HP OpenView Reporter

for the Microsoft Windows operating system

Software Version: A.03.70

DB Schema Document

Legal Notices

Warranty

Hewlett-Packard makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

A copy of the specific warranty terms applicable to your Hewlett-Packard product can be obtained from your local Sales and Service Office.

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause in DFARS 252.227-7013.

Hewlett-Packard Company
United States of America

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Copyright Notices

Trademark Notices

Java™ is a trademark of Sun Microsystems, Inc.

Microsoft® is a U.S. registered trademark of Microsoft Corporation.

Windows NT® is a U.S. registered trademark of Microsoft Corporation.

Windows® and MS Windows® are U.S. registered trademarks of Microsoft Corporation.

Netscape™ and Netscape Navigator™ are U.S. trademarks of Netscape Communications Corporation.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

Oracle7™ is a trademark of Oracle Corporation, Redwood City, California.

OSF/Motif® and Open Software Foundation® are trademarks of Open Software Foundation in the U.S. and other countries.

Pentium® is a U.S. registered trademark of Intel Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

This manual's title page contains the following identifying information:

- Software Version number, which indicates the software version
- Document release date, which changes each time the document is updated
- Software release date, which indicates the release date of this version of the software

To check for recent updates or to verify that you are using the most recent edition, visit the following URL:

http://ovweb.external.hp.com/lpe/doc_serv/

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

You can visit the HP OpenView support web site at:

http://www.hp.com/managementsoftware/support

This Web site provides contact information and details about the products, services, and support that HP OpenView offers.

HP OpenView online software support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valuable support customer, you can benefit by using the support site to:

- Search for knowledge documents of interest
- Submit enhancement requests online
- Download software patches
- Submit and track progress on support cases
- Manage a support contract
- Look up HP support contacts
- Review information about available services
- Enter discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to:

http://www.hp.com/managementsoftware/access_level

To register for an HP Passport ID, go to:

http://www.managementsoftware.hp.com/passport-registration.html

Contents

1 Introduction

The Reporter Database Schema

This document describes the hp OpenView Reporter products Reporter Database schema. This database schema is of interest to owners of the following products:

1. hp OpenView reporter.
2. Products that use OpenView embedded reporting, such as OpenView Internet services and OpenView operations for Windows.

How to Use This Document

This document is provided to give advanced users information about the nature of the database used to store their data and, if they so choose, perform customizations that go beyond configurations supported through the vendor. For supportable configurations, refer to the documentation provided with the product.

Recommendations in the use of this document

- Do not apply changes to production systems until you have thoroughly tested them on evaluation systems.
- Avoid adding data directly to the database. The interactions between applications and the database go beyond the scope of this document and cannot be predicted based on knowledge of the schema alone.

This document is provided for users who need to surpass limitations on supportable configurations of OpenView products. As such, users are on their own in dealing with the results of making modifications based on this document.

Additional Applications

For greater access to the data, the user may wish to consider acquiring the following applications for use in working with the Reporter Database:

- Crystal Decisions Crystal Reports Professional Edition 10
- SQL Server 2000, if the user is relying on a SQL Server 2000 database or the MSDE database provided under the default configuration of OpenView reporter.
- SQL Server 2005, if the user is relying on a SQL Server 2005 database or the MSDE database provided under the default configuration of OpenView reporter.

- Oracle 9i / 10g can be used to provide access to the database if the user is accessing a Reporter Database using Oracle 9i / 10g.

2 Collection Tables

Collections

Collections are entered into the Reporter database through the Gather process. This is usually through the `Gather.exe` or `GatherCODA.exe` utility, though it is possible for other OpenView products to store data as collections.

What is in Collections Tables

Collections store performance data, usually from one of OpenView's performance agents. For a full list of performance agents, please refer to the HP OpenView web site, *http://www.openview.hp.com/*. Collection tables have a one-to-one correspondence with metric lists, and the name of the metric list used to collect data is the same as that of the database table used to store data.

Collections tables have a variable number of fields, but the first five fields are common to all collections. These are necessary to the proper storage and maintenance of a collection table, and are entered by the collecting utility, rather than being defined as part of a metric list. Additional fields in the database are the result of data being collected as metrics in the metric list, and may vary in order, number, name and data type, but the first five are fixed.

Fields Common to all Collections

Column	Data Type	Notes
ID	AUTONUMBER	ID is an auto-numbered field unique to each record in the table
SYSTEMNAME	VARCHAR(200)	SYSTEMNAME is the system name as provided by the agent
DATETIME	DATETIME	DATETIME is the local time format for the start of the measurement interval
GMT	DATETIME	GMT is the universal time format for the start of the measurement interval
SHIFTNAME	VARCHAR(12)	SHIFTNAME is the name of the shift corresponding to the start of the measurement interval.

In addition, there is a field commonly associated with performance data collections, but is not necessarily collected. The INTERVAL metric must be part of the metric list that corresponds to the table for this column to exist. The interval represents the amount of time that the data represents. If the agent was only active for part of the interval, this field will have less than its normal value. The INTERVAL metric stores the amount of time that the agent was

collecting data for the predefined interval in seconds. Metric lists configured to gather one point every hour should have a default INTERVAL value of 3600 seconds. It is possible to have an INTERVAL greater than the predefined interval only under unusual circumstances (for example, during Daylight Saving Time, when an hour is repeated, it may appear as a '7200 second hour').

Additional Collection Table Fields

Column	Data Type	Notes
INTERVAL	INT	INTERVAL is the duration of the measurement interval in seconds

How to Use Collections Tables

The SYSTEMNAME field is most frequently used in joining collections tables or selecting a subset of data from them. DATETIME and SHIFTNAME can be used to restrict the amount of data presented in a report through the Reporter GUI, but are generally not used in joins with other tables.

When joining with another collection table, link the SYSTEMNAME to the SYSTEMNAME. When joining to an inventory table, such as the GROUPS or SYSTEMS tables, link the SYSTEMNAME in the collection table to the SYSTEMID in the inventory table. This should be done even if the inventory table has a field called SYSTEMNAME, as in the case of the SYSTEMS table.

