you receive a message explaining that the database failed to upgrade and that you need to manually upgrade the database. Instructions for manually upgrading the database follow. You can complete these steps multiple times.

To manually upgrade the Enterprise Configuration Discovery and Tracking server, complete all of the following steps, but skip the last two steps. To manually upgrade the Domain Configuration Discovery and Tracking server, complete all of the following steps:

1. Stop the installation process.
2. Restore your backup database. Do not restore your backup server files.
3. Change to the directory where the Configuration Discovery and Tracking server is installed. For example, the /opt/IBM/cmdb/dist/bin directory.
4. Use one of the following procedures:
   a. (For Linux, Solaris, and AIX operating systems) To upgrade from 5.1.1 or 5.1.1.1, use one of the following procedures:
      - If you use a DB2 database, run the following commands to upgrade the database:

        ```
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/
          DB210_ModSchema.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/DB220_Meta.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/DB230_Indexes.sql
        ```

      - If you use an Oracle database, run the following commands to upgrade the database:

        ```
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/
           Oracle10_ModSchema.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/Oracle20_Meta.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/Oracle30_Indexes.sql
        ```

      Errors are reported in the command prompt window. If you have a lot of discovery data in the database, the commands can take a long time to run.
   b. (For Linux, Solaris, AIX, and Linux on System z operating systems) To upgrade from 5.1.1.2, use one of the following procedures:
      - If you use a DB2 database, run the following commands to upgrade the database:

```
db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/
  DB210_ModSchema.sql
db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/DB230_Indexes.sql
db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/DB220_Meta.sql
```

- If you use an Oracle database, run the following commands to upgrade the database:

```
db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/
   Oracle10_ModSchema.sql
db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/Oracle30_Indexes.sql
```

Errors are reported in the command prompt window. If you have a lot of discovery data in the database, the commands can take a long time to run.

c. (For Windows operating systems) To upgrade from 5.1.1.2-TIV-ITADDM-LA0001, use one of the following procedures:

- If you use a DB2 database, run the following commands to upgrade the database:

```
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\
  DB210_ModSchema.sql
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\DB230_Indexes.sql
```

- If you use an Oracle database, run the following commands to upgrade the database:

```
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\
   Oracle10_ModSchema.sql
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\Oracle30_Indexes.sql
```

Errors are reported in the command prompt window. If you have a lot of discovery data in the database, the commands can take a long time to run.

5. If you are upgrading a Domain Configuration Discovery and Tracking server, not an Enterprise Configuration Discovery and Tracking server, use one of the following commands:

For Linux, Solaris, AIX, and Linux on System z operating systems:
/opt/IBM/cmdb/dist/bin/template-upgrade.sh

For Windows operating systems:
C:\IBM\cmdb\dist\bin\template-upgrade.bat

6. If you are upgrading a Domain Configuration Discovery and Tracking server, not an Enterprise Configuration Discovery and Tracking server, use one of the following commands to create the reporting views:

For Linux, Solaris, AIX, and Linux on System z operating systems:
For a DB2 database: /opt/IBM/cmdb/dist/bin/db-upgrade.sh /opt/IBM/cmdb/dist/support/bin/make_db2_reporting_views.sql

For Windows operating systems:
For a DB2 database: c:\IBM\cmdb\dist\bin\db-upgrade.bat c:\IBM\cmdb\dist\support\bin\make_db2_reporting_views.sql

Roll back for the Configuration Discovery and Tracking feature Fix Pack 003

To roll back Fix Pack 003 for the Configuration Discovery and Tracking feature, complete the following steps:

1. Use the following command to stop the server:

```
$COLLATION_HOME/bin/control stop
```

2. Restore the backup files for the Configuration Discovery and Tracking server. (You created a backup of these files before you installed the fix pack.) To restore the backup files for the server, complete the following steps:

a. Go to the /opt/IBM directory.

b. Rename the current installation directory to *cmdb.bak*.
c. Unzip the backup files. For example, use the following command:
   `unzip cmdb.zip`

3. For Linux, Solaris, AIX, and Linux on System z operating systems, use the following command to change file ownership to non-root: `chown <non-root_user_id>: <non-root_user_group> <collation_home>` For example, `chown cmdbuser : cmdbgrp /opt/ibm/cmdb/dist`
4. Restore the database. Use one of the following procedures:
   - To restore a DB2 database, complete the following steps:
     a. Log in as the DB2 database instance owner. For example, you could use the **db2inst1** ID.
     b. Run the following command:
        `db2 restore database cmdb`
        Replace *cmdb* with the name of your database.
   - To restore an Oracle database, refer to the Oracle documentation for instructions.
5. Start the Configuration Discovery and Tracking server.

Dropping a database and clearing the View Manager Disk cache

If you are dropping your database and you have the View Manager Disk cache enabled, you have to manually delete the old disk caching directories.

To do this, after you drop the database, go to the directory where you have configured the com.collation.view.cache.disk.path property and delete the contents of the following sub-directories:

- /graph
- /tree

Note: The default directory is dist/var/viewmgr.

Section AB. Using the Tivoli Change and Configuration Management Database installation wizard

If you originally installed the Configuration Discovery and Tracking feature using the Tivoli Change and Configuration Management Database installation wizard, complete the following steps to install Fix Pack 003 using the Tivoli Change and Configuration Management Database installation wizard.

Before you install Fix Pack 003, complete the following tasks:

1. If running in an enterprise environment, perform an incremental sync of all domains before installing Fix Pack 003.
2. Create a backup of the database. To create a backup of a DB2 database, complete the following steps:
   a. Stop the Configuration Discovery and Tracking server.
   b. Use one of the following procedures:
      - For Linux, Solaris, AIX, and Linux on System z operating systems, log in as the DB2 database instance owner. For example, you could use the **db2inst1** ID.
      - For Windows operating systems, open the DB2CMD command prompt.
   c. Run the following command:

`db2 backup database cmdb`

Replace *cmdb* with the name of your database.

If you use an Oracle database, refer to the Oracle documentation for instructions about how to back up the database.

3. If you created any custom database views or triggers, drop them. If you do not drop the custom database views and triggers, the installation process can not complete successfully.
4. Create a complete backup of the Configuration Discovery and Tracking server files. Depending on your operating system, use the zip or tar command to zip the entire directory. For Linux, Solaris, AIX, and Linux on System z operating systems, you want to tar the /opt/IBM/cmdb directory. For Windows operating systems, you want to zip the C:\ibm\cmdb directory.
5. If you have made changes to the cmdb-context.xml file, save a copy of the cmdb-context.xml file. For Linux, Solaris, AIX, and Linux on System z operating systems, the file is usually located in the /opt/IBM/cmdb/dist/deploy-tomcat/ROOT/WEB-INF/cmdb-context.xml directory. For Windows operating systems, the file is usually located in the c:\ibm\cmdb\dist\deploy-tomcat\ROOT\WEB-INF\cmdb-context.xml directory.

   After the installation process is complete, you need to merge the contents of the cmdb-context.xml file with the contents of the cmdb-context.xml file installed with Fix Pack 003 before starting the Configuration Discovery and Tracking server.

   In addition, contact support to determine if the changes to the cmdb-context.xml file can alternatively be defined in the collation.properties file to prevent this issue in the future.
6. Complete this step if you installed Fix Pack 002. If you did not install Fix Pack 002, skip this step. This step is for upgrading a Domain Configuration Discovery and Tracking server and a Enterprise Configuration Discovery and Tracking server. If the Configuration Discovery and Tracking server uses an Oracle backend database, complete the following steps:
   a. Change to the dbscripts directory.
   b. Edit the fp3versiongrants.sql file.
   c. Replace *taddmusr* with *primary oracle database user* setup for Configuration Discovery and Tracking.
   d. Run the fp3versiongrants.sql script using the *archive oracle user* as the login.

Installing Fix Pack 003 for the Configuration Discovery and Tracking feature using the IBM Tivoli CCMDB installation program

To install Fix Pack 003 using the Tivoli Change and Configuration Management Database installation wizard, complete the following steps:

1. Change to the Tivoli Change and Configuration Management Database directory within the fix pack directory and launch the Tivoli Change and Configuration Management Database installation wizard for your operating system:
   - Linux: setupLinux.bin
   - AIX: setupaix
2. The installation wizard detects the previous release and prompts you to upgrade. Follow the panels to enter the required information and start the installation.
3. Upon completion of the installation, the wizard displays a success or failed status message.

4. If you saved a backup copy of the cmdb-context.xml file, merge the contents of the backup copy with the contents of the cmdb-context.xml file installed as a part of Fix Pack 003. For Linux, Solaris, AIX, and Linux on System z operating systems, the file is usually located in the /opt/IBM/cmdb/dist/deploy-tomcat/ROOT/WEB-INF/cmdb-context.xml directory. For Windows operating systems, the file is usually located in the c:\ibm\cmdb\dist\deploy-tomcat\ROOT\WEB-INF\cmdb-context.xml directory. Complete this step before starting the Configuration Discovery and Tracking server.
5. If running in an enterprise environment, for all domains that were upgraded, run the purge recyclebin command using SQLPlus if the backend database is Oracle.
6. If running in an enterprise environment, check that dist/sync/deleteImplicitRelationsOnFullSync does not exist. If the file exists, delete it. Perform a full synchronization to synchronize the new database schema changes. If the file does not exist, a pop-up message is displayed when performing a full synchronization. Click **OK** to ensure implicit relationships are not deleted.

By default, the installation process for the Tivoli Change and Configuration Management Database installation wizard feature sets the **com.collation.topomgr.generateExplicitRelationship** parameter to *true*. However, if you use only the Configuration Discovery and Tracking feature and this parameter is set to *true*, performance is not optimal.

If you plan to add the Process Management and Integration Platform feature of IBM Tivoli CCMDB, leave the parameter set to *true*. In addition, leave the parameter set to *true* if you plan to use the TBSM DLA.

If you do not plan to add the Process Management and Integration Platform feature of IBM Tivoli CCMDB and do not plan to use the TBSM DLA, change the **com.collation.topomgr.generateExplicitRelationship** parameter to *false*. This parameter is located in the collation.properties file (/opt/IBM/cmdb/dist/etc/collation.properties).

If the installation fails, the log files contain information that you can use to troubleshoot the problem.

After you install the Configuration Discovery and Tracking server, you may have to clear the Java Webstart cache on each client machine where you log in to the Product Console. You must clear the Java Webstart cache on each client machine where you log in to the Product Console after an upgrade.

To clear the Java Webstart cache, complete the following steps:

1. Go to the directory where Java 1.5 is installed.
2. Change directories to the \jre\bin directory.
3. Double-click on **javaws.exe**.
   a. In the Java Application Cache Viewer window, click **Edit → Preferences**.
   b. In the Java Control panel, under the Temporary Internet Files section, click **Delete Files**, check all boxes, and click **OK**.
   c. Click **OK**.

Locating log files

The software generates log files that contain message information. You can access this information for general use or when requested by IBM Software Support.

The Fix Pack 003 installation log files are located in the following directories: /CCMDB_1111install.log and /logs/fixpackAll.log

Upgrading the database manually

If the database failed to upgrade during the fix pack installation process, you receive a message explaining that the database failed to upgrade and that you need to manually upgrade the database. Instructions for manually upgrading the database follow. You can complete these steps multiple times.

To manually upgrade the Enterprise Configuration Discovery and Tracking server, complete all of the following steps, but skip the last two steps. To manually upgrade the Domain Configuration Discovery and Tracking server, complete all of the following steps:

1. Stop the installation process.
2. Restore your backup database. Do not restore your backup server files.
3. Change to the directory where the Configuration Discovery and Tracking server is installed. For example, the /opt/IBM/cmdb/dist/bin directory.
4. Use one of the following procedures:
   a. (For Linux, Solaris, and AIX operating systems) To upgrade from 5.1.1 or 5.1.1.1, use one of the following procedures:
      - If you use a DB2 database, run the following commands to upgrade the database:

        ```
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/
          DB210_ModSchema.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/DB220_Meta.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/DB230_Indexes.sql
        ```

      - If you use an Oracle database, run the following commands to upgrade the database:

        ```
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/
           Oracle10_ModSchema.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/Oracle20_Meta.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.1/Oracle30_Indexes.sql
        ```

      Errors are reported in the command prompt window. If you have a lot of discovery data in the database, the commands can take a long time to run.
   b. (For Linux, Solaris, AIX, and Linux on System z operating systems) To upgrade from 5.1.1.2, use one of the following procedures:
      - If you use a DB2 database, run the following commands to upgrade the database:

        ```
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/
          DB210_ModSchema.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/DB230_Indexes.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/DB220_Meta.sql
        ```

      - If you use an Oracle database, run the following commands to upgrade the database:

        ```
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/
           Oracle10_ModSchema.sql
        db-update.sh /opt/IBM/cmdb/dist/etc/schema_upgrade/5.1.1.2/Oracle30_Indexes.sql
        ```

      Errors are reported in the command prompt window. If you have a lot of discovery data in the database, the commands can take a long time to run.
   c. (For Windows operating systems) To upgrade from 5.1.1.2-TIV-ITADDM-LA0001, use one of the following procedures:
      - If you use a DB2 database, run the following commands to upgrade the database:

```
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\
  DB210_ModSchema.sql
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\DB230_Indexes.sql
```

- If you use an Oracle database, run the following commands to upgrade the database:

```
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\
   Oracle10_ModSchema.sql
db-update.bat c:\IBM\cmdb\dist\etc\schema_upgrade\5.1.1.2\Oracle30_Indexes.sql
```

Errors are reported in the command prompt window. If you have a lot of discovery data in the database, the commands can take a long time to run.