GLOBAL

The GLOBAL table is the product of the GLOBAL metric list, which collects the GLOBAL class from the OpenView Performance Agent (formerly MeasureWare) and the OpenView Operations embedded performance component (also known as Coda). This table is used to provide system performance metrics at regular intervals, and is used in the default reports for CPU, disk, memory, and network data.

What is in the GLOBAL Table

This table provides the five fields common to all collections and the INTERVAL metric. In addition, 12 metrics are collected by default under reporter A.03.0x. Other fields may be added to the table by adding metrics to the GLOBAL metric list. Columns will be created in the database when the metric has been successfully gathered from an agent.

GLOBAL Table Fields

Column	Data Type	Notes
GBL_CPU_TOTAL_TIME	FLOAT	The time in seconds the CPU was not idle.
GBL_CPU_TOTAL_UTIL	FLOAT	The percentage of the interval the CPU was not idle.
GBL_DISK_PHYS_IO	INT	The number of physical disk I/O operations performed against local drives (non-network file systems).
GBL_FS_SPACE_UTIL_PEAK	FLOAT	The percentage of used to total disk space on the most full file system on the local system (non-network file systems).
GBL_MEM_PAGEOUT_RATE	FLOAT	The rate of memory page-outs to the disk per second.
GBL_NET_IN_PACKET_RATE	FLOAT	The rate of packets successfully received per second.
GBL_NET_OUT_PACKET_RATE	FLOAT	The rate of packets successfully sent per second.
GBL_ALIVE_PROC	INT	The number of alive processes with respect to the amount of time in the interval. An alive process is one that exists on the system.

Column	Data Type	Notes
GBL_ACTIVE_PROC	INT	The number of active processes with respect to the amount of time in the interval. An active process is one that uses CPU resources on the system.
GBL_STARTED_PROC	INT	The number of processes started during the interval.
GBL_RUN_QUEUE	FLOAT	Please see the agent documentation.
GBL_SWAP_SPACE_UTIL	FLOAT	Please see the agent documentation.
GBL_COLLECTOR	VARCHAR(18)	Please see the agent documentation.

In addition to metrics currently collected under OpenView reporter A.03.01, two metrics were collected under previous versions and may be present in databases created under obsolete versions of reporter.

Additional GLOBAL Table Fields

Column	Data Type	Notes
GBL_MEM_PAGE_REQUEST	INT	The number of page requests to or from the disk during the interval.
GBL_NET_PACKET_RATE	INT	The number of successful packets per second sent and received during the interval.

How to Use the GLOBAL Table

See the section on Collections tables in general for joining this table.

TRANSACTIONS

The TRANSACTIONS table is the product of the TRANSACTIONS metric list, which collects the TRANSACTIONS class from the OpenView Performance Agent (formerly MeasureWare). This class is not currently available from the OpenView Operations embedded performance component (also known as Coda). This table is used to store ARM transaction monitoring data.

What is in the TRANSACTIONS Table

This table provides the five fields common to all collections and the INTERVAL metric. In addition, information that is specific to ARM transactions is stored in 9 fields by default. A variable number of other metrics may be available from data sources that supply this metric, depending on the version of ARM supported. Fields may be added to the database by adding metrics to the TRANSACTIONS metric list. Columns will be created in the database when the metric has been successfully gathered from an agent.

TRANSACTIONS Table Fields

Column	Data Type	Notes
TT_NAME	VARCHAR(60)	Transaction name.
TT_APP_NAME	VARCHAR(60)	Transaction application name.
TT_COUNT	INT	The number of completed transactions during the interval.
TT_WALL_TIME_PER_TRAN	FLOAT	The average transaction time in seconds.
TT_ABORT	INT	The number of transactions that were aborted.
TT_ABORT_WALL_TIME_PER_TRAN	FLOAT	The average transaction time in seconds (aborted transactions).
TT_SLO_COUNT	INT	The number of completed transactions that violated the service level objective.
TT_SLO_PERCENT	INT	The percentage of completed transactions that violated the service level objective.
TT_SLO_THRESHOLD	FLOAT	The service level objective in seconds.

How to Use the TRANSACTIONS Table

See the section on Collections tables in general for joining this table.

APPLICATION

The APPLICATION table stores information about application performance as described in an agent's Parm file.

What is in the APPLICATION Table

This table provides the five fields common to all collections and the INTERVAL field. In addition, 6 other metrics are collected by default. Fields may be added to the database by adding metrics to the APPLICATION metric list. Columns will be created in the database when the metric has been successfully gathered from an agent.

APPLICATION Table Fields

Column	Data Type	Notes
APP_NAME	VARCHAR(20)	The name of the application.
APP_CPU_TOTAL_TIME	FLOAT	The total time in seconds that CPU was active with this application.
APP_DISK_PHYS_IO	INT	The total number of I/O processes performed by this application.
APP_MEM_VIRT	INT	The amount of virtual memory used by this application (in KB) at the end of the interval.
APP_ALIVE_PROC	INT	The number of alive processes with respect to the amount of time in the interval.
APP_ACTIVE_PROC	FLOAT	The time in seconds the application had windows in the active state.
APP_COMPLETED_PROC	INT	

How to Use the APPLICATION Table

See the section on Collections tables in general for joining this table.

UPTIME

The UPTIME table is new to OpenView reporter A.03.01. It is used to store agent uptime information in a way independent of other metric lists. Previous versions of OpenView reporter used the GLOBAL metric list and table for this calculation, which limited the length

of time agent uptime data could be stored to the length of time the user was willing to store the entire GLOBAL table.

What is in the Uptime Table

This table provides the five fields common to all collections and the INTERVAL field. The INTERVAL field is used to store agent uptime for extended durations. While it would have been possible to use any metric list's INTERVAL field to calculate agent uptime, using a separate metric list solely for uptime calculations allows users to store and report on uptime data for a number of days independent of the RETAINDAYS settings of other metric lists.

How to Use the Uptime Table

See the section on Collections tables in general for joining this table.

DOWNTIME

The DOWNTIME table is unique collection that is does not correspond directly to a metric list. It relates to the UPTIME metric list and Reporter's RepMaint utility uses data from the UPTIME table to update the DOWTIME table. This calculation is performed nightly, after data is collected but before reports are generated, and represents the amount of time out of a shift that the agent was up or down. Because data is only available through the agent, it's impossible for Reporter to determine when the system may have been up but the agent was down. Because the DOWNTIME table is not a direct result of a metric list, it is impossible to extend the table by adding more metrics to a metric list.

What is in the DOWNTIME Table

This table provides the five fields common to all collections. The DATETIME field is used to store the date, but because shifts span a number of hours, it should not be used to show a time within that date.

DOWNTIME Table Fields

Column	Data Type	Notes
SHIFTTIME	INT	The duration of the shift in minutes.
DOWNTIME	INT	The time in minutes for which data is unavailable during the shift.

How to Use the DOWNTIME Table

The DOWNTIME table can be used to present availability information on a per shift basis. For per interval availability, the UPTIME table should be used. This table is used to generate Agent Uptime report.

SYSDOWNTIME

The SYSDOWNTIME is unique collection table that has some fields which directly associated with metric list and one is calculated field which does not correspond directly to a metric list. It is calculated at time of gathering.

What is in the SYSDOWNTIME Table

This table provides the five fields common to all collections. The DATETIME field is used to store the date, but because shifts span a number of hours, it should not be used to show a time within that date.

DOWNTIME Table Fields

Column	Data Type	Notes
SHIFTTIME	INT	The duration of the shift in minutes.
GBL_SYSTEM_UPTIME_SECOND	DATETIME	This is system uptime in seconds.
GBL_BOOT_TIME	INT	This is time when system last boot up
SYSDOWNMINS	INT	This is total system downtime in minutes

How to Use the SYSDOWNTIME Table

The SYSDOWNTIME table can be used to present availability information on a per shift basis. This table is used to generate SYSTEM Uptime report.

3 Inventory Tables

Inventories

Inventories are entered into the Reporter database through the Discovery process. This is usually through the `Discovery.exe` or `Discover_ITO.exe` utility, though it is possible for other OpenView products to store data as collections.

What is in Inventory Tables

Inventories store lists of systems and associate these systems by name with information about the system provided by the agent or attributed to the system by the user or through an OpenView application. The SYSTEMID field is the only column common to all inventory tables.

Inventory Table Fields

Column	Data Type	Notes
SYSTEMID	VARCHAR(200)	The name of the system.

How to Use Inventory Tables

The SYSTEMID field is most often used in joining inventory tables with collections and other tables. The SYSTEMID in an inventory table corresponds to the SYSTEMNAME in a collection. The SYSTEMNAME in the inventory table should be ignored for the purpose of joining with collections and other tables in the Reporter Database.

SYSTEMS

The SYSTEMS table is the primary table for information about the system in the inventory as discovered through the performance agent.

What is in the SYSTEMS Table

The SYSTEMS table stores information about the system's software and hardware at the time it was last discovered. The SYSTEMS table will preserve only one instance of a system's configuration. Unlike performance data collections, changes over time are not tracked with multiple records. In addition to the SYSTEMID field, 21 other fields are stored in the SYSTEMS table.

SYSTEMS Table Fields

Column	Data Type	Notes
NETWORKID	VARCHAR(50)	A system's network. Often "Microsoft Windows Network", "ITO" or null.
DOMAINID	VARCHAR(50)	A system's domain.
SYSTEMID	VARCHAR(200)	The name of the system.
SYSTEMNAME	VARCHAR(200)	The agent's name for the system.
PROXYNAME	VARCHAR(200)	Null, unless the system is proxied.
DISCOVERYDATETIME	DATETIME	The last time the discovery process updated this system's data.
OSNAME	VARCHAR(50)	Operating system name, such as "NT" or "HP-UX".
OSRELEASE	VARCHAR(50)	The operating system major version, such as "4.0" for NT or "B.11.00" for HP-UX.
OSVERSION	VARCHAR(50)	The operating system minor version, such as "SvcPk 2" for NT or "A" for HP-UX.
MACHINETYPE	VARCHAR(50)	The processor type, such as "80686" or "9000 PA-RISC"
AGENT	VARCHAR(50)	Deprecated. This is the name of the primary agent. Please see the DATASOURCES table for retrieving data about all agents available.
CPUS	INT	The number of CPUs on the system.
DISKS	INT	The number of disks on the system.
NETWORKS	INT	The number of network interfaces on the system.
MEMORY	INT	The amount of physical memory (in KB) on the system.
SWAP	INT	The system's swap space (in KB).

Column	Data Type	Notes
DATASOURCES	INT	The number of data sources that are available from the primary agent.
GMTOFFSET	INT	The number of minutes from GMT that the system's local time is offset.
DSTCODE	INT	Whether the system observes DST (Daylight Saving Time).
EXCLUDE	INT	Whether the system should be omitted from data collection and reporting.
ITOAGENT	VARCHAR(50)	Deprecated. The secondary/event agent on the system.
IPADDRESS	VARCHAR(40)	The system's IP address. Note that this field is not currently in use, aside from systems discovered under ITO.
LONGHOSTNAME	VARCHAR(1024)	This holds the system name if the length of systemname field is greater than 200 characters. (In this case the systemname field will have uniquely generated GUID.)If the System name is less than 200 characters then this field would never be used.
OVDBNAME	VARCHAR(1024)	This Field holds name of ITO management server.

How to Use SYSTEMS Table

The SYSTEMID field is most often used to report on systems within a group. Other aspects of the system may be used to filter data for presentation.

DATASOURCES

The DATASOURCES table contains information about potential collections available from a system through performance agents.

What is in the DATASOURCES Table

The DATASOURCES table stores the SYSTEMID and 6 other fields. Multiple records may exist for a system, depending on the availability of data sources, classes, and agents.

DATASOURCES Table Fields

Column	Data Type	Notes
ID	AUTONUMBER	A generated record ID.
SYSTEMID	VARCHAR(200)	The name of the system.
DATASOURCE	VARCHAR(50)	The datasource name.
CLASS	VARCHAR(128)	The class name.
FIRSTTIME	DATETIME	The first time the class was found.
LASTTIME	DATETIME	The last time the class was found.
AGENT	VARCHAR(50)	The agent providing the class.

How to Use DATASOURCES Table

The SYSTEMID field is most often used to generate single-system or group reports.

GROUPS

The GROUPS table contains information about groups of systems.

What is in the GROUPS Table

The GROUPS table contains the SYSTEMID and 4 other fields.