5. If you are upgrading a Domain Configuration Discovery and Tracking server, not an Enterprise Configuration Discovery and Tracking server, use one of the following commands:

   For Linux, Solaris, AIX, and Linux on System z operating systems:
   /opt/IBM/cmdb/dist/bin/template-upgrade.sh

   For Windows operating systems:
   C:\IBM\cmdb\dist\bin\template-upgrade.bat

6. If you are upgrading a Domain Configuration Discovery and Tracking server, not an Enterprise Configuration Discovery and Tracking server, use one of the following commands to create the reporting views:

   For Linux, Solaris, AIX, and Linux on System z operating systems:
   For a DB2 database: /opt/IBM/cmdb/dist/bin/db-upgrade.sh
   /opt/IBM/cmdb/dist/support/bin/make_db2_reporting_views.sql

   For Windows operating systems:
   For a DB2 database: c:\IBM\cmdb\dist\bin\db-upgrade.bat
   c:\IBM\cmdb\dist\support\bin\make_db2_reporting_views.sql

Roll back Fix Pack 003 when using the Tivoli Change and Configuration Management Database installation wizard

Roll back for Fix Pack 003 when using the Tivoli Change and Configuration Management Database installation wizard is not supported.

Dropping a database and clearing the View Manager Disk cache

If you are dropping your database and you have the View Manager Disk cache enabled, you have to manually delete the old disk caching directories.

To do this, after you drop the database, go to the directory where you have configured the com.collation.view.cache.disk.path property and delete the contents of the following sub-directories:

- /graph
- /tree

Note: The default directory is dist/var/viewmgr.

Section B. Installing Fix Pack 003 for the Process Management and Integration Platform feature

This section provides fix pack installation instructions for the Process Management and Integration Platform feature.

Installing Fix Pack 003 for the Process Management and Integration Platform feature

To install this fix pack for the Process Management and Integration Platform feature, complete the steps described in this section.

Before installing Fix Pack 003, install the latest IBM WebSphere® Application Server Java SDK Interim Fix for V6.0. Visit http://www.ibm.com/support/docview.wss?rs=180&uid=swg24011133 to download and install the most recent interim fix for your operating system.

Note: It is important to make a make a copy of the disk image of the system on which you are planning to install the fix pack. There is no automated uninstall feature supplied with this fix pack. If the installation fails, you have to restore the system to its previous working state using the copy of the disk image prior to attempting the installation again.

In order to upgrade successfully, you must ensure that there are no tasks running. Follow these steps to prepare to upgrade:

1. Terminate all running tasks:
   a. Point your browser to http://<*hostname*>9080/bpc, where <*hostname*> is the host name on which IBM® WebSphere® Portal Server is installed.
   b. Log in as wasadmin.
   c. Click **Task Instances > Administered By Me**.
   d. Select **All tasks**.
   e. Click **Terminate**.
2. Stop all Task Module applications that have human tasks or BPEL processes running:
   a. Log on to the WebSphere Administration console as wasadmin.
   b. Click **Applications > Enterprise Applications**.
   c. Click each of the following applications and repeat steps d through k to stop all BPEL and human tasks associated with the applications:
      - CTGCMChMTasksModuleApp
      - CTGWETasksModule
   d. Click on the CTG application link.
   e. Go to **Related Items > EJB Modules** and click on each jar file link.
   f. Click **Business processes**.
   g. Select and stop any task entries that are listed.
   h. Click **Save**.
   i. Click **Human tasks**.
   j. Select and stop any task entries that are listed.
   k. Click **Save**.

To install the fix pack, complete the following steps:

1. Ensure that the WebSphere Application Server and Tivoli Directory Server are running.
2. Open a command prompt and source the DB2 profile used for the process database. For example:

```
Windows: set DB2INSTANCE=ctginst1
Linux: ./home/ctginst1/sqllib/db2profile
```

From this same command prompt, start the WebSphere Portal Server.

3. Change to the IBM Tivoli CCMDB directory within the fix pack directory and launch the IBM Tivoli CCMDB installation wizard for your operating system:

```
Windows: setupwin32.exe
Linux: setupLinux.bin
AIX: setupaix
```

4. The installation wizard detects the previous release and prompts you to upgrade. Follow the panels to enter the required information and start the installation. You must enter user ids and passwords for the CMDB database (if it is a DB2 database), the process database, WebSphere Application Server, and WebSphere Portal Server.
5. Upon completion of the installation, the wizard displays a success or failed status.

   Note: The Fix Pack installation log files are located in: <CCMDB_install_dir>/CCMDB_1112_install.log and <CCMDB_install_dir>/logs/fixpackAll.log.

After completing the upgrade, verify that the correct shared libraries are associated with each application, and that all the human tasks and business process associated with each application are running. Follow these steps to verify the shared libraries:

1. Log on to the WebSphere Administration console as wasadmin.
2. Click on the name of each application listed in step 3 and click **Additional properties > Libraries**.
3. Check whether the appropriate libraries are listed as associated:

Table 2. Applications and associated shared libraries

Application	Required libraries
CTGPMCMChangeHistoryCIApp	CTGFoundationCommonLib, CTGPMCfgMgmtLib, CdbApiClientLibrary
CTGCMChMTasksModuleApp	CTGWE_DYNFLOW, CdbApiClientLibrary
CTGWETasksModuleApp	CTGWE_CHGMGMT, CTGWE_DYNFLOW, CdbApiClientALibrary

4. If any library in this table is not listed, click **Add** and choose each missing library from the list to add it.
5. If you added any shared libraries to any application, click **OK** and then click **Save** when you are finished.

Follow these steps to verify that all human tasks and business processes are started:

1. Log on to the WebSphere Administration console as wasadmin.
2. Click **Applications > Enterprise Applications**.
3. Click each of the following applications and repeat steps 4 through 11to stop all BPEL and human tasks associated with the applications:
   - CTGCMChMTasksModuleApp
   - CTGWETasksModule
   - CTGPMCMChangeHistoryCIApp
4. Click on the CTG application link.
5. Go to **Related Items > EJB Modules** and click on each jar file link.
6. Click **Business processes**.

7. Select and start any task entries that are in the Stopped state.
8. Click **Save**.
9. Click **Human tasks**.
10. Select and start any task entries that are in the Stopped state.
11. Click **Save**.

Roll back for the Process Management and Integration Platform feature Fix Pack 003

Roll back for the Process Management and Integration Platform feature Fix Pack 003 is not supported.

Dropping a database and clearing the View Manager Disk cache

If you are dropping your database and you have the View Manager Disk cache enabled, you have to manually delete the old disk caching directories.

To do this, after you drop the database, go to the directory where you have configured the com.collation.view.cache.disk.path property and delete the contents of the following sub-directories:

- /graph
- /tree

Note: The default directory is dist/var/viewmgr.

Chapter 6. Configuring Configuration Discovery and Tracking

You do not have to complete any specific configuration tasks before using Configuration Discovery and Tracking. This section provides instructions for optional configuration tasks that can be useful in your environment.

Configuring for sensors and DLA data merging

The data merging logic in Fix Pack 003 implements a course-grained priority system for incoming data.

Data from different sources, including discovery systems, is assigned a numerical priority that determines whether the incoming data can overwrite existing data.

When data from two different sources describes the same object, the priority of the sources is compared. Only if the incoming data source has a lower priority number, is it allowed to overwrite existing data.

By default, the DLA loading process has the highest priority (the lowest priority = 0) in the system. This means that DLA data can always overwrite data discovered by Configuration Discovery and Tracking sensors. In many cases, this is the best setting for the system.

In some cases, it may be desirable to allow Configuration Discovery and Tracking sensors to overwrite DLA data. For example, when loading data from the z/OS DLA with WebSphere attribute information populated, the priority property would be set in this case so the WebSphere attributes populated by the WebSphere discovery sensor will merge with the DLA data.

If this the case, the DLA loading process priority can be lowered by setting the property `com.collation.bulkload.data.priority-10` in the *$COLLATION_HOME*/etc/collation.properties file. After setting this property, restart the Configuration Discovery and Tracking server.

From this point on, until the priority value is changed, DLA data has approximately equivalent priority as Configuration Discovery and Tracking discovered data, allowing either discovered data or DLA data to overwrite each other's values. Data coming from the APIs of Configuration Discovery and Tracking continue to have a priority of 0, regardless of the property setting in the properties file.

In cases where Configuration Discovery and Tracking discovered data has different details about an object than DLA data, it is possible for some details from the Configuration Discovery and Tracking to be lost. This is a result of the merging and consolidation logic that takes place in the product. The only solution for this situation currently is to alter the loading priority of DLAs so as not to overwrite discovered data.

Configuring Linux for System z systems

For Linux for System z environments, if you have trouble starting the Configuration Discovery and Tracking server, perform some configuration tasks to enable the start up process.

The following list provides information about errors that can be reported in log files during the start up process and ways to solve the problems. To see the errors, the log level needs to be set to *DEBUG*.

- Error: In the tomcat.log file:

  ```
  rmid: (WARNING) restart service throws:java.rmi.activation.ActivationException:
  timeout creating child process
  ```

  Solution: In the collation.properties file, increase the default value of the `com.collation.jini.rmidtimeout = 30000` property to *300000*. Start the Configuration Discovery and Tracking server.
- Error: In the ClientProxy.log file:

  ```
  jini.JiniServiceFactory - retry limit reached breaking
  ```

  Solution: In the collation.properties file, increase the default value of the `com.collation.jini.service.timeout = 10` property to *100*. Start the Configuration Discovery and Tracking server.
- Error: When starting the bulk load process, by calling the loadidml.sh file, either the ″Segmentation fault″ or ″Option too large: -Xmx1024M″ message is displayed.

  Solution: In the bulkload.properties file, decrease the default value of the `com.ibm.cdb.bulk.allocpoolsize = 1024` property. The next allowable size is *768*. If you go any lower than *512*, the server does not run. Start the Configuration Discovery and Tracking server.
- Error: When running a Linux for System z system with the minimum memory requirement of 2 GB, the Configuration Discovery and Tracking server does not start when requesting a Java heap size of 1024 MB.

  Solution: In the <*CMDB_HOME*>/dist/deploy-tomcat/ROOT/WEB-INF/cmdb-context.xml file, go to the <!– Shared VM - Proxy –> section. Decrease the default size of the `jvmArgs -Xmx1024` property from *-Xmx1024* to *-Xmx768*. Start the Configuration Discovery and Tracking server.
- Error: When you configure a system, the system is unable to create another JVM.

  Solution: Some Linux for System z servers do not allow you to reserve large blocks of memory. In the <*CMDB_HOME*>/dist/deploy-tomcat/ROOT/WEB-INF/cmdb-context.xml file, decrease the JVM property (-Xmx) from *1024* to *512*. Start the Configuration Discovery and Tracking server.

Configuring WebSphere for discovery when security is enabled

There are two steps to configuring WebSphere for discovery when security is enabled: access list configuration and certificate set up.

Access list configuration

If security is disabled, user accounts are not needed. If security is enabled, you need the following items configured:

- WebSphere Administrator user ID and password
- Client-side SSL certificate, includes two files: trust and key stores, including their passphrases. The default passphrase is *WebAS*.

If the user has the *Administrator* role, Configuration Discovery and Tracking discovers all configuration data, including security configurations. WebSphere

documentation and support say that the *Monitor* role privilege works for everything, but security configuration. This information has not been validated extensively in the field.

If you have disabled security, that does not mean that you are not using SSL. You need to see if you are prompted for a password when you connect to the WebSphere Application Server Administrator Console. If you need only a username to log in to the Administrator Console, security is disabled. If you need a username and password to log in to the Administrator Console, security is enabled.

Certificate set up

When you set up the access list, including the username, password, and scope limitations, you can click **SSL Settings** to set up the certificates.

The certificates need to be on the machine running the console, which is not necessarily the Configuration Discovery and Tracking server. These certificates can be retrieved from the WebSphere Application Server.

Depending on the version of WebSphere installed, you can find the files DummyClientTrustFile.jks and DummyClientKeyFile.jks in the <*WebSphere Root Directory*>/profiles/default/etc directory. If the files are not located in this directory, do a file search for the dummy client certificate files.

The default passphrase is *WebAS*.

Saving Domain Manager reports

If you want to save Configuration Discovery and Tracking server and Enterprise Configuration Discovery and Tracking server Domain Manager reports in the PDF, CSV, or XML formats, update the **com.collation.domain.pdfreport.enabled** parameter.

To update the **com.collation.domain.pdfreport.enabled** parameter, complete the following steps:

1. Open the collation.properties file (/opt/IBM/cmdb/dist/etc/collation.properties) in a text editor.
2. Change the **com.collation.domain.pdfreport.enabled** parameter from *false* to *true*.
3. Save and close the collation.properties file.
4. Restart Configuration Discovery and Tracking.

Bidirectional support is not provided for saving Domain Manager reports in the PDF, CSV, or XML formats. In addition, the instructions provided in this section have been tested with the English locale; the instructions may not work with non-English locales.

Updating the collation.properties file

Throughout this document there are references to the collation.properties file. This file is located in the *$COLLATION_HOME*/etc/ directory.

Many properties are specified in the collation.properties file, including the platform properties. The defaults for platform properties are specified in the

collation.properties file. When working with discovery profiles, you can change a platform property for a profile. If you do not change the platform properties, the default properties, in the collation.properties file, are used.

Upgrading the IBM CMDB Enterprise JavaBean Application Programming Interface

If you are upgrading the CMDB Enterprise JavaBean (EJB) Application Programming Interface (API) for Fix Pack 003, make sure the following property is added to the <*CMDB_WebSphere*>/etc/collation.properties file:

```
com.ibm.client.type=ejbclient
```

<*CMDB_WebSphere*> is the directory where the CMDB_WebSphere.zip file is installed.

See the *Configuration Discovery and Tracking SDK Developer's Guide* for information on installing and configuring the CMDB EJB API.

Chapter 7. Limitations, known problems and workarounds

The following sections contain information about the limitations and problems that apply to this fix pack.

Limitations

This section describes limitations in this fix pack. Where applicable and known, workarounds are identified.

Discovering SAP for Windows operating systems is not supported

Problem: Currently, the SAP sensor is not supported on the Configuration Discovery and Tracking server running on the Windows operating systems.

Workaround: No workaround is identified. Contact support for the estimated time of availability.

Support for application descriptors by the Microsoft System Management Server sensor

Limitation: The Microsoft System Management Server (SMS) sensor does not support application descriptors in Fix Pack 003.

Workaround: To check on the status of SMS sensor support for application descriptors, contact support.

Support for federation of external sources

Limitation: The federation of external sources is not supported by Fix Pack 003.

Workaround: To check on the status of support for the federation of external sources, contact support.

Support for the self monitoring tool

Limitation: The self monitoring tool is not supported by Fix Pack 003.

Workaround: To check on the status of support for the self monitoring tool, contact support.

Known problems and workarounds

The following sections identify problems that might occur during the use of this fix pack. Where available, workaround solutions are provided for the problems.

Anchors do not start

Problem: The anchor does not start if the sshd daemons are not configured to allow port forwarding. Anchors rely on local port forwarding to be enabled in the remote anchor server.

Workaround: The configuration file for the sshd daemon should be located in either the /etc/ssh or /usr/local/etc/ssh directory. Open the sshd_config file and ensure the **AllowTcpForwarding** option is set to *yes*. The default is *yes*, but there is a chance that when an IBM server is installed, the option is changed to *no*.

If you use SSH2, SSH2 adds four **TcpForwarding** options. These options restrict forwarding for groups and users. These options must be enabled or absent.

Application Summary Details report does not contain data

Problem: The Application Summary Details report in the Enterprise Configuration Discovery and Tracking server Domain Manager Portal is empty.

Workaround: Complete the following steps:

1. Edit the dist/etc/domainquery file.
2. Add the following line to the bottom of the file:

   ```
   collection:name_x,groupname_x
   ```
3. Perform a full synchronization for the Enterprise Configuration Discovery and Tracking server.

Cannot install the Configuration Discovery and Tracking server to a non-ASCII directory

Problem: The Configuration Discovery and Tracking server installation path string must be only ASCII characters. If the path contains non-ASCII characters, the Configuration Discovery and Tracking server will not start, even if the installation was successful.

Workaround: Use only ASCII characters in your Configuration Discovery and Tracking server installation path.

Configuration Discovery and Tracking cannot be installed to path containing a space

Problem: On the supported Windows operating systems, you cannot install Configuration Discovery and Tracking to a path containing a space (″ ″). There is a problem with GigaSpaces, the component which starts up all of our services. The GigaSpaces component uses third party code that has a problem handling spaces in the path name.

Workaround: Install Configuration Discovery and Tracking in a path with a name that does not contain spaces.

Creating configuration item fails (IBM Tivoli CCMDB)

Problem: When using IBM Tivoli CCMDB to create a configuration item, a message reports that the configuration item creation process failed.

Workaround: Use IBM Tivoli CCMDB to create a configuration item, but complete only the required fields. Do not type any data in the non-required fields. The configuration item should be created. When the item is displayed on the user interface, edit the configuration item to include additional information.

Discovery of Windows 2000 machines can intermittently fail

Problem: The discovery of Windows 2000 machines can intermittently fail.

Workaround: Rediscover only the failing machines. To check on the status of a fix for this problem, contact support. Do not apply this fix pack if your production installation requires consistent Windows 2000 discovery.

Error when launching the Configuration Discovery and Tracking Product Console from IBM Tivoli CCMDB

Problem: After Fix Pack 003 is installed, manually give the servlet authorization to all users from the WebSphere Application Server Administrator Console.

1. Log in to the Web Sphere Application Server Administrator Console.
2. Click **Applications** → **Enterprise Applications**.
3. Open UICfgMgmt_XXXXXX.
4. Check **Map security roles to users/groups**.
5. Check **All authenticated**.
6. Click **OK**.
7. Click **Save**.
8. Restart the Portal Server.

If this set of steps does not launch the Configuration Discovery and Tracking Product Console, you receive the following error message:

```
Error 500: java.rmi.AccessException: CORBA NO_PERMISSION...
```

Workaround: To fix the problem, manually give CDBAPI_EAR all user access. Complete the following steps:

1. Log in to the WebSphere Application Server Administrator Console.
2. Click **Applications** → **Enterprise Applications**.
3. Open CDBAPI_EAR.
4. Check **Map security roles to users/groups**.
5. Check **Everyone**.
6. Click **OK**.
7. Click **Save**.
8. Restart CDBAPI_EAR.

Explicit relationship information on discovered resources is not provided

Problem: If the **com.collation/topomgr.generateExplicitRelationship** parameter in the collation.properties file is set to *false* (which is the default), explicit relationship information on resources discovered by Configuration Discovery and Tracking is not gathered. If you change the parameter to *true* and Configuration Discovery and Tracking receives a query for explicit relationship information, only information

about resources that are discovered after the parameter was set to *true* is provided. However, explicit relationship information cannot be provided on resources discovered while the flag was set to *false*.

Workaround: If the **com.collation/topomgr.generateExplicitRelationship** parameter was set to *false* when you started discovering resources, and you change the flag to *true*, you must run the explicitrel.sh script or call the **generateExplicitRelationships** API. This gathers explicit relationship information on resources that were discovered when the flag was set to *false*. Therefore, when Configuration Discovery and Tracking receives a query for explicit relationship information, all information will be provided regardless of the setting of **com.collation/topomgr.generateExplicitRelationship** parameter when the resources were discovered.

This can take a long time to run, especially with a DB2 database. If the transaction log grows too large, the explicitrel.sh script can be run with the **autocommit** parameter, or, if using the **generateExplicitRelationships** API, the com.collation.explicit.autocommit property can be set to *true*. This increases the length of time it takes to run, but prevents the transaction log size issue by running as multiple transactions. If the script or the API fails with **autocommit**, it must be rerun to success or the database must be restored.

Extended attributes are not displayed in the Details panel of an object

Problem: Extended attributes are not displayed in the Details panel of an object if the values for the attributes are not set.

Workaround: Set values for the extended attributes and they will be displayed in the Details panel.

Fix Pack 003 installation process hangs

Problem: The Fix Pack 003 installation process hangs when a non-root user does not have write permissions to the directory where the user wants to install the product.

Workaround: Complete the following steps:

1. Stop the installation process.
2. Obtain write permissions for the directory where you want to install the product, or select a directory where you have write access.
3. Restart the installation process.

IP Range information missing for the Details panel of the SMS Site

Problem: When you log in to the Product Console and select **Show Details** for the SMS Site, on the Boundaries tab, the IP Range is not displayed.

Workaround: No workaround is identified.

Multiple copies of same computer system are displayed

Problem: On the Product Console, multiple copies of the same computer system are displayed when both of the following actions occur:

- You run the Tivoli Monitoring 6.1 Fix Pack 003 and later DLA and use the bulk load program to load the IDML book to Configuration Discovery and Tracking.
- You perform a sensor discovery of systems where Tivoli Monitoring agents or servers reside.

If you query the computer systems using the API, multiple copies of the same computer system are also displayed.

Workaround: Use the reconciliation tool (the ldfxidml.sh script) instead of the bulk load program (the loadidml.sh script) to load Discovery Library books into the Configuration Management Database. The readme files for the reconciliation tool (*$COLLATION_HOME*/etc/ldfxidml.readme) and bulk load program (*$COLLATION_HOME*/etc/loadidml.readme) provide more information.

Network configuration on Linux for System z systems does not create packets that Nmap can read

Network configuration on Linux for System z systems does not create packets that Nmap can read.

Problem: The StackScan sensor uses Nmap to gather data about the targets for credential-less discovery. If Nmap is not working properly, the StackScan sensor will not leverage Nmap and only use the hostscanner capability. It can run without errors, but the Linux for System z system running the StackScan sensor returns the following message:

```
stored - 0 ComputerSystems in the database
```

If you type the `nmap <hostname>` command for any system other than the local host, the following message is displayed:

```
Note: Host seems down. If it is really up, but blocking our ping probes, try -P0...
```

Workaround for SUSE Linux for System z: The network needs to run with the following option:

```
QETH_OPTIONS='fake_ll=1'
```

This option should be added to the configuration file for the NIC. Depending on the NIC that is used, the exact name of the file changes.

Contact your system administrator for the exact name of the configuration file used by your system.

The configuration file should be in the /etc/sysconfig/hardware directory. The file name could be hwcfg-qeth-bus-ccw-0.0.5000.

Workaround for RedHat Linux for System z: The network needs to run with the following option:

```
OPTIONS='fake_11=1'
```

This option should be added to the configuration file for the NIC. Depending on the NIC that is used, the exact name of the file changes.

Contact your system administrator for the exact name of the configuration file used by your system.

The configuration file should be in the /etc/sysconfig/network-scripts directory. The file name could be ifcfg-eth0.

Verify that the alias in /etc/modprobe.conf contains the following:

```
alias eth0 qeth
```

On Windows operating systems, changing com.collation.websphere.root.dir for WebSphere discovery could cause problems

Problem: When the Configuration Discovery and Tracking server is running on Windows and the `com.collation.websphere.root.dir` property in the collation.properties file is changed to point to the WebSphere jar files on the system, the path should escape out the backslashes.

For example, if the path to the WebSphere jars is C:\IBM\WebSphere\AppServer\, then the path should look like this: `com.collation.websphere.root.dir=C:\\IBM\\WebSphere\\AppServer\\`

If there is a space in the path like this: C:\Program Files\IBM\WebSphere\AppServer\, an error message is displayed.

Workaround: Copy <*WebSphere Root Directory*>/lib, <*WebSphere Root Directory*>/java/jre/lib/ext, and <*WebSphere Root Directory*>/properties to a directory without a space in the path.

Process/thread limit on Linux operating systems

Problem: On a Linux operating system, each thread is considered a process, so you may have an issue with the process/thread (nproc) limit.

Workaround: When Configuration Discovery and Tracking is installed on a Linux operating system, there is an additional entry that should be added to the /etc/security/limits.conf file. To solve this issue, you have to add two lines to the /etc/security/limits.conf file at installation time. Complete the following steps to fix this issue:

In the /etc/security/limits.conf file, add these two lines to the file:

```
coll hard nproc 16000
coll soft nproc 16000
```

Replace *coll* with the Configuration Discovery and Tracking service account user.

Query of either the Hardware Management Console (HMC), IVM, or Virtual I/O Server (VIOS) objects in the Domain Manager does not work

Problem: When you attempt to query either the Hardware Management Console (HMC), Integrated Virtualization Manager (IVM), or Virtual I/O Server (VIOS) objects in the Domain Manager, the query fails.

Workaround: When querying for HMC, VIOS or IVM from the Domain Manager, you have to query for ComputerSystem Objects providing the proper object type. The following query is an example of how to use the Domain Manager to query for HMC Systems whose displayName contains 'ibm':

```
localhost.ComputerSystem.type equals 'HMC' && localhost.ComputerSystem.displayName
contains 'ibm'
```

Running Nmap on AIX 5.3

Problem: Running credential-less discovery on an AIX Configuration Discovery and Tracking server guesses the target operating system with low confidence, incorrectly, or not at all.

The binary Nmap improves the results, but Nmap is not available for AIX 5.3.

Workaround: Complete the following steps:

1. Go to http://aixpdslib.seas.ucla.edu/packages/nmap.html.
2. Download Nmap 3.0. The Web site states the version of Nmap is for AIX 4.2. This is the version of Nmap you want to run on AIX 5.3.
3. Put the file in the root directory (/).
4. Run the nmap.3.00.tar.Z command.
5. Run the tar xvf nmap.3.00.tar command.

The executable Nmap is located in the /usr/local/bin directory. If not already included, change the root profile to include the /usr/local/bin directory.

You also need to load the DLPI module. Complete the following steps:

1. Use the strload -q -d dlpi command to check if it is active or not.
2. If this command does not respond with *dlpi: yes*, then enable it using the strload -f /etc/dlpi.conf command.