GROUPS Table Fields

Column	Data Type	Notes
ID	AUTONUMBER	A generated record ID.
GROUPNAME	VARCHAR(255)	The name of the group.
SYSTEMID	VARCHAR(200)	The name of the system.
OWNER	VARCHAR(10)	
OWNER_GUID	VARCHAR(200)	

How to Use GROUPS Table

The SYSTEMID field is most often used in conjunction with the GROUPNAME field. If a join is performed between the SYSTEMID in the GROUPS table and the SYSTEMID (in an inventory table) or SYSTEMNAME (in a collection), the GROUPNAME can be selected in the WHERE clause of the SQL statement, allowing reports to be generated for systems in a group. OpenView reporter will automatically replace any default value in the selection statement in a Crystal Report template with the selected group when generating a group report.

CUSTOMER_SYSTEMS

The CUSTOMER_SYSTEMS table contains information about customers and systems.

What is in the CUSTOMER_SYSTEMS Table

The CUSTOMER_SYSTEMS table contains the SYSTEMID and 2 other fields.

GROUPS Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
CUSTOMER_NAME	VARCHAR(50)	The name of the customer.
SYSTEMID	VARCHAR(200)	The name of the system.

How to Use CUSTOMER_SYSTEMS Table

The SYSTEMID field is most often used in conjunction with the CUSTOMER_NAME field. If a join is performed between the SYSTEMID in the CUSTOMER_SYSTEMS table and the SYSTEMID (in an inventory table) or SYSTEMNAME (in a collection), the CUSTOMER_NAME can be selected in the WHERE clause of the SQL statement, allowing reports to be generated for systems belonging to a customer. OpenView reporter will automatically replace any default value in the selection statement in a Crystal Report template with the selected customer when generating a report for a customer.

4 Configuration Tables

Configuration

Configuration data are entered into the Reporter database directly by a user through the Reporter GUI or indirectly through the RepLoad utility and SRP files.

What is in Configuration Tables

Configuration tables store information about OpenView reporter.

How to Use Configuration Tables

Generally, configuration tables should not be used in generating reports. They are for the OpenView reporter product's use, and their schema is subject to undocumented change from a user's perspective.

AUTOGROUP

The AUTOGROUP table stores rules used for the automatic creation of groups. One row in the AUTOGROUP table corresponds to a single rule.

How to Use AUTOGROUP Table

Discovery applications (Discovery, Discover_ITO and Discover_Neutron) assign systems to groups based on the criteria in the AUTOGROUP table.

What is in the AUTOGROUP Table

Discovery applications (Discovery, Discover_ITO and Discover_Neutron) assign systems to groups based on the criteria in the AUTOGROUP table. The AUTOGROUP table uses total of 23 fields to store this information.

AUTOGROUP Table Fields

Column	Data Type	Notes
ID	AUTONUMBER	A generated record ID.
DESCRIPTION	VARCHAR(50)	The description of the autogroup rule.
GROUPNAME	VARCHAR(255)	The name of the group.

Column	Data Type	Notes
ALLTRUE	BIT	A Boolean value to indicate whether any (F) or all (T) criteria must be true.
NETWORKID	VARCHAR(255)	The pattern to match the network.
DOMAINID	VARCHAR(255)	The pattern to match the domain.
SYSTEMID	VARCHAR(200)	The pattern to match the system name.
SYSTEMNAME	VARCHAR(200)	The pattern to match the short name provided by the agent.
PROXYNAME	VARCHAR(255)	The pattern to match for the proxy.
OSNAME	VARCHAR(255)	The pattern to match for the OS.
OSRELEASE	VARCHAR(255)	The pattern to match for the OS release.
OSVERSION	VARCHAR(255)	The pattern to match for the OS version.
MACHINETYPE	VARCHAR(255)	The pattern to match for the machine type.
AGENT	VARCHAR(255)	The pattern to match for the agent.
CPUS	VARCHAR(50)	The pattern to match for the number of CPUs.
DISKS	VARCHAR(50)	The pattern to match for the number of disks.
NETWORKS	VARCHAR(50)	The pattern to match for the number of network interfaces.
MEMORY	VARCHAR(50)	The pattern to match for the amount of memory (in KB).
SWAP	VARCHAR(50)	The pattern to match for the amount of swap space (in KB).
DATASOURCES	VARCHAR(50)	The pattern to match for the number of data sources.
GMTOFFSET	VARCHAR(50)	The pattern to match for the timezone offset from GMT (in minutes).

Column	Data Type	Notes
ITOAGENT	VARCHAR(255)	The pattern to match for the ITO agent.
IPADDRESS	VARCHAR(255)	The pattern to match for the IP address.
DATASOURCE	VARCHAR(255)	The pattern to match for a data source.
DATACLASS	VARCHAR(255)	The pattern to match for a data class.

COMPLETED

The COMPLETED table stores the results of tasks run by the Reporter Scheduler.

How to Use COMPLETED Table

The COMPLETED table can be used to report on the success or failure of scheduled tasks, as is done in the Completed Activities report.

What is in the COMPLETED Table

Only the Scheduler should update the COMPLETED table. Please note that if a report is being generated by a scheduled task, the task controlling report generation (for example, RepCrys) will be listed as running while report generation takes place. When the report is complete, this information will be out of date. As a result, you may wish to filter out running tasks. The completed table uses 8 fields to store information about a task.

COMPLETED Table Fields

Column	Data Type	Notes
ID	AUTONUMBER	A generated record ID.
STARTTIME	DATETIME	The date and time a task was started.
STOPTIME	DATETIME	The date and time a task was stopped.
RUNSECONDS	INT	The number of seconds that the task ran before it stopped.
PROGRAM	VARCHAR(50)	The program's name.
PARAMETERS	VARCHAR(255)	The command line parameters.
STATUS	INT	The program's exit status.

Column	Data Type	Notes
SPECIAL	INT	Reserved field.

CONFIG_OPTIONS

The CONFIG_OPTIONS table stores miscellaneous settings used by OpenView reporter.

How to Use CONFIG_OPTIONS Table

The CONFIG_OPTIONS table can be used to report on the settings selected by users or established through report packages.