The host machine reverts back to default behavior after a reboot. The default behavior is defined in the /etc/pse.conf file. If you want to permanently enable DLPI, uncomment any lines in this file that contain references to dlpi.

To run Nmap, from the command line, run the following command:

```
nmap <hostname or ip>
```

There are many options available, use `nmap -h` to show all options.

For the StackScan sensor to run on an AIX Configuration Discovery and Tracking server, you need to grant sudo NOPASSWD access to the non-root Configuration Discovery and Tracking user.

Session sensor may not start on some targets

Problem: When running a discovery without profiles or running discovery with a profile that includes the Ping sensor and the StackScan sensor, the Session sensor may not start on some targets.

Workaround: Run the discovery using a profile. Select either the Ping sensor or the StackScan sensor. Do not select both Ping and StackScan sensors. If you do not have Nmap in your environment, select the Ping sensor.

A System p topograph icon is not created when you discover a System p machine

Problem: If you discover a System p™ machine and then look at the topology graph, the System p icon is not displayed. Once you discover a non-System p machine, and then view the topology graph, the System p machine that you discovered earlier displays.

Workaround: Discover a non-System p machine in order to get a System p machine icon to display in the topology graph.

Timeout while discovering WebSphere MQ on Windows operating systems

Problem: A timeout might occur during the discovery of WebSphere MQ V5.3 or V6 on the Windows operating systems because Windows WMI (Windows Management Instrumentation) runs very slowly.

Workaround: You can increase the timeout interval by adding the following line to the collation.properties file in the Agent Settings section:

```
com.collation.discover.agent.MQServerAgent.timeout=X
```

where *X* is the timeout value in milliseconds. The default value is 600000.

Unable to start discovery using discovery profiles

Problem: The ProfileManager is unable to get discovery profiles from the database and throws the following exception:

```
java.lang.Exception [PLATFORM.JINI.E.0] The application cannot
find the Jini Service.
```

Workaround: Complete the following steps:

1. Open the collation.properties file.
2. Find the com.collation.jini.service.retries property. Set the value to *200*.
3. Restart the server.

Uninstallation process on Windows completes without error, but database is not dropped

Problem: During the installation process, the following message is displayed:

```
The database already exists. Do you want to continue?
```

Workaround: There are two options:

1. Manually drop the database using the following command: db2 drop database <db_name>. Continue with the install.
2. Continue with the installation process and complete the following steps after the installation process is complete:
   a. Stop the Configuration Discovery and Tracking server.
   b. Open a DB2 Windows command prompt and run the following command: dist/bin/support/bin/make_db2_db <db_name>

      The command drops the existing database and creates a new database.

c. Restart the Configuration Discovery and Tracking server.

Using IDD StackScan sensor with SELinux does not work

Problem: On systems that have Security-enhanced Linux (SELinux), the IDD StackScan sensor runs, but the operating system of the each discovered target is not stored.

This problem occurs if the security level on the Configuration Discovery and Tracking server is high.

Workaround: To determine whether your system has SELinux, and at which level it is set, complete the following steps:

1. On a Linux command line on the system running the Configuration Discovery and Tracking server, log on as *root*.
2. Type *getenforce* and press Enter.

If the *getenforce* command returns a value of *permissive* or *disabled*, the StackScan sensor will run properly.

If the *getenforce* command returns a value of *enforcing*, you must change the security level in one of the following ways:

- Temporarily change the security level to *permissive* in order to decrease the security level.
- Permanently change the security level to one of the following values:
  - *Disable*, in order to disable SELinux.
  - *Permissive*, in order to decrease the security level.

To temporarily lower the security level, complete the following steps:

1. On a Linux command line, type `setenforce 0` and press Enter to set SELinux to *permissive*.
2. Run another discovery using the StackScan sensor.
3. After the discovery has completed, you can enable SELinux by typing `setenforce 1` and pressing Enter to set SELinux to *enforcing*.

To permanently change the security level to *disable* or *permissive*, edit the /etc/sysconfig/selinux file and reboot the machine.

WebSphere discovery does not work on Windows operating systems

Problem: When the Configuration Discovery and Tracking feature runs on a Windows machine, discovery fails for WebSphere discovery when the `com.collation.websphere.performance.setting` property is set to *true*:

Workaround: By default, this property is set to *false*. Do not change the default value for the `com.collation.websphere.performance.setting` property and the problem does not occur.

Chapter 8. New functionality provided with Fix Pack 003 for the Configuration Discovery and Tracking feature

This section identifies and explains the new functionality that is available for the Configuration Discovery and Tracking feature in Fix Pack 003.

New sensors provided with Fix Pack 003 for the Configuration Discovery and Tracking feature

The following table identifies the new sensors that are delivered with the Configuration Discovery and Tracking feature in Fix Pack 003.

Table 3. New sensors provided by Fix Pack 003 for the Configuration Discovery and Tracking feature

Sensor	Gather configuration information from
Oracle Application Server for Windows (10.1.3)	Oracle Application Servers that run on Windows servers
SMS (2003)	Microsoft System Management Server
Tru64 (OS only; Version 5.x)	Operating system-only sensor and detects version 5.1.

The following sensors are currently undergoing hardening in controlled environments:

- Veritas Cluster
- Veritas Storage Manager
- Solaris Zones
- Sunfire

These sensors are being tested and refined. These sensors are not generally available.

Using the Oracle Application Server for Windows sensor (10.1.3)

Minimum required jars at 10.1.3

For a successful discovery of Oracle Application Server, the following list of Oracle Application Server jar files need to be installed on the Configuration Discovery and Tracking server. These jar files can be copied under the Configuration Discovery and Tracking installation directory. You can save these files in another directory and point to the files using an entry in the collation.properties file.

- j2ee/home/lib/ejb.jar
- j2ee/home/lib/adminclient.jar
- j2ee/home/lib/javax77.jar
- j2ee/home/lib/jmxcluster.jar
- j2ee/home/lib/jmx_remote_api.jar
- j2ee/home/lib/jmxri.jar
- j2ee/home/oc4jclient.jar
- opmn/lib/argus.jar
- opmn/lib/ons.jar

- opmn/lib/opmnconfig.jar
- opmn/lib/optic.jar
- opmn/lib/repositorycheck.jar

If com.collation.oracleapp.root.dir is set to a relative directory (for example, lib/oracleapp), these files should be in *$COLLATION_HOME*/lib/oracleapp.

If com.collation.oracleapp.root.dir is set to an absolute directory (for example, /home/oracle/product/10.1.3/Oracle_AS), these files should be in /home/oracle/product/10.1.3/Oracle_AS.

In both cases, the files must be readable by the *collation* user.

The following property, in the collation.properties file, must also be set:

```
com.collation.platform.os.ignoreLoopbackProcesses=true
```

Permissions

If permissions on the client box are not set appropriately, the Oracle Application Server sensor can fail when discovering Oracle Application Server on UNIX-based clients. To ensure permissions are correct, run the following commands on UNIX clients when you log in as root or as the Oracle Application Server owner:

```
chmod -R 755 $oracleHome/jdk
chmod -R 755 $oracleHome/opmn/bin
chmod -R 755 $oracleHome/opmn/conf
chmod -R 755 $oracleHome/opmn/lib
chmod -R 755 $oracleHome/Apache/Apache/
chmod 711 $oracleHome/Apache
chmod 711 $oracleHome/Apache/Apache
chmod 711 $oracleHome
chmod 711 $oracleHome/j2ee
chmod 711 $oracleHome/j2ee/home
chmod 711  $oracleHome/opmn
chmod 711 $oracleHome/lib
chmod 755 $oracleHome/lib/*
```

Using the Microsoft System Management Server sensor

For the Microsoft System Management Server (SMS) sensor discovery, there is no new authentication data introduced. Windows authentication defined for discovery Windows Computer System is used.

Requirements

There must be a SMS WMI provider running on each discovered SMS node.

Additions to collations.properties

The SMS sensor introduces three new properties to the collation.properties file:

com.collation.discover.agent.SMSServerAgent.GetQueries=true
: Enables and disables discovering SMS queries. The default value is *true*.

com.collation.discover.agent.SMSServerAgent.GetClients=true
: Enables and disables discovering information about SMS clients. The default value is *true*.

com.collation.discover.agent.SMSServerAgent.MaxNrClients=100
The maximum number of clients to collect information about, when GetClients is set to *true*. The default value is *100*.

Using the Tru64 UNIX sensor

The HP Tru64 operating system is a UNIX variant, but has some unique commands and some unique flags and options for common commands. The Tru64 sensor is currently designed and implemented to work on Tru64 v5.1 or later.

The Tru64 sensor requires two open source packages to be installed on the target Tru64 machine, sudo and lsof. Both commands should be in the paths available to the non-root user. In addition, both packages should be built on the specific Tru64 machine in question. The most tested versions are sudo-1.6.8p9 and lsof-4.78, however, other versions are likely to work, except in the case where the specific package does not support Tru64. Refer to the distributor Web sites or package README files for a list of restrictions.

For sudo to work properly, the /etc/sudoers file should be edited with the visudo command to contain the following line:

```
<non-rootuser> ANY = NOPASSWD: /sbin/hwmgr
```

For example, if the non-root user is *taddmusr*, to enable that user to run /sbin/hwmgr on any machine, it would look like the following example:

```
taddmusr ANY = NOPASSWD: /sbin/hwmgr
```

Or, to enable this privilege only on the target machine, for example, if the target machine was named *target*:

```
taddmusr target = NOPASSWD: /sbin/hwmgr
```

In addition, there are two commands that must be located in their default locations on the Tru64 resource. One command is /sbin/hwmgr (this command is just like the aforementioned sudoers file) and the other command /usr/sbin/ifconfig.

Without these command in the default locations, the Tru64 sensor is unlikely to work, and the most likely indicators will be one or more errors during discovery with the text `storage error` and the Tru64 resource will only be displayed as an `Other IP Device`. Recheck the locations and permissions on the dependencies and rerun the discovery.

New feature for the MQ sensor

The MQ sensor is able to pick up application descriptors.

The appdescriptors directory with appdescriptors files must be created inside the MQ Queue Manager data directory. For example, on Linux, this directory is /var/mqm/qmgrs/test5/appdescriptors. The test5 directory is a Queue Manager name. Read access to that directory must be granted for the *mqm* user used to discover that manager.

The module level for MQ is a QueueManager.

Enhancement to credential-free discovery allows for shallow discovery of applications

You can use the Level 2 discovery profile to enable shallow discovery of applications running on a target system using only the system credentials by adding the following property to the collation.properties file:

```
com.collation.internalTemplatesEnabled=true
```

If you have this property set to *true*, you receive a CustomAppServer object representing the application running on the target machine. You do not need to provide application credentials to enable this property.

Enterprise Configuration Discovery and Tracking server full synchronization

A full synchronization, by default, only deletes configuration items for the domain before synchronizing them up again.

It does not delete all of the Enterprise Configuration Discovery and Tracking server local relationships with the configuration items to the domain. If you need to delete all references to the domain configuration items, add an empty file called dist/sync/deleteImplicitRelationsOnFullSync to the Enterprise Configuration Discovery and Tracking server before you do a full synchronization with the domain. This ensures that all implicit relationships to the domain configuration items are deleted. They are not recreated after the synchronization and need to be recreated in the Enterprise Configuration Discovery and Tracking server.

Configuring the View Manager Cache type

The View Manager supports two types of caching, in-memory and disk caching. When you enable in-memory caching, views are cached inside the View Manager process. You must use this for small data sets where the time taken to rebuild the cache after a server is restarted is not a major consideration.

When you enable disk caching, views are cached to disk. You must use this for large data sets. Because this is persistent storage, it remains accessible through the Configuration Discovery and Tracking server when you restart the server.

To optimize performance, complete the following steps to enable disk caching:

1. Shut down the CDT server.
2. Edit the *$COLLATION_HOME*/etc/collation.properties file.
3. Set the com.collation.view.cache.disk property to *true*.
4. Restart the server.

If you want to leave the View Manager Disk cache enabled, and you want to change the location of the cache to a different directory other than the default directory, complete the following steps:

1. Shut down the CDT server.
2. Edit the *$COLLATION_HOME*/etc/collation.properties file.
3. Set the com.collation.view.cache.disk.path.property to the preferred path. The path you set has to be a relative path from the $COLLATION_HOME directory.

The default path is set to the var/viewmgr directory, which means that the disk cache will be located at the *$COLLATION_HOME*/var/viewmgr directory.

4. Remove the existing disk cache, if there is any.
5. Restart the server.
6. Verify that the disk cache gets created at your new location.

Configuring the Configuration Discovery and Tracking server to prebuild certain views into the View Manager Disk cache

To enable the Configuration Discovery and Tracking server to prebuild certain views into the cache on a CDT server restart and after a discovery, complete the following steps:

1. Go to the *$COLLATION_HOME*/etc/collation.properties file.
2. Set one of the following properties to *true* (default value is *true*):
3.

com.collation.view.prebuiltcache.treeview.physical.enabled=true
Navigation tree for physical infrastructure

com.collation.view.prebuildcache.treeview.components.enabled=true
Navigation tree for software components tree

com.collation.view.prebuildcache.treeview.application.enabled=true
Navigation tree for applications tree

com.collation.view.prebuildcache.treeview.bizservice.enabled=true
Navigation tree for business services tree

com.collation.view.prebuildcache.graph.appinfrastructure.enabled=true
Application Software Infrastructure Topology

com.collation.view.prebuildcache.graph.physicalinfrastructure.enabled=true
Physical Infrastructure Topology

The following properties define trees that are used to generate screens in various wizards, such as to define an application, business service, or dependencies. They must be enabled when the performance of the above wizards is important.

com.collation.view.prebuildcache.treeview.define.application.enabled=true

com.collation.view.prebuildcache.treeview.define.bizservice.enabled=true

com.collation.view.prebuildcache.treeview.define.appserver.enabled=true

com.collation.view.prebuildcache.treeview.define.appserverclusters.enabled=true

com.collation.view.prebuildcache.treeview.define.appserverhosts.enabled=true

com.collation.view.prebuildcache.treeview.define.appserverclusterserviceshosts.enabled=true

Configuring the Configuration Discovery and Tracking server for the View Manager cache to use access control for views

If the user wants the views to be ACL aware, for example, if you want a user that has logged in to the console to only view objects that he has access to view, complete the following steps:

1. Go to the *$COLLATION_HOME*/etc/collation.properties file.

2. Set the com.collation.view.accesscontrol.enabled=true. The default for this property is *false*.

When access control for views is enabled, no view caching is available. For large views, performance can be decreased.

Dropping a database and clearing the View Manager Disk cache

If you are dropping your database and you have the View Manager Disk cache enabled, you have to manually delete the old disk caching directories.

To do this, after you drop the database, go to the directory where you have configured the com.collation.view.cache.disk.path property and delete the contents of the following sub-directories:

- /graph
- /tree

Note: The default directory is dist/var/viewmgr.

Using the WebSphere seed sensor

The WebSphere sensor is usually initiated by an Operating System sensor. To discover WebSphere resources on platforms that do not have an Operating System sensor, for example, z/OS®, the WebSphere sensor has been enhanced to support discovery initiated from seed files.

Before running the WebSphere seed sensor

There are some prerequisite steps you need to complete before running the WebSphere seed sensor:

1. If you are not loading an IDML book, skip this step. If you are loading an IDML book before running the WebSphere discovery, IDML must populate WebSphereCell, WebSphereNode, and WebSphereServer instances using the AppServerKeyName naming rule. This naming rule uses the primarySAP implicit attribute that is derived from an accessedVia relationship to the BindAddress.

   If the attributes are not included in the IDML book. the merge of data does not occur between DLA discovered data and discovery sensor data.
2. Create the WebSphere seed file. If you are going to discover WebSphere on a z/OS system, use the utility provided on the IBM Tivoli Open Process Automation Library Web site to generate the seed files automatically from the IDML books created from the z/OS DLA. If you are going to discover WebSphere on a non-mainframe system, manually create the seed file. Use the following file naming conventions when creating the seed file:
   a. If you want the file to be included as part of the discovery, the file name must end with a .xml extension.
   b. The file name must adhere to the following format:

      ```
      <cellname>_<fqdn>_<port>.xml
      i.e. c1_0.0.0.0_2809.xml
      ```

   The following is an example of the file format:

   ```
   <IDML_WAS_SEED>
     <WAS_ROOT_DIR>/opt/WebSphere/AppServer</WAS_ROOT_DIR>
     <WAS_VERSION>6.0.2.7</WAS_VERSION>
   ```

```
  <SOAP_CONNECTOR_PORT>8880</SOAP_CONNECTOR_PORT>
  <RMI_CONNECTOR_PORT>2809</RMI_CONNECTOR_PORT>
  <JMX_LISTEN_IP_ADDRESS>0.0.0.0</JMX_LISTEN_IP_ADDRESS>
  <HOST_MAPPINGS>
    <HOST_MAPPING>
      <HOST_NAME>wasserver.compnay.com</HOST_NAME>
      <PRIMARY_IP_ADDRESS>0.0.0.0</PRIMARY_IP_ADDRESS>
      <IP_ADDRESS>0.0.0.0</IP_ADDRESS>
    </HOST_MAPPING>
  </HOST_MAPPINGS>
</IDML_WAS_SEED>
```

WAS_ROOT_DIR
Directory path of where the WebSphere Application Server is installed.

WAS_VERSION
The version of WebSphere Application Server. Can be found in the WebSphere Application Server product file; in the *<WebSphere Root Directory>*/properties/version directory.

SOAP_CONNECTOR_PORT
The port number is retrieved from the serverindex.xml file for the *SOAP_CONNECTOR_ADDRESS* endpoint name. For example, *<WebSphere Root Directory>*/profiles/*<app server or dmgr>*/conf/cells/*<cell name>*/nodes/*<node name>*.

If the resource is a deployment manager, use the serverindex.xml file where the following is specified: `serverType="DEPLOYMENT_MANAGER"`.

If the resource is standalone, use the serverindex.xml file where the following is specified: `serverType="APPLICATION_SERVER"`

RMI_CONNECTOR_PORT
The port number is retrieved from the same serverindex.xml file used to find the soap port, where the endpoint name is *BOOTSTRAP_ADDRESS*.

JMX_LISTEN_IP_ADDRESS
The IP address that is used to connect through JMX. Usually this is the same IP address as the WebSphere server.

HOST_MAPPINGS
A list of mappings between hostname and IP address for the WebSphere Application Server or Deployment Manager and each distributed Node Agent.

HOST_MAPPING
One host mapping consisting of a host name, primary IP address, and IP address.

HOST_NAME
The fully qualified domain name.

PRIMARY_IP_ADDRESS
The primary IP address that the host name resolves to.

IP_ADDRESS
The IP address that the host name resolves to, if different from the primary IP address.

3. Place the .xml files in the *$COLLATION_HOME*/var/dla/zos/was directory. If the directory does not exist, create the directory. The scope of discovery is

controlled by the files in this directory. If discovery of a particular WebSphere server is no longer desired, the file must be removed from this directory or renamed without the .xml extension.

4. Edit the sensor configuration file, IdmlFileUDS.xml, located in the *$COLLATION_HOME*/etc/discover-sensors directory, or skip this step and create a new sensor configuration file when you run the WebSphere seed sensor. If you edit the sensor configuration file, use the following suggestions:

   <filename>
   The directory where the WebSphere XML seed files are located.

   <scope>
   The IP address of the Configuration Discovery and Tracking server where the WebSphere XML seed files reside.

5. If you plan to load an IDML book prior to WebSphere discovery, add the following property to the collation.properties file (in the *$COLLATION_HOME*/etc/ directory): `com.collation.bulkload.data.priority=10`

Running the WebSphere seed sensor

To run the WebSphere seed sensor, complete the following steps:

1. Start the Configuration Discovery and Tracking server.
2. Open the Product Console.
3. Add the IP address of the Configuration Discovery and Tracking server where the WebSphere seed file resides to the scope.
4. Add the access credentials of the Configuration Discovery and Tracking server into the Access List, and, if security is enabled for the WebSphere server being discovered, add the credential entry for the WebSphere server.
5. Complete the following steps to create a new configuration file for the IdmlFileUDS sensor:
   a. In the **Discovery Profiles** window, click **IdmlFileUDS**.
   b. Click **New**.
   c. Type the sensor configuration name and description.
   d. Select the **Enable this configuration and disable selected configuration** checkbox.
   e. Double-click */data/latest/dist/var/dla/zos/was* and type the location of the WebSphere XML seed files.
   f. Double-click *0.0.0.0* and type the IP address of the Configuration Discovery and Tracking server.
   g. Save the configuration file.
6. Create a discovery profile that includes the following sensors:
   - AnchorSensor
   - IdmlFileUDS (If you created a sensor configuration using the instructions in the previous step, select the configuration you created.)
   - PortSensor
   - PingSensor
   - SessionSensor
   - WebSphereIdmlSeedSensor
   - WebSphereCellSensor
   - GenericServerSensor
   - WebSphereNodeSensor

- WebSphereSensor

7. Run the discovery and select the scope to include the Configuration Discovery and Tracking server as well as the discovery profile you created.

Troubleshooting the WebSphere sensor

The following list provides information about errors that can be reported and ways to solve the problems.

- Error: If security is enabled on the WebSphere Application Server, there could be incorrect credentials or password. This type of error message results if there are no WebSphere credentials in the ACL:

  ```
  ERROR cdb.WebSphereAgentDelegate - [WebSphereAgentDelegate.E.1] discover()
  failed with exception : java.lang.Exception: Unable to connect to the
  WebSphere server at 9.48.158.37:8,880 - ADMC0016E: The system cannot create
  a SOAP connector to connect to host 9.48.158.37 at port 8880..
  ```

  Solution: Add WebSphere credentials in the ACL.

- Error: The following error is displayed if there is a credential entry for the WebSphere server, but the certificates are not correct, or have not been entered through the ACL, or the password is incorrect.

  ```
  ERROR cdb.WebSphereJMXUtils - An error occurred, unable to establish a
  repository connection using the credentials raleigh-was60:
  com.ibm.websphere.management.exception.AdminException:
  javax.management.JMRuntimeException: ADMN0022E: Access is denied for
  the getServerConfig operation on FileTransferServer MBean because of
  insufficient or empty credentials.
  ```

  Solution: Correct the certificates, enter the certificates through the ACL, or specify the correct passphrase.

- Error: The WebSphere sensor does not start for one of the following reasons:
  - The process is not running.
  - The process running does not match the template for WebSphere, possibly because the command line was truncated. For Linux, Solaris, AIX, and Linux on System z operating systems, the command line must contain the word *WsServer*. For Windows operating systems, the command line must contain *IBM WebSphere Application Server*.
  - No PID was found for the WebSphere Application Server. WebSphere Application Server could not have started as a service on Windows.

  Solution: There are multiple reasons why the WebSphere Application Server does not start or is corrupted. Check the system log and WebSphere server start logs for error messages. For Windows operating systems, for some versions of WebSphere, the WebSphere server must be started as a service instead of through the command line interface.

- Data store exceptions, including the following:
  1. The schema was not updated accurately for the version of Configuration Discovery and Tracking installed. The solution is to drop and recreate the database.
  2. The data in the database is corrupt. The solution is to clean out the data by using topology clear.
  3. The database tuning script has not been run prior to Configuration Discovery and Tracking schema creation.
  4. The data persisted from the sensor is too large and is persisting duplicate objects. This error can be determined by looking at what model object the topology manager had trouble persisting. At times, the problem is that too

many objects are trying to persist to the same table for the database. To see this error, the log level needs to be set to *DEBUG*.

5. Long strings try to persist into the database. The TopologyManager logs, the error message is printed. To see this error, the log level needs to be set to *DEBUG*. The solution is to identify which table and column contains the problem, update the schema, and recreate the database.

- Error: The server is down for one of the following reasons:
  - Configuration Discovery and Tracking runs when WebSphere servers are in maintenance and a discovery does not complete. The local anchor log can display the following error message:

    ```
    INFO cdb.AnchorServer[main] - [AnchorServer.I.0] server no longer
    accepting new connections
    ```

  - Another error message can report that the query cannot be completed.

  Solution: Check WebSphere server health.
- Error: An incorrect jar version is installed. The following error message can be displayed:

  ```
  class not found exception: javax/management/AttributeNotFoundException
  ```

  Solution: Install the correct jar version.
- The **com.collation.websphere.deep.discovery.setting** property must be set to *true* if Web services need to be discovered. This property only works for WebSphere Application Server, version 6.x. If Web services are required to be discovered, the WebSphere, version 6.0 or WebSphere, version 6.1 Thin Client jars must be used and the property must be set to *true*.
- The **com.collation.websphere.performance.setting** property should not be set to *true* because of performance issues. By default, this property is set to *false*.
- Error: The server cannot connect because it cannot read the serverindex.xml file. The following error message is displayed:

  ```
  ERROR j2ee.WebSphereAgent - [WebSphereAgent.E.17] getConnectorPorts()
   - unable to read server index file
  C:\Program Files\WebSphere\AppServer\config\cells\slapc14\nodes\slapc14\
  serverindex.xml, terminating discovery
  ```

  Solution: This error can occur because of permission issues with the file. It is possible that the user given as part of the computer system credentials, for the WebSphere server, does not have access to the file system.
- JDBC dependency connections are only made with those resources found under the Application Server level, not at the cell and node levels.

Downloading and installing the WebSphere Application Server seed utility for the z/OS DLA

If you want to discovery WebSphere on z/OS, you can use the WebSphere Application Server seed utility for the z/OS DLA to create the XML seed files from the IDML book. If you do not use the utility, you have to manually create the XML seed files.

You can download the WebSphere Application Server seed utility for the z/OS DLA from the IBM Tivoli Open Process Automation Library Web site.

To download the utility, complete the following steps:

1. Go to the IBM Tivoli CCMDB section of the IBM Tivoli Open Process Automation Library Web site: http://catalog.lotus.com/wps/portal/tccmd.
2. Enter the following search criteria: *TADDM WAS seed utility for the z/OS DLA*.

3. Click **TADDM WAS Seed Utility for the z/OS DLA**.
4. Follow the instructions on the Web site to download the utility.
5. After downloading the utility, read the documentation for installation instructions.