What is in the CONFIG_OPTIONS Table

The contents of the Reporter GUI's Options dialog box (File | Configure | Options) are found in the CONFIG_OPTIONS table. In addition, information about the packages snapped in through the Report Packages dialog box (File | Configure | Report Packages) are stored in this table.

COMPLETED Table Fields

Column	Data Type	Notes
KEY_NAME	VARCHAR(20)	The name of the configuration option.
DATA_TYPE	INT	The type of the configuration option.
DATA_INTVALUE	INT	The option value, if the type is an integer.
DATA_STRING	DATETIME	The option value, if the type is a date time string.
DATA_CHAR128	VARCHAR(128)	The option value, if the type is an integer.
DATA_DATETIME	DATETIME	The option value, if the type is a date time value.

CUSTOMER_MASTER

The CUSTOMER_MASTER table stores information about a customer, and is shared among OpenView products.

How to Use CUSTOMER_MASTER Table

The CUSTOMER_MASTER table can be used to report on the available customers and products associated with them.

What is in the CUSTOMER_MASTER Table

The customer name is available from the CUSTOMER_MASTER table, as well as the encrypted customer password used for directory access. The 26 lettered product fields correspond to present or future OpenView products.

CUSTOMER_MASTER Table Fields

Column	Data Type	Notes
CUSTOMERID	INT	A generated record ID.
CUSTOMER_NAME	VARCHAR(50)	The name of the customer.
VIEW_PSWD	VARCHAR(50)	The customer's password.
PRODUCT_A	INT	An OpenView product indicator.
PRODUCT_B	INT	An OpenView product indicator.
PRODUCT_C	INT	An OpenView product indicator.
PRODUCT_D	INT	An OpenView product indicator.
PRODUCT_E	INT	An OpenView product indicator.
PRODUCT_F	INT	An OpenView product indicator.
PRODUCT_G	INT	An OpenView product indicator.
PRODUCT_H	INT	An OpenView product indicator.
PRODUCT_I	INT	An OpenView product indicator.
PRODUCT_J	INT	An OpenView product indicator.
PRODUCT_K	INT	An OpenView product indicator.
PRODUCT_L	INT	An OpenView product indicator.
PRODUCT_M	INT	An OpenView product indicator.

Column	Data Type	Notes
PRODUCT_N	INT	An OpenView product indicator.
PRODUCT_O	INT	An OpenView product indicator.
PRODUCT_P	INT	An OpenView product indicator.
PRODUCT_Q	INT	An OpenView product indicator.
PRODUCT_R	INT	An OpenView product indicator.
PRODUCT_S	INT	An OpenView product indicator.
PRODUCT_T	INT	An OpenView product indicator.
PRODUCT_U	INT	An OpenView product indicator.
PRODUCT_V	INT	An OpenView product indicator.
PRODUCT_W	INT	An OpenView product indicator.
PRODUCT_X	INT	An OpenView product indicator.
PRODUCT_Y	INT	An OpenView product indicator.
PRODUCT_Z	INT	An OpenView product indicator.

DATABASES

The DATABASES table stores information about databases used by hp OpenView reporter.

How to Use DATABASES Table

The DATABASES table is used by RepCrys to access databases other than the Reporter Database for report generation.

What is in the DATABASES Table

The DATABASES table stores the OpenView reporter name for a database, as well as login information and the name of the ODBC DSN.

DATABASES Table Fields

Column	Data Type	Notes
DATABASENAME	VARCHAR(50)	The database name it is used within the Reporter Database, as in the REPORTDEFINITIONS table.
DATABASEUSER	VARCHAR(50)	The user name to log on to the database.
DATABASEPASSWORD	VARCHAR(50)	The (encrypted) password to log on to the database.
DATABASESERVER	VARCHAR(50)	The ODBC DSN configured that will provide access to the database.

DISCOVERDOMAINS

The DISCOVERDOMAINS table stores information about the discovery area.

How to Use DISCOVERDOMAINS Table

The DISCOVERDOMAINS table is used to instruct Reporter to discover all systems in a domain.

What is in the DISCOVERDOMAINS Table

The DISCOVERDOMAINS table stores the network name and the domain to be discovered.

DISCOVERDOMAINS Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
NETWORK_NAME	VARCHAR(50)	The name of the network.
[DOMAIN]	VARCHAR(50)	The name of the domain.

DISCOVERSYSTEMS

The DISCOVERSYSTEMS table stores information about the discovery area.

How to Use DISCOVERSYSTEMS Table

The DISCOVERSYSTEMS table is used to instruct Reporter to discover a single system.

What is in the DISCOVERSYSTEMS Table

The DISCOVERSYSTEMS table stores the network and domain name, if provided, as well as the system ID.

DISCOVERSYSTEMS Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
NETWORK_NAME	VARCHAR(50)	The name of the network.
[DOMAIN]	VARCHAR(50)	The name of the domain.
SYSTEM	VARCHAR(200)	The name of the system.

DLL_VERSIONS

The DLL_VERSIONS table stores information about the runtime environment used by OpenView reporter.

How to Use DLL_VERSIONS Table

The DLL_VERSIONS table provides information about the existence or absence of DLLs and their versions. This is used primarily for troubleshooting and reporting on the state of the OpenView reporter system.

What is in the DLL_VERSIONS Table

The DLL_VERSIONS table stores twelve fields describing found or missing files. Most of this information can be retrieved using the properties menu for files in Windows Explorer.

DLL_VERSIONS Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
FILENAME	VARCHAR(40)	The file's name.
FILEVERSION	VARCHAR(64)	The ASCII file version.
BINFILEVERSION	VARCHAR(64)	The binary file version.
PATHNAME	VARCHAR(255)	The file's path location.

Column	Data Type	Notes
FILETIMESTAMP	DATETIME	The date and time of the file's creation.
DESCRIPTION	VARCHAR(64)	The description provided in the file's properties.
COMPANYNAME	VARCHAR(64)	The company provided in the file's properties.
PRODVERSION	VARCHAR(64)	The ASCII product version.
BINPRODVERSION	VARCHAR(64)	The binary product version.
LINKTIMESTAMP	DATETIME	The date and time the file was linked.
DISCOVERYDATE	DATETIME	The date and time the file was found or found missing.

HOLIDAYS

The HOLIDAYS table stores information about scheduled holidays.