Chapter 9. Functionality provided with Fix Pack 002 for the Configuration Discovery and Tracking feature

This section identifies and explains the new functionality that is available for the Configuration Discovery and Tracking feature in Fix Pack 002.

Configuring Configuration Discovery and Tracking

You do not have to complete any specific configuration tasks before using Configuration Discovery and Tracking. This section provides instructions for optional configuration tasks that can be useful in your environment.

Saving Domain Manager reports

If you want to save Configuration Discovery and Tracking server and Enterprise Configuration Discovery and Tracking server Domain Manager reports in the PDF, CSV, or XML formats, update the **com.collation.domain.pdfreport.enabled** parameter.

To update the **com.collation.domain.pdfreport.enabled** parameter, complete the following steps:

1. Open the collation.properties file (/opt/IBM/cmdb/dist/etc/collation.properties) in a text editor.
2. Change the **com.collation.domain.pdfreport.enabled** parameter from *false* to *true*.
3. Save and close the collation.properties file.
4. Restart Configuration Discovery and Tracking.

Bidirectional support is not provided for saving Domain Manager reports in the PDF, CSV, or XML formats. In addition, the instructions provided in this section have been tested with the English locale; the instructions may not work with non-English locales.

Using sensors

The information in this section does not apply to the IDD StackScan sensor. The StackScan sensor uses credential-less discovery and is described in the next section.

Sensors work by emulating a user running locally to gather (discover) information. Like an agent, the sensor is local to the monitored host, but only for a brief period of time. Because the sensor is local to the monitored host for a limited amount of time, the sensor is able to gather discovery-related information without incurring the typical agent costs of installation and maintenance.

The sensor uses secure network connections, encrypted access credentials, and host-native utilities. In this way, a sensor is safe to use and provides the same data acquisition that a user (or an agent) could acquire by residing locally on a resource. The sensor asks, figuratively, the host and the applications how they are configured and who they are talking to. The sensor differs from the agent as the agent takes the role of an outsider and tries to infer data from what is observed. With a sensor, there is no need to install anything permanently on each host that requires ongoing maintenance.

The sensor functions in the following way:

- A secure connection is made to the host about how to be discovered using a secure protocol, for example, SSH.
- SSH uses credentials that enable read-only access to system information (without root access), but does not necessitate access to application data or customer information.
- Once the SSH session is created, the sensor invokes native operating system utilities that return all the running processes on this host, their environments, and arguments. The sensor also captures all of the open ports and the processes that are communicating on the ports.
- Using these native utilities, the sensor is able to create an inventory of the host. The inventory includes the following details about the host: configuration, software and patches installed, applications it is hosting, and running processes.
- Differently configured sensors, using appropriate credentials, can log onto whatever application are found on this host, for example, Oracle or WebSphere. This application-specific sensor, after logging into the application, can ask the application how it is configured. This information is important for gathering an accurate configuration. For example, Oracle can use the concept of a running configuration (using `spfile`) that differs significantly from the configuration described in a configuration file.
- Once the host and application topology and configuration details are acquired, they are sent back to a central management server and the SSH session is ended.
- After the session is ended, the connection is stopped with nothing left on the discovered host to be maintained.

Because this same sensor-driven process is simultaneously occurring on other hosts, the process at the other end of each port is captured. Using correlation back at the central server, a topology can be created that accurately portrays the dependencies between applications.

The data captured by the sensor is reformatted according to schema in the CIM model (Common Data Model) and stored in a configuration management database.

Using the IDD StackScan sensor

The IDD StackScan sensor provides credential-less discovery using stack classification for a less intrusive mapping of the installed operating system and open ports on a computer system. The StackScan sensor can collect the type of operating system, the active IP interface, and the open ports.

Use discovery profiles to enable (and disable) the StackScan sensor. Discovery profiles are part of Fix Pack 002. Information about discovery profiles is provided in the Using discovery profiles section.

The StackScan sensor labels each discovered computer system with a confidence level for the operating system. If the operating system confidence level is higher than a threshold, the computer system is displayed under the appropriate operating system category. If the operating system confidence level is below the threshold, the operating system is modeled as a general computer system. The threshold is configured between zero and 100. The sensor configuration attribute is **confidenceThreshold**.

To enable and disable the StackScan sensor, and to set the **confidenceThreshold** attribute, use a discovery profile.

For supported Windows operating systems, the StackScan sensor needs raw socket support enabled on the operating system (for example, the Windows 2003 server) where the Configuration Discovery and Tracking server is running. If the operating system does not provide raw socket support, the StackScan sensor cannot work; data collection does not occur.

Configuring sudo access control

The StackScan sensor requires sudo access control to collect discovery information. For Windows operating systems, sudo access control is not needed.

To configure sudo access, complete the following steps for the Configuration Discovery and Tracking server and anchor hosts:

1. From a command prompt window, use the su command to switch to root authority on the local host.
2. Type the visudo command.
3. Type the following line in the /usr/local/etc/sudoers or /etc/sudoers file:

   ```
   <TADDM_USER>ALL=(ALL) NOPASSWD:ALL
   ```

   <TADDM_USER> is the non-root user ID that is used by the Configuration Discovery and Tracking server.

Installing Nmap

If you use the StackScan sensor, the use of Nmap is optional. If you want to increase the accuracy of your results, use Nmap.

Nmap is an open source network exploration tool and security scanner. When using the StackScan sensor, you should install Nmap as a part of the Configuration Discovery and Tracking installation process.

Nmap must be installed on the same operating system where your Configuration Discovery and Tracking server is installed. If you use an anchor server, Nmap needs to be installed on the anchor server too.

For the AIX and Windows operating systems, you need to download the latest version of Nmap from the following Web site: www.insecure.org

The supported Linux operating systems includes a version of Nmap. This can be an older version. Download the newest version from www.insecure.org.

For Solaris, you need to download and install a corresponding version of openssl before you install Nmap. Complete the following steps to install the correct openssl and Nmap software:

1. Go to www.sunfreeware.com.
2. Install the corresponding version of openssl and the appropriate Nmap software:
   - For Nmap 4.20, you need to have the libssl.so.0.9.8 library installed on the Configuration Discovery and Tracking server first.
   - For Nmap 3.93, you need to have the libssl.so.0.9.7 library installed on the Configuration Discovery and Tracking server first.
3. Use the following command to validate the command before proceeding with any tasks:

   ```
   nmap -v -sS -O
   ```

   The output should not contain errors.

Use the instructions to install Nmap from a source RPM instead of compiling Nmap from the source code.

If you do decide to install Nmap by compiling the source code, specify /usr/bin as the default directory. For example, to specify /usr/bin as the default directory, you could use the following option for the command:

```
./configure –prefix=/usr/bin
```

Note: Before installing Nmap for any operating system, you may want to go to the Configuration Discovery and Tracking support Web site for any late, breaking news about your specific operating system and Nmap versions. http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliApplicationDependencyDiscoveryManager.html

Sensors

The following table lists the sensors that are delivered with Configuration Discovery and Tracking:

Table 4. Sensors provided by Configuration Discovery and Tracking

Sensor	Gather configuration information from
ActiveDirectory	Windows Active Directory
AixComputerSystem	AIX computer systems
AlteonPort	Alteon network equipment (port information)
AlteonSnmp	Alteon network equipment (SNMP information)
AlteonVlan	Alteon network equipment (vlaninformation)
Anchor	Anchor hosts
ApacheServer	Apache Web Servers, for example, IBM HTTP Server
BigIPPort	BigIP network equipment (port information)
BigIPSnmp	BigIP network equipment (SNMP information)
BigIPVlan	BigIP network information (vlaninformation)
BridgeSnmp	Generic network bridge - basic SNMP information
BridgeSnmp2	Generic network bridge - detailed SNMP information
Cdp	Cisco Discovery Protocol sensor
Checkpoint	CheckPoint firewalls
CheckpointSnmp	CheckPoint Firewalls (SNMP information)
CiscoPort	Cisco network equipment (port information)
CiscoTelnet	Cisco network equipment (using the TelNet protocol)
CiscoVlan	Cisco network equipment (vlan information)
CiscoWorksFile	CiscoWorks File sensor
CiscoWorksFileUDS	CiscoWorks File Universal Data Service sensor
CiscoWorksUDS	CiscoWorks Universal Data Service sensor
Citrix	Citrix Presentation Server for Windows
CustomAppServer	Servers identified by a Custom Server template defined in the Configuration Discovery and Tracking database
CustomComputerSystem	Computers identified by a Computer System template defined in the Configuration Discovery and Tracking database

Table 4. Sensors provided by Configuration Discovery and Tracking (continued)

Sensor	Gather configuration information from
Db2	IBM DB2 UDB
Dns	DNS Server or Service running on the client computer
EntityMIB	Basic SNMP information from any device
ExtremeVlan	Extreme network equipment (VLAN information)
F5 3-DNS load balancer	The 3-DNS Controller from F5 Networks; discovers load balancing configuration information, for example, IP addresses that are configured
GenericComputerSystem	Any computer in the infrastructure
GenericServer	Any server in the infrastructure (defined by having an active listening port)
HostResources	Hardware resources in a computer
HpUxComputerSystem	Computers running an HPUX operating system
IDD StackScan	Credential-less discovery using stack classification for a less intrusive mapping of the installed operating system and open ports on a computer system.
IIsWebService	Microsoft® Internet Information Server
IpDevice	Any IP device in the infrastructure
IpInterface	Any IP interface defined in the infrastructure
IPlanetServer	SUN IPlanet Java Servers
IpRange	Any IP device within the scope
IPSOComputerSystem	SOC IP discovery
JBoss	JBoss Application Servers
LanManagerSnmp	Lan Manager SNMP discovery
LinuxComputerSystem	Any computer running a Linux distribution
LotusDominoDetail	Lotus® Domino® Server (detailed information)
LotusDominoDomain	Lotus Domino Server (domain information)
LotusDominoInitial	Lotus Domino Server (basic information)
MQServer	IBM WebSphere MQ
NetscreenSnmp	Netscreen firewall SNMP sensor
NFSServer	Network File Server servers
NokiaSnmp	Nokia network equipment (SNMP information)
OpenVmsComputerSystem	Computers running an OpenVMS operating system
Oracle	Oracle database servers
OracleApp	Oracle Application Servers
Ping	IP Interface (status information)
Pix	Cisco PIX Firewall sensor
PortScan	Open ports on any IP device
Session	Listening-port session information
SMBServer	Server Message Block File servers, for example, Windows file servers and Linux Samba servers
SMIS	Storage Management Initiative Specification
SMS	Microsoft System Management Server

Table 4. Sensors provided by Configuration Discovery and Tracking (continued)

Sensor	Gather configuration information from
SnmpMib2	Detailed SNMP information
SqlServer	Generic relational database servers
StackScan	Credential-less discovery using stack classification
Storage	Computers (gathers information related to the storage subsystem)
SunSparcComputerSystem	Sun SPARC Computer Systems
Sybase	Sybase database servers
SybaseIQ	Sybase databases
System p	Current IBM RISC/UNIX-based server and workstation product line
VMware	VMware ESX Server using a Linux kernel that loads additional code
WebLogic	BEA WebLogic Application Servers
WebSphere	IBM WebSphere Application Servers
WebSphereCell	IBM WebSphere Application Servers (cell information)
WindowsComputerSystem	Computers running a Windows operating system

Operating system sensors

The following table lists the operating system sensors used by the Configuration Discovery and Tracking discovery process:

Table 5. Operating system sensors

Operating system sensors	Versions supported
Sun Solaris	2.8, 2.9, 2.10
Linux	Red Hat Linux, 3.x and 4.x SUSE Linux, 9.x, 10.x
AIX	5.x
HP-UX	11.0, 11i V.1 (B.11.11)
Windows	2000, with Service Pack 4 2003, with Service Pack 1
OpenVMS	7.x
* System p	HMC 5.2.x, IVM 1.2.ix
** VMware ESX	2.5x, 3.0

* In order to discover the System p system (and its logical partitions), you must provide the IP address of the management console that is managing the System p system. The SystemP sensor supports two management consoles. Using SSH, the SystemP sensor establishes a session with the management console. The session is established for a user. This user must be created in the following respective management console:

- Hardware Management Console (HMC)
- Integrated Virtualization Manager (IVM).

– For an HMC management console, a user based on the **hmcoperator** role is needed.
– For an IVM management console, a user with the **View Only** role is needed.

For example, create a new role called *taddmViewOnly* based on **hmcoperator**. In addition, the following command line tasks must be assigned to the new role:

Managed System
lshwres and lssyscfg

Logical Partition
lshwres, lssyscfg, and viosvrcmd

HMC Configuration
lshmc

** For VMware ESX, in order to get more configuration information about the guest operating system, the *getguestinfo ip* value must be set to the IP address for the guest operating system, or the IP address is included in the scope of the resources to discover. Otherwise, the host name for the guest operating system is not known and is displayed as *unknown* in the GUI and XML output.

Web server sensors

The following table lists the Web server sensors used by the Configuration Discovery and Tracking discovery process:

Table 6. Application sensors for Web servers

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
Apache 1.x	X	X	X	X	X	X	
Apache 2.x	X	X	X	X	X	X	
iPlanet 4.x	X	X	X	X	X		
iPlanet / Sun One (iPlanet) 6.x	X	X	X	X	X		
IIS 5.x						X	
IIS 6.x						X	
IBM HTTP Web server 6.x	X	X	X	X	X	X	

Application server sensors

The following table lists the application servers sensors used by the Configuration Discovery and Tracking discovery process:

Table 7. Sensors for application servers

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
WebSphere 5.1x		X	X	X	X	X	
WebSphere 6.x		X	X	X	X	X	

Table 7. Sensors for application servers (continued)

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
WebLogic 7.x	X	X	X	X	X		
WebLogic 8.x	X	X	X	X	X	X	
JBoss 4.x			X	X	X	X	
Lotus Domino 6.0x		X				X	
Lotus Domino 6.5x		X				X	
Oracle App 10.1.3			X	X	X	X	
Citrix 3.x, (Enterprise edition only)						X	
Citrix 4.x (Enterprise edition only)						X	

Database sensors

The following table lists the databases sensors used by the Configuration Discovery and Tracking discovery process:

Table 8. Application sensors for databases

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
Oracle 8i	X	X	X	X	X	X	
Oracle 9i	X	X	X	X	X	X	X
Oracle 10g/i	X	X	X	X	X	X	X
DB2 7.x		X	X	X	X	X	
DB2 8.x		X	X	X	X	X	
MS SQL 2000						X	
Sybase ASE 12.x					X		
Sybase IQ 12.x					X		

Packaged application sensors

The following table lists the packaged application sensors used by the Configuration Discovery and Tracking discovery process:

Table 9. Application sensors for packaged applications

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
SAP (SLD) 6.40	X	X	X	X	X	X	
SAP (SLD) 7.x	X	X	X	X	X	X	
SAP (CCMS) 4.6C	X	X	X	X	X	X	
SAP (CCMS) 4.6D	X	X	X	X	X	X	
SAP (CCMS) 6.x	X	X	X	X	X	X	

Messaging sensors

The following table lists the messaging sensors used by the Configuration Discovery and Tracking discovery process:

Table 10. Application sensors for messaging applications

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
WebSphere MQ 5.3	X	X*	X	X	X**	X***	
WebSphere MQ 6.0	X	X*	X	X	X**	X***	

* Only AIX 5.3.

** Only Solaris 2.8 and 2.9.

*** In addition, Windows XP Professional.

Services sensors

The following table lists the services sensors used by the Configuration Discovery and Tracking discovery process:

Table 11. Sensors for Services applications

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
MS SMS 2003						X	
MS Active Directory, 2000, 2003						X	
Sun ONE Directory Server, 5.x					X	X	
WFS (Samba), 3.x						X	

Firewall sensors

The following table lists the firewall sensors and the versions of the operating systems that they support:

Table 12. Sensors for Firewalls

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
Firewall (CiscoPIX 6.x)	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Firewall (NetScreen - All versions)	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Firewall (Check Point FireWall-1)	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Storage devices

The following table lists the host bus adaptor sensors used by the Configuration Discovery and Tracking discovery process:

Table 13. Storage devices

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
Storage devices attached via HBA (Emulex)					X		
Storage devices attached via Qlogic			X	X			

Network sensors

The following table lists the network sensors used by the Configuration Discovery and Tracking discovery process:

Table 14. Sensors for Networks

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
CiscoWorks 2000		X					
CiscoWorks 4.0.5					X	X	
Load balancer (Alteon 3.x)							
Router (Cisco Systems IOS 10.3+)							

Table 14. Sensors for Networks (continued)

Sensor and version	HPUX 11.0, 11i	AIX 5.x	Red Hat AS 3.0, 4.0	SuSE 9, 10	Solaris 2.8-2.10	Win 2K/2003	OpenVMS
Switch (Cisco Systems IOS 10.3+, CatOS)							
Switch (Extreme - Summit® 48/48i)							
Switch (HP ProCurve - SNMP supported)							

Introducing the z/OS Discovery Library Adapter

The z/OS Discovery Library Adapter (DLA) discovers mainframe resources and relationships, including the following items:

- z/OS configuration details (for example, PARMLIB members)
- zSeries® machine details (for example, serial number, processing capacity, and LPAR information)
- Address space details (for example, JobName, PGM, allocations, and related storage volumes)
- DB2 subsystems (for example, DataSharingGroup, databasees, and tablespaces)
- IMS™ subsystems (for example, transactions, programs, and databases)
- MQ subsystems (for example, sender channels and receiver channels)
- CICS® regions (for example, system initialization table settings, transactions, programs, and files)
- WebSphere Application Servers (for example, node, cell, and configuration files)

The z/OS DLA stores this resource and relationship information in an XML file. After discovery, the z/OS DLA sends the XML files, using FTP, to the Discovery Library File Store. You can use the bulk load program to load the XML files into IBM Tivoli CCMDB.

Downloading the z/OS DLA

You can download the z/OS DLA from the IBM Tivoli Open Process Automation Library Web site.

To download the z/OS DLA, complete the following steps:

1. Go to the IBM Tivoli CCMDB section of the IBM Tivoli Open Process Automation Library Web site: http://catalog.lotus.com/wps/portal/tccmd.
2. Enter the following search criteria: *Discovery Library Adapter for z/OS*.
3. Click **Discovery Library Adapter for z/OS**.
4. Follow the instructions on the Web site to download the z/OS DLA.
5. After downloading the z/OS DLA, read the z/OS DLA documentation for installation instructions.

Importing the z/OS Discovery Library books to Configuration Discovery and Tracking

The output file for the z/OS DLA is, by default, in the Discovery Library file format. For example, ZOSDISC100.pthomo1.au.ibm.com.2006-01-22T23.35.06Z.xml.

You need to transfer the XML file to the Discovery Library File Store and load the file in the Configuration Discovery and Tracking feature with the Discovery Library bulk load program. See the ″Discovery Library bulk loader″ section for more information.

Using information from the z/OS DLA with Configuration Discovery and Tracking

You can view all of the related information from the z/OS DLA with the Product Console.

To view information from the z/OS DLA with the Product Console, complete the following steps:

1. Go to the sidebar and find the Discovered Components section.
2. From the list, select **Physical Infrastructure**.
3. In the tree, select **Systems Tier** → **Sysplexes**. Find the z/OS resource you want to work with.
4. In the Details panel you can view information about and work with the z/OS resource.

Discovery Library bulk loader

The bulk load program (the loadidml.sh script) loads Discovery Library books into the Configuration Management Database.

The loadidml.sh script reads the books, imports the data into the Configuration Management Database, and logs the results in the results directory for the bulk load program. In addition, the bulk load program logs error messages in the *$COLLATION_HOME*/log/bulkloader.log file.

To ensure data consistency, only one bulk load program can run at a time. The bulk load program is designed to be run at the Configuration Discovery and Tracking server. To ensure proper authorizations, the bulk load program must be run by the same user ID that runs the Configuration Discovery and Tracking server processes.

All the directories that you use to store log and result files need to exist prior to running the loadidml.sh script. These directories may be customized through manipulation of the configuration settings defined in the *$COLLATION_HOME*/etc/bulkloader.properties file.

To configure and run the bulk load program, complete the following steps:

1. Check the *$COLLATION_HOME*/etc/bulkload.properties file for accuracy. You should not need to change anything in the file if you want to accept the defaults.

   com.ibm.cdb.bulk.workdir=bulk
   This is the directory the bulk load program uses to copy files before loading them. The default directory is relative to the top-level directory

of the directory referenced by the *$COLLATION_HOME* variable. This variable is usually `/opt/IBM/cmdb/dist`.

com.ibm.cdb.bulk.workdir.cleanup=false
Specifies whether the working directory should be cleaned up after the load completes.

com.ibm.cdb.bulk.processedfiles.cleanup=30
Number of days to keep files in the processed files list.

com.ibm.cdb.bulk.retrycount=5
Number of times to retry loading a file if a discovery is currently in progress.

com.ibm.cdb.bulk.retrydelay=600
Number of seconds in between retries while a discovery is in progress.

com.ibm.cdb.bulk.resultsdir=bulk/results
Directory to write the results files created during the load of a file. The default directory is relative to the top-level directory referenced by the *$COLLATION_HOME* variable. This variable is usually `/opt/IBM/cmdb/dist`.

com.ibm.cdb.bulk.stats.enabled=false
Whether statistics gathering of the bulk load program are performed. Turning on statistics decreases performance and increases log and result file sizes.

Do not change any other settings in the properties file not specifically mentioned here. Other settings may be ignored.

2. Verify that the working directory and the results directory mentioned in the bulkload.properties file are valid.

   The working and results directory must exist before running the bulk load program or the bulk load program does not run. The bulk load program does not automatically create these directories.

   The directories should be created using the same user that starts and stops the Configuration Discovery and Tracking server. If the bulk load program does not have permissions to read and write from the working and results directories, the bulk load program cannot run.

   If you want to use different directories, you have to create these directories manually and update the properties file.

3. Run the bulk load program. Log in as the user used to start and stop the Configuration Discovery and Tracking server. You can use the following commands:

```
su - taddm
export $COLLATION_HOME=/opt/IBM/cmdb/dist
cd dist/bin
./loadidml.sh -f <path_to_idml_file> -h <hostname> -u <userid> -p
<passwd> -o -b <bidirectional format on or auto>
```

   where

   -f <path_to_idml_file>
   This flag is required. The flag points to a fully-qualified path to the input file or a directory that contains input XML files. The directory where the input file is placed must not be the same as the working directory of the bulk load program. If a shared directory is used to stage input file, or, if files are copied to a local directory, the directory

where input files are staged cannot be the same as the working, results, or log directory of the bulk load program.

-h <hostname>
This flag specifies the hostname of the Configuration Discovery and Tracking server.

-u <userid>
This flag specifies the userid to be used to authenticate with the Configuration Discovery and Tracking server.

-p <password>
This flag specifies the password used to authenticate with the Configuration Discovery and Tracking server.

-o This flag instructs the bulk load program to override the processed files file and load the IdML files.

-b This flag specifies if bidirectional support is enabled, disabled, or automatically configured. Choices for the bidirectional flag are *on* and *auto*. When the bidirectional flag is *on*, you can configure the bidirectional parameters for each Management Software System using the predefined bidirectional profiles. When the bidirectional flag is *auto*, the bidirectional transformation is enabled and the bidirectional format is detected automatically.

If you are using SSH, you should not choose *on* for the bidirectional flag. When you choose *on* for the bidirectional flag and use SSH, the bulk load bidirectional configuration window is not displayed. Without completing the fields in the bulk load bidirectional configuration window, you cannot configure the bidirectional parameters.

The `-u` and `-p` parameters are optional. A userid and password should only be supplied if the userid has the correct permissions (full update and read privileges) and is defined in the Configuration Discovery and Tracking server as a valid user. The `-h` parameter should only be used in case you experience problems resolving the IP name of the Configuration Discovery and Tracking server.

If the bulk load program does not run, read the messages on the Product Console and the messages in the bulkload.properties file. The log file is located in the *$COLLATION_HOME*/log directory.

4. After the bulk load program runs, check the results file for problems during the bulk load program. The results file is located in the resultsdir directory configured in the bulkload.properties file. Use the following command to navigate to the directory with the bulkload.properties file:

   ```
   cd $COLLATION_HOME/bulk/results
   ```

   Look for a file with a .results extension and named the same as the IdML. If, for example, the name of the imported IdML file is `test.xml`, the name of the results file is test.results. Important entries in this file are marked with SUCCESS and FAILURE tags. Percentage successful messages are also recorded if statistics are enabled. FAILURE tags are for individual objects and do not necessarily indicate a failure of the entire file.
5. To process the same book again after the first initial load, either use the `-o` flag, or remove the specific entry from the processedfiles.list file. The processedfiles.list file is located in the working directory specified in the bulkload.properties file.
6. If the bulk load program indicates another bulk load program is running and you know this is not the case, go to the working directory and delete the

.bllock file and run the bulk load program again. The .bllock file is a hidden file because it starts with a period. Deleting this file should only be done if you are sure that another bulk load program is not already running.

You can also delete this file if you pressed <control-c> when running the loadidml.sh script.

You should also read the information in the bulkload.log file. The log file can contain details about messages that are displayed.

7. (Optional) To refresh the reports displayed in the Product Console, for example, the Inventory Summary report, with the data you added using the bulk load program, go to the dist/bin directory and run the following command:

   ```
   cdm.sh
   ```

   If you are using an Enterprise Configuration Discovery and Tracking server, go to the dist/bin directory and run the following command:

   ```
   ecmdb.sh
   ```

The following list describes the return codes for the bulk load program.

At a command line, you see the following return codes and their messages.

0 The program ran all the way through. This does not mean that everything was loaded. Check the results file for that information.

1 Some error occurred but it is unknown. Check the bulkload.log file in the log directory to see if there is more information.

2 A basic environment property needed to run the bulk load program is not set.

3 A command line parameter was supplied that is not valid. It is either the parameter itself or the data supplied with the parameter that is not correct. Correct the command and try again.

4 The user ID or password was not correct and the bulk load program could not log in. This happens when an incorrect -u -p parameter was supplied to the bulk load program.