How to Use HOLIDAYS Table

The HOLIDAYS table is used in calculating downtime.

What is in the HOLIDAYS Table

The HOLIDAYS table stores the name of the holiday in the description and the day that it occurs.

HOLIDAYS Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
HOLIDAYDATE	DATETIME	The name of the holiday.
DESCRIPTION	VARCHAR(50)	A description of the holiday

METRICLISTINFO

The METRICLISTINFO table stores information about metric lists used in the collection of data from performance agents.

How to Use METRICLISTINFO Table

The METRICLISTINFO table is used in collecting data to the database and maintaining the associated table.

What is in the METRICLISTINFO Table

The METRICLISTINFO table contains 5 fields related to the collection and maintenance of the metric list's data.

METRICLISTINFO Table Fields

Column	Data Type	Notes
METRICLISTNAME	VARCHAR(128)	The name of the metric lists. This is also used as the name of a table to be created and store collected data.
DATASOURCE	VARCHAR(50)	The default data source. This can be '*' or be part of a name in conjunction with '*' to perform simple pattern matching.
CLASS	VARCHAR(128)	The name of the metric class. Often but not always, this is also the name of the metric list.
SUMMARIZATION	INT	The summarization interval requested of the agent.
RETAINDAYS	INT	The number of days collected data is kept. If the data's DATETIME field is older than the current day by more than the number of days in this field, the data will be culled by RepMaint.

METRICLISTMETRICS

The METRICLISTMETRICS table stores metrics associated with metric lists.

How to Use METRICLISTMETRICS Table

The METRICLISTMETRICS table is used in collecting data.

What is in the METRICLISTMETRICS Table

The METRICLISTMETRICS table contains the metric list and a metric that belongs to it.

METRICLISTMETRICS Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
METRICLISTNAME	VARCHAR(128)	The name of the metric list.
METRICNAME	VARCHAR(128)	The name of the metric.

METRICSPERGROUP

The METRICSPERGROUP table stores the association between metric lists and groups.

How to Use METRICSPERGROUP Table

The METRICSPERGROUP table is used in collecting data.

What is in the METRICSPERGROUP Table

The METRICSPERGROUP table stores the name of the metric list and the group. In addition it stores the default data source and a Boolean switch to indicate whether this is proxied data (in which case, the system name is stored differently).

METRICSPERGROUP Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
GROUPNAME	VARCHAR(255)	The name of the group.
METRICLISTNAME	VARCHAR(128)	The name of the metric list.
DATASOURCE	VARCHAR(50)	The default data source. This can be '*' or be part of a name in conjunction with '*' to perform simple pattern matching.
PROXY	BIT	True or False.

METRICSPERSYSTEM

The METRICSPERSYSTEM table stores metrics associated with metric lists and systems.

How to Use METRICSPERSYSTEM Table

The METRICSPERGROUP table is used in collecting data.

What is in the METRICSPERSYSTEM Table

The METRICSPERSYSTEM table stores the name of the metric list and the system. In addition it stores the default data source and a Boolean switch to indicate whether this is proxied data (in which case, the system name is stored differently).

METRICSPERSYSTEM Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
SYSTEMID	VARCHAR(200)	The name of the system.
METRICLISTNAME	VARCHAR(128)	The name of the metric list.
DATASOURCE	VARCHAR(50)	The default data source. This can be '*' or be part of a name in conjunction with '*' to perform simple pattern matching.
PROXY	BIT	True or False.

REPORTDEFINITIONS

The REPORTDEFINITIONS table stores information about the available reports.

How to Use REPORTDEFINITIONS Table

The REPORTDEFINITIONS table.

What is in the REPORTDEFINITIONS Table

The REPORTDEFINITIONS table.

REPORTDEFINITIONS Table Fields

Column	Data Type	Notes
REPORTNAME	VARCHAR(50)	The unique report name.
CATEGORY	VARCHAR(50)	The heading under which the report should appear in the links pages.

Column	Data Type	Notes
TEMPLATENAME	VARCHAR(255)	The crystal reports template used to report on all systems. This is the file name and path relative to the data directory.
GROUPTEMPLATENAME	VARCHAR(255)	The crystal reports template used to report on a group. This is the file name and path relative to the data directory.
SYSTEMTEMPLATENAME	VARCHAR(255)	The crystal reports template used to report on a single system. This is the file name and path relative to the data directory.
HTMLDIRECTORY	VARCHAR(255)	The output directory used to store the resulting HTML page and associated image files. Subdirectories are made to hold system and group versions of the report. A customer-segmented report will be placed in a subfolder to the Webpages directory based on the customer name, then the report name, then a group or system name if applicable.
PRINTERNAME	VARCHAR(50)	Unused at the present time.
DESCRIPTION	VARCHAR(80)	The text that will appear as a link to the report.
DATERANGE	VARCHAR(50)	The date range period as selected through the Reporter GUI. A null in this field will take the default value.
DATABASENAME	VARCHAR(50)	The database to be provide the data for the report See the DATABASES table for more information.
SHIFT	VARCHAR(50)	The shift to be used for all instances of this report. See the SHIFT table for more information.

Column	Data Type	Notes
PAGE_BREAKS	VARCHAR(1)	The letter "T" if the report should be split into multiple pages with links at the bottom.
RUNTIME	INT	The amount of time in minutes before the report will be timed out (killed).

REPORTFAMILIES

The REPORTFAMILIES table associates reports with report families, used to organize and present the reports.

How to Use REPORTFAMILIES Table

The REPORTFAMILIES table is used in providing links to generated reports.

What is in the REPORTFAMILIES Table

The REPORTFAMILIES table contains the report and the family to which it belongs.

REPORTFAMILIES Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
REPORTNAME	VARCHAR(50)	The name of the report.
FAMILY	VARCHAR(50)	The name of the report family.

REPORTSPERGROUP

The REPORTSPERGROUP table associates reports with groups.

How to Use REPORTSPERGROUP Table

The REPORTSPERGROUP table is used in generating reports. If an association exists between a report and a group, the report is generated for that group. If no association is found, the report is not generated. A GROUPNAME of "ALL" is required for a report to be generated for all systems.

What is in the REPORTSPERGROUP Table

The REPORTSPERGROUP table contains the report name and group name.