5 The xml file being processed contained errors but the bulk load program continued to process the file.

6 The xml file being processed contained errors and caused the bulk load program to stop the processing of the file.

7 The xml parser failed to parse the xml file and the bulk load program processing stopped.

8 The API server returned an error but the bulk load program was able to recover and continue.

9 The API server returned an error and the bulk load program stopped processing.

10 Only one copy of the bulk load program can run at a time. A copy was already running so this copy can not run.

11 A discovery is processing and the bulk load program is locked out and can not run. Based on what is configured in the properties file, the bulk load program tries to run again, but if this error is returned, it has exhausted the retry attempts.

12 A discovery is processing and the bulk load program is locked out and can

not run. Based on what is configured in the properties file, the bulk load program tries to run again, but if this error is returned, it has exhausted the retry attempts.

13 There is an property specified in the input file for the bulk load program that is not valid.

14 The file was already processed as recorded in the processedfiles.list file in the working directory of the bulk load program. Either use the `-o` override parameter to force processing of the file or edit the processedfile.list and remove the entry for this file from the list.

15 The API server was not started and the bulk load program could not connect.

Topology updates for collections and business services

The Product Console contains a graphical user interface that is used to provide details on discovered components (such as collections and business services) and topologies (such as business applications, application infrastructure, and physical infrastructure).

This section describes enhancements that provide a more detailed view of collections and business services.

Exploring collections

You can display detailed information on the relationship topology and physical topology of a collection.

This section provides description of the new Collections menu items and instructions on how to display topology information for collections. The following information can be displayed:

- Relationship topology of a collection, which includes the following subcategories:
  - Expansion of child collections
  - Logical topology of collections
- Physical topology of a collection

The following tables lists the new Collections pop-up menu items in the Discovered Components section:

Table 15. Collections pop-up menu items in the Discovered Components section

Menu item	Description
Show Relationship Topology	Displayed detailed information about relationships between the components in the collection.
Show Physical Topology	Displays the topology of hardware used by the collection.

When you display the collection relationship topology only the parent collection is expanded. However, a collection might contain a child collection as one of its components. Therefore, if a parent collection has a child collection as one of its components, that child collection is not expanded automatically in the graph. You can use the graph that is displayed in the main window to view the following information:

- Relationship topology of the child collection

- Expansion of a child collection
- Logical relationships between components of the parent collection

When you display the physical topology of a collection, only components that cannot have their own physical topology are expanded automatically. Components such as business services, business applications, and collections can have their own physical topology and therefore are not expanded automatically. You can use the graph that is displayed in the main window to display the physical topology for these types of components.

The following tables lists the new Collections pop-up menu items in the main window:

Table 16. Collections pop-up menu items in the main window

Menu item	Description
Expand	Expands the child collection and shows explicit relationships between the components.
Show Logical Relationships	Displays the logical relationships between the components in the collection.
Show Relationship Topology	Displayed detailed information about relationships between the components in the collection.
Show Physical Topology	Displays the topology of hardware used by the collection.

Overview of the relationship topology of a collection

You can display a graph of the relationship topology of any collection.

After you display the graph of the topology, you can expand child components. You can also display logical relationships between components of the parent collection and child collections.

Viewing the relationship topology of a collection

To view the relationship topology of a collection, complete the following steps:

1. In the Product Console, click **Collections** in the **Discovered Components** list. A list of the members of each collection is displayed.
2. Right-click on either a collection or a component a collection for which you want to show topology. A pop-up menu is displayed.
3. From the pop-up menu, click **Show Relationship Topology**. A collection relationship topology graph is displayed in the main window.
4. If the collection displayed in the main window contains a child collection, right-click on the child collection and click **Show Relationship Topology**. A collection relationship topology graph for the child collection is displayed in the main window.

Introducing the expansion of a collection

When you expand the view of a collection, you can also see the explicit relationships between members of the collection.

The explicit relationships are relationships that are created between components of the collection (such as RunsOn and Federates) and they are shown in the graph by a solid black lines between the components.

Viewing the expansion of a collection

To expand the view of a child collection within a parent collection, complete the following steps:

1. In the Product Console, click **Collections** in the **Discovered Components** list at the bottom left side of the window. A list of the collections and their corresponding components is displayed.
2. Right-click on the collection that you want to expand. A pop-up menu is displayed.
3. From the pop-up menu, click **Show Relationship Topology**. A collection relationship topology graph is displayed in the main window.
4. Right-click on the child collection. A pop-up menu is displayed.
5. From the pop-up menu, click **Expand**. A collection relationship topology graph of the child collection is displayed in the main window. The graph also contains the explicit relationships between components of the collection, as shown by a solid black line between components.

Introducing logical relationships in a collection

You can use the relationship topology graph displayed in the main window to show logical relationship between two or more selected components in a collection.

If one of the selected components is a child collection that has not been expanded, that collection is be expanded when the logical relationships are displayed. In addition to the logical relationships between components, the explicit relationships between components are also shown.

Viewing the logical relationships in a collection

To view logical relationships in a collection, complete the following steps:

1. In the Product Console, click **Collections** in the **Discovered Components** list at the bottom left side of the window. A list of the members of each collection is displayed.
2. Right-click on the collection for which you want to display the logical relationships. A pop-up menu is displayed.
3. From the pop-up menu, click **Show Relationship Topology**. A collection relationship topology graph is displayed in the main window.
4. Select two or more components for which you want to view the logical relationships and right-click. A pop-up menu is displayed.
5. From the pop-up menu, click **Show Logical Relationships**. Dotted red lines are drawn between the components that have logical relationships.

Overview of the physical topology of a collection

The physical topology of a collection shows the topology of the hardware used by the collection.

When you display the physical topology of a collection, components within that collection that have their own physical topology (business application, business service, or collection) are not be expanded in the topology graph. For example, if your collection contains a business application, that business application is displayed as one component. You can expand that business application within the topology graph that is displayed in the main window.

Viewing the physical topology of a collection

To view physical topology of a collection, complete the following steps:

1. In the Product Console, click **Collections** in the **Discovered Components** list at the bottom left side of the window. A list of all the collections is displayed.
2. Right-click on the collection that you want to display the topology. A pop-up menu is displayed.
3. From the pop-up menu, click **Show Physical Topology**. A physical topology graph is displayed in the main window.
4. If the collection displayed physical topology graph contains a component that is a business application, business service, or child collection, right-click on the desired component and click **Show Physical Topology**. A physical topology graph for the selected component is displayed in the main window.

Exploring business services

You can display detailed information on the physical topology of a business service.

The Business Services pop-up menu contains one entry: Show Physical Topology. Click **Show Physical Topology** to display the topology of hardware that is used by the business server.

When you display the physical topology of a business service, only components that cannot have their own physical topology are expanded automatically. Components such as business services, business applications, and collections can have their own physical topology and therefore are not expanded automatically. You can use the graph that is displayed in the main window to display the physical topology for these types of components.

Overview of the physical topology of a business service

The physical topology of a collection shows the topology of the hardware used by the business service.

When you display the physical topology of a business service, components within that business service that have their own physical topology (business application, business service, or collection) are not be expanded in the topology graph. For example, if your business service contains a business application, that business application is shown as one component. From the topology graph, you can expand that business application individually.

Viewing the physical topology of a business service

To view physical topology of a business service, complete the following steps:

1. In the Product Console, click **Business Services** in the **Discovered Components** list. A business service overview is displayed.
2. Right-click on the business service for which you want to display the topology. A pop-up menu is displayed.
3. From the pop-up menu, click **Show Physical Topology**. A physical topology graph of the business service is displayed in the main window.
4. If the business service displayed in the physical topology graph contains a component that is a business application, business service, or child collection, right-click on the desired component and click **Show Physical Topology**. A physical topology graph for the selected component is displayed.

Using discovery profiles

Discovery profiles help you discover your IT environment.

Configuration Discovery and Tracking discovers and collects configuration information for the entire application infrastructure, identifying deployed software components, physical servers, network devices, virtual LAN, and host data used in a runtime environment. Using a discovery profile, you take control of what you discover.

For example, you configure individual sensors, manage multiple configurations of the same sensor, pick the appropriate configuration based on a set of criteria, and manage sets of configuration of different sensors to be applied on a single run.

When you run a discovery, you must select a profile. If no profile is selected, the discovery is run against the default profile, which is Level 3 discovery. The default profile can be changed. Click **Edit** → **Preferences** and select another profile.

Creating discovery profiles

To create discovery profiles, complete the following steps. When creating discovery profiles, default profiles, default sensors, and default sensor configurations are not editable.

1. In the **Discovery** drawer of the Product Console, click **Discovery Profiles**.
2. In the Discovery Profiles window, click **New**.
3. Type the profile name. The profile name must be unique.
4. Type a description for the new profile. The description is displayed on the user interface with the Sensor Configuration and Platform Properties pages.
5. When you create a new profile, you can use an existing profile as a basis. From the **Clone existing profile** list, select an existing profile or select *None*.

   There are three levels of discovery profiles to choose from:

   Level 1 Discovery
   This profile can be used to perform credential-less discovery. It can be used to discover active computer systems in the runtime environment.

   Level 2 Discovery
   This profile can be used to discover detailed information about the active computer systems in the runtime environment.

   You can use the Level 2 profile to enable shallow discovery of applications running on a target system using only the system credentials by adding the following property to the collation.properties file:

   ```
   com.collation.internalTemplatesEnabled=true
   ```

   If you have this property set to *true*, you receive a CustomAppServer object representing the application running on the target machine. You do not need to provide application credentials to enable this property.

   Level 3 Discovery
   This profile can be used to discover the entire application infrastructure, deployed software components, physical servers, network devices, virtual LAN, and host data used in a runtime environment.
6. Click **OK**. The discovery profile is created and listed with the other existing profiles. The profiles are listed beside the Sensor Configuration and Platform Properties pages. If you cannot see the profiles, look for a splitter bar beside Sensor Configuration page. Move the splitter bar to see the list of profiles. When you select a profile, the details for the profile are displayed on the Sensor Configuration and Platform Properties pages.

7. On the Sensor Configuration page, select a sensor and you can create, enable, and configure sensors. When you configure a sensor, double-click the value that you want to edit.

   You can add scope restrictions to a sensor. A scope restriction means that when a discovery is performed using a profile, the sensor runs only on the scope configured with this scope restriction. For example, if you want the *WebSphereSensor* for the *ProfileTest* profile to run on the *WebSphereDiscovery* scope set, create a new sensor configuration based on the WebSphere sensor and configure a scope restriction of *WebSphereDiscovery*. When you run the discovery using the *ProfileTest* profile, select the appropriate scope sets (including *WebSphereDiscovery*) and the WebSphere sensors runs only on the WebSphere discovery scope set.
8. On the Platform Properties page, you can add, edit, or delete properties for a platform.
9. Click **Save**.

Changing discovery profiles

To change discovery profiles, complete the following steps:

1. In the **Discovery** drawer of the Product Console, click **Discovery Profiles**.
2. In the Discovery Profiles window, select the profile you want to change. The profiles are listed beside the Sensor Configuration and Platform Properties pages. If you cannot see the profiles, look for a splitter bar beside Sensor Configuration page. Move the splitter bar to see the list of profiles. When you select a profile, the details for the profile are displayed on the Sensor Configuration and Platform Properties pages.
3. On the Sensor Configuration page, select a sensor and you can create, enable, configure, and delete sensors. When you configure a sensor, double-click the value that you want to edit. When you delete sensors, the software does not allow you to delete default sensors.
4. On the Platform Properties page, you can add, edit, or delete properties for a platform.
5. Click **Save**.

Deleting discovery profiles

To delete a discovery profile, complete the following steps:

1. In the **Discovery** drawer of the Product Console, click **Discovery Profiles**.
2. In the Discovery Profiles window, select the profile that you want to delete. You cannot delete a default profile. The profiles are listed beside the Sensor Configuration and Platform Properties pages. If you cannot see the profiles, look for a splitter bar beside Sensor Configuration page. Move the splitter bar to see the list of profiles.
3. Click **Delete**. A confirmation message is displayed.

Scheduling discovery profiles

To create a schedule for a discovery profile, complete the following steps:

1. In the **Discovery** drawer of the Product Console, click **Schedule**.
2. In the Schedule window, click **Add**.
3. On the Details page, complete the following steps:
   a. Type a name.
   b. Select a start date and time.

c. Select the option for repeating the discovery.

4. On the Scope page, complete the following steps:
   a. Select a scope.
   b. For the scope, set the options.
   c. Select a profile.
5. Click **OK**.

Running a discovery using profiles

To run a discovery using a profile, complete the following steps:

1. In the **Discovery** drawer of the Product Console, click **Overview**.
2. In the Overview window, click **Run Discovery**.
3. Select the scope elements and components.
4. Select a profile.
5. Click **OK**.

Appendix. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

If you are viewing this information softcopy, the photographs and color illustrations might not appear.

Trademarks

AIX, DB2, IBM, the IBM logo, IBMLink, Informix, OS/2, OS/390, OS/400, Tivoli, the Tivoli logo, Tivoli Enterprise Console, and TME are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe, Acrobat, Portable Document Format (PDF), and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel Inside (logos), MMX, Celeron, Intel Centrino, Intel Xeon, Itanium, Pentium and Pentium III Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Other company, product, and service names may be trademarks or service marks of others.

Index

R

S

T

V

W

Z

Printed in USA