REPORTSPERGROUP Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
GROUPNAME	VARCHAR(255)	The name of the group.
REPORTNAME	VARCHAR(50)	The name of the report.

REPORTSPERSYSTEM

The REPORTSPERSYSTEM table associates reports with systems.

How to Use REPORTSPERSYSTEM Table

The REPORTSPERSYSTEM table is used in generating reports. If an association exists between a report and a system, the report is generated for that system. If no association is found, no single-system version of the report is generated. The system may still appear in reports for all systems or a group of systems, but an instance for that report by itself is not generated.

What is in the REPORTSPERSYSTEM Table

The REPORTSPERSYSTEM table contains the report name and system name.

REPORTSPERSYSTEM Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
SYSTEMID	VARCHAR(200)	The name of the system.
REPORTNAME	VARCHAR(50)	The name of the report.

REPORTSPERCUSTOMER

REPORTSPERCUSTOMER table is used to determine what all reports to be generated for a given customer.

How to Use REPORTSPERCUSTOMER Table

When user creates/adds a customer to Reporter an entry is made into this table for the created/added customer and all the reports, which could be assigned to this customer. From the Reporter GUI when a report is assigned to this customer, a flag in REPORTSPERCUSTOERM table is turned on indicating that this report is to be generated for the given customer.

What is in the REPORTSPERCUSTOMER Table

The REPORTSPERCUSTOMER table contains the name of the customers, list of reports that could be assigned to this customer and a flag that tells what all reports to be generated for this customer.

Primary key is the combination of CUSTOMER_NAME and REPORTNAME.

REPORTSPERCUSTOMER Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
CUSTOMER_NAME	VARCHAR (50)	Customer name.
REPORTNAME	VARCHAR (127)	Report name.
GENERATE_REPORT	INT	Flag that tells whether this report to be generated for this customer or not (value of 1 tells generate report and a value of 0 tells don't generate report).

SCHEDULE

The SCHEDULE table tracks scheduled tasks for the Reporter Scheduler.

How to Use SCHEDULE Table

The Reporter Scheduler uses the SCHEDULE table to determine when and how a task should be run.

What is in the SCHEDULE Table

The SCHEDULE table contains ten fields related to scheduled tasks.

SCHEDULE Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.

Column	Data Type	Notes
JOBTIME	DATETIME	The time the task is scheduled to run.
PROGRAM	VARCHAR(50)	The name of the executable. This may include the path to locate the file, either absolute or relative to the data directory, if the file is not on the environment PATH.
PARAMETERS	VARCHAR(255)	Command line parameters used in running the scheduled task.
LASTTIME	DATETIME	The last time the task was run.
STATUS	INT	The return value for the executable.
JOBRESET	INT	The interval at which the task should repeat.
RUNTIME	INT	The length of time in minutes the task is allowed to run before it times out.
MAXCONCURRENTPROGRAMS	INT	The number of programs that can be run at once. If a program has a concurrency of one, it is the only one that can be run at that time. If it has a concurrency of two, two can be run, but will not if the other program has a concurrency of one. A concurrency of zero is a special case; a program with a concurrency of zero will override normal concurrency controls and will run regardless of any other running tasks.
SPECIAL	INT	Unused at this time.

SHIFT

The SHIFT table stores information about shifts used to divide days of the week into shifts.

How to Use SHIFT Table

The SHIFT table is used to assign a shift to collected data during database maintenance. It is also used in downtime calculation.

What is in the SHIFT Table

The SHIFT table contains the name of the shift, its start and stop times, and the day of the week for which a particular start and stop time apply. Each day of the week will have an entry, if it is part of a named shift. Time that does not fall under a named shift becomes part of the OFFSHIFT. Shifts are universal to the Reporter Database may not overlap, as only one shift can be applied to a record.

SHIFT Table Fields

Column	Data Type	Notes
ID	INT	A generated record ID.
SHIFTNAME	VARCHAR(12)	Shift name.
SHIFTDOW	INT	Day of week.
STARTTIME	INT	Start time in minutes from midnight.
STOPTIME	INT	Stop time in minutes from midnight.

5 Database ER Diagram

Relationship between tables (installed by default) in Reporter Database

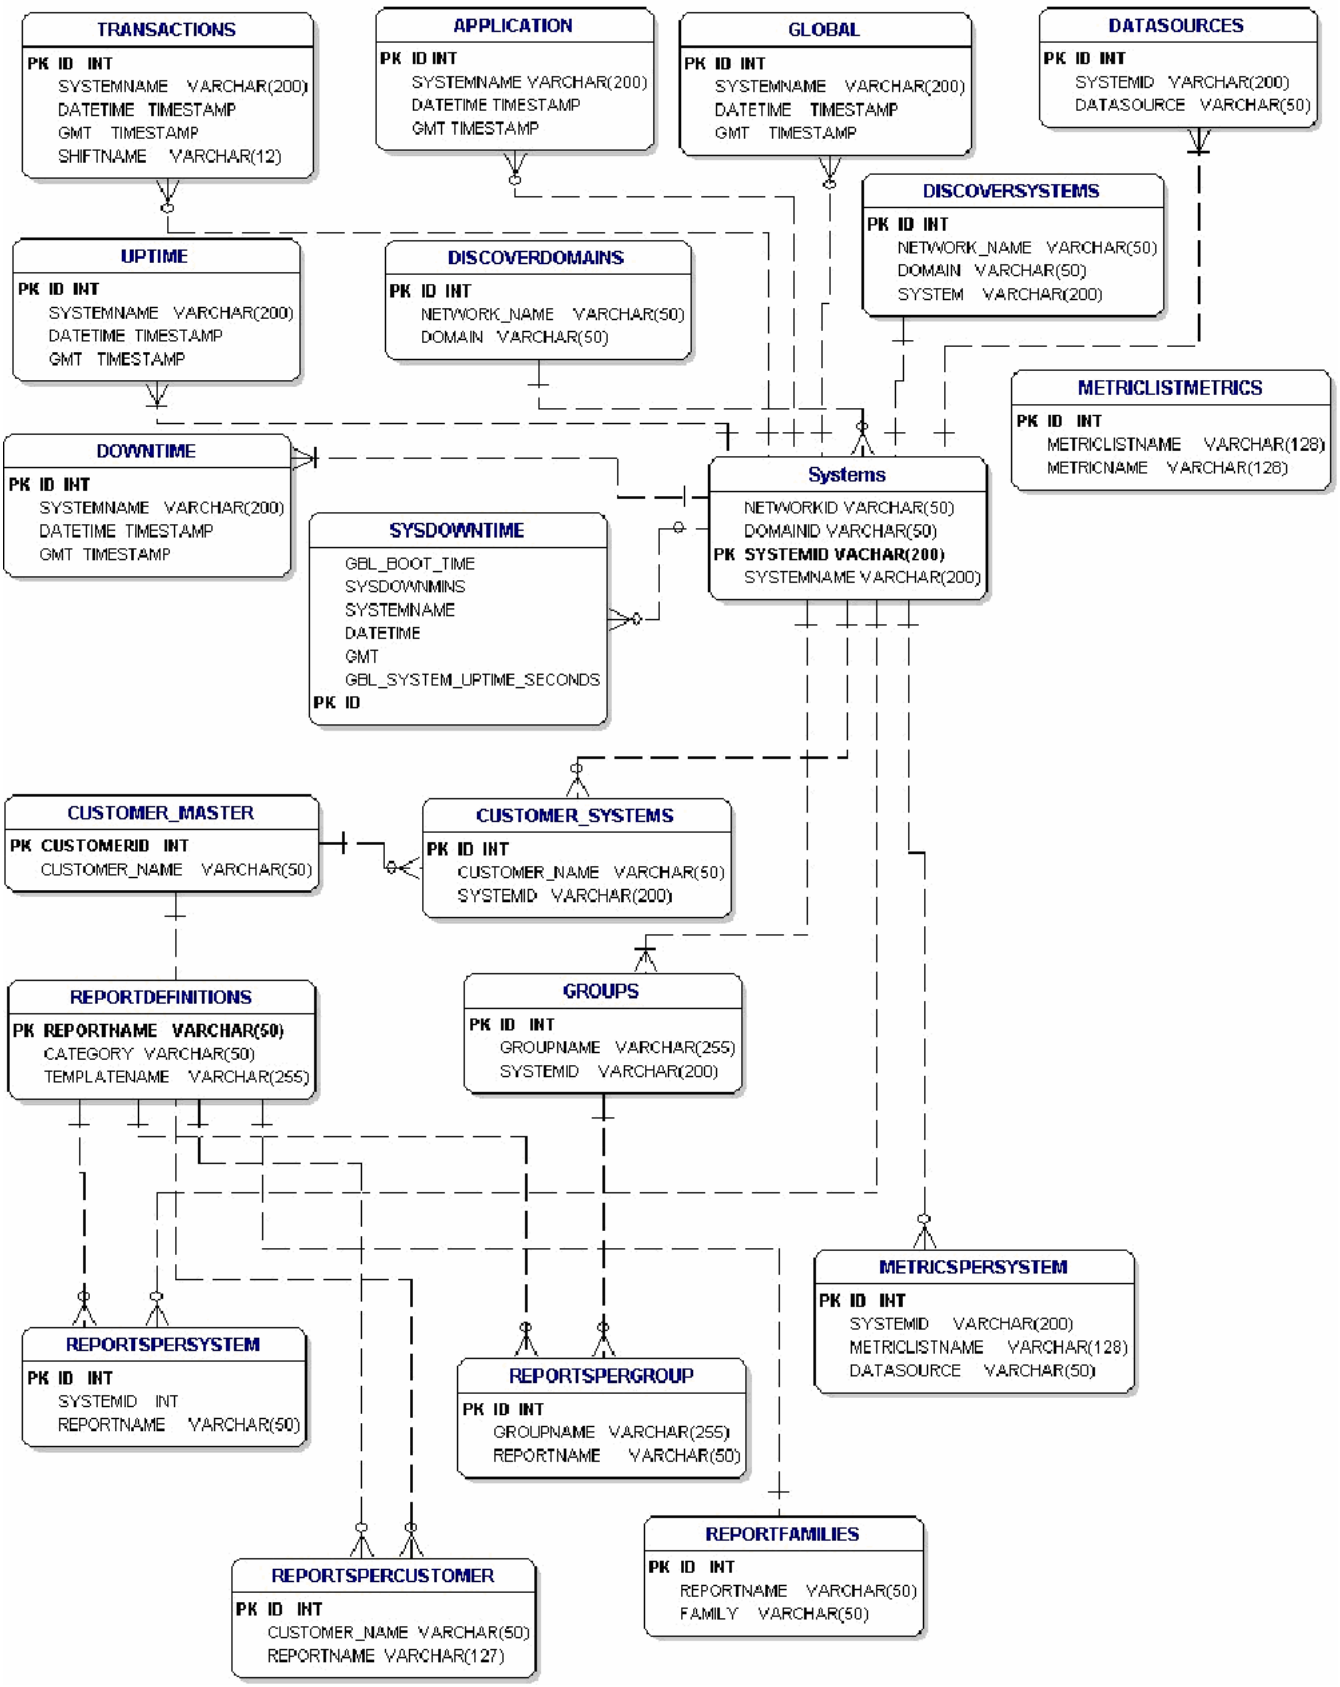

Configuration Tables

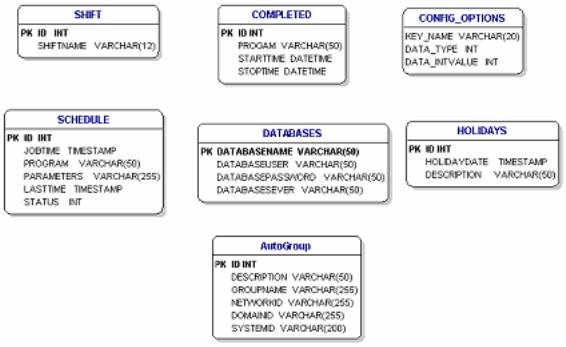

Note:

(a) Every Collection table is related to 'Systems' table in the same way as 'Application', 'Transaction' and 'Global' tables are related to 'Systems' table.

(b) The relationships explained in this diagram are NOT enforced by way of database procedures or triggers. As can be seen, most of the tables have an auto generated INTEGER as Primary Key and integrity of data is enforced by Reporter modules that insert/delete/update these tables.