Chapter 27: System Administration

SEER*DMS is a Web-based application that interacts with the registry's Oracle database. The registry's information technology (IT) staff are responsible for a variety of tasks related to maintaining the integrity and security of the database, providing ad hoc programming support to registry staff, and maintaining the hardware systems that support SEER*DMS. The first sections of this chapter provide a broad overview of the duties of the registry system administrators and a cross-reference to more specific information related to these topics.

System Tasks are available within SEER*DMS to execute processes in batch mode. Management and IT staff should develop registry-specific policies and procedures for the use of these utilities. The System Administration page in SEER*DMS allows you to monitor system logs and review registry-specific configuration settings and algorithms.

In this chapter, you'll learn about

- Technical Skills Required to Maintain SEER*DMS
- Overview of IT Responsibilities
- SEER*DMS Technical Support
- System Tasks
    - System Tasks to Build CTCs from Non-abstract Records
    - Patient Set Edits Task
    - Identify Patients for Active Follow-up (AFUP)
    - Task to Purge Death Certificate Records
    - Task to Purge Supplemental Records
    - Task to Rematch Unlinked Records
    - Aborting a System Task
- System Administration Page
    - Configuration
    - Environment
    - Hibernate
    - Logs
    - Memory
    - System Files
    - System Properties
    - Threads

Technical Skills Required to Maintain SEER*DMS

In addition to proficiency in navigating within the MS Windows environment, the collective skill sets of your staff should include the skills described below.

Server Management
The registry's network administrator should have the skills to manage day to day operational issues such as server startup and shutdown, backup and recoveries, etc. Instructions written specifically for your registry's system configuration are posted in the System Administration section of your registry folder on the SEER*DMS Web Portal.

Ability to Write SQL Queries
IT staff must have the ability to write and optimize SQL queries and have a full understanding of the SEER*DMS database structure. If the registry has limited SQL experience, formal SQL training should be considered. The SEER*DMS Web portal (http://seer.cancer.gov/seerdms/portal) contains sample SQL as well as links to SQL references. SQL for the system reports is embedded in the report XML that may be viewed via the System Administration page.

Understanding of Data Structures
In order to put the SQL to good use, the IT staff must also have a full understanding of the SEER*DMS database. This would include a mapping to the legacy system, a mapping to the data fields displayed in the SEER*DMS editors, and a complete understanding of all table relationships. Diagrams are available on the SEER*DMS Web portal and data mappings are provided in the Field Mappings section of the Help menu.

Ability to Create Reports using an External Reporting Package
The SEER*DMS application includes numerous reports that were determined to be useful to all registries and could be implemented generically. If the registry requires additional reports, the registry's IT Staff must be able to generate and maintain the reports in an external reporting package that interfaces with the Oracle database.

Proficiency in an External Programming or Scripting Tool
Registry IT Staff must have the ability to restructure data files into SEER*DMS supported file formats and to create registry-specific extracts. To do this in a reliable and repeatable fashion, the IT staff must be able to write and maintain programs or scripts to perform these tasks. For example, registries often need to write ad hoc programs to create text files in the Generic Supplemental Format containing data sent by non-medical organizations. The files sent by these organizations are typically text files in non-standard formats. Any programming tool may be used for this purpose. However, many external programming tasks require connectivity to the Oracle database. Therefore, the IT staff must select and learn to use a tool that can interface directly with the Oracle database.

Several scripts, written in Perl, are delivered with SEER*DMS. If the registry wishes to utilize these scripts, it is the responsibility of the IT staff to maintain, update, or re-write them as necessary.

Overview of IT Responsibilities

The following summarizes the responsibilities of the registry's IT staff. Information and instructions related to these tasks are available in the *SEER*DMS Technical Reference* and various sections of this User's Manual, as indicated below.

1. Maintain valid licenses and maintenance agreements for Oracle, Linux, and any other server-side software required to implement SEER*DMS at your registry. The specific list of software required is provided in the registry-specific section of the SEER*DMS Portal (http://seer.cancer.gov/seerdms/portal).

2. Using the SEER*DMS system itself, registry IT staff will:

   a. Execute the System Tasks described in this chapter. These include Build CFOs, Build DCOs, Patient Set Edits, etc.

b. Extract data from the SEER*DMS database and save the data in the format required for a data submission or for use in an analytic tool (see *Chapter 24: Creating Reports and Extracting Data*).

c. Monitor the system logs displayed in the SEER*DMS System Administration tool and report problems to IMS staff.

d. Import data files as described in *Chapter 5: Importing Data Files*.

e. Provide technical support to registry staff, as needed. This requires a general knowledge of the system features described throughout this manual and the registry-defined workflow detailed in the SEER*DMS Technical References.

3. Using registry-defined strategies and systems, registry IT staff will:

   a. Maintain on-site and off-site backups of the SEER*DMS servers.

   b. Incorporate SEER*DMS servers into the registry-wide disaster recovery plan.

4. Using an external reporting package such as Crystal reports, IT staff will:

   a. Create new external reports as requested by registry staff. This will require a working knowledge of SQL and the reporting software.

   b. Acquire and install the reporting software on the PCs of all registry staff who need to run external reports.

5. Using Windows tools, registry IT staff will:

   a. Configure the registry's personal computers and maintain current versions of client software required for SEER*DMS as described in *Chapter 3: Using SEER*DMS*.

   b. Provide on-site technical support for all PC and network-related issues.

   c. Write, test, and execute routines to pre-process data files that are not formatted per the registry-defined specifications for SEER*DMS import files.

   d. Write, test, and execute routines to extract data from the SEER*DMS database according to registry-specific requirements. Standard extracts for SEER, NAACCR, and other submissions are available within the SEER*DMS application.

   e. Write, test, and execute routines to process extracted data files (running edits, making data files available to staff, submitting data files to NCI and other organizations).

   f. Monitor external logs and report problems to IMS staff.

   g. Manage system disaster recovery procedures.

6. Using an external SQL tool, registry IT staff will:

   a. Fill ad hoc requests for data, reports.

   b. Create and maintain registry-defined Oracle tables. A schema is available on the same server as the data warehouse for registry-defined tables. The registry may use this schema to create tables for research, linkage to SEER*DMS data warehouse tables, and other purposes. Tables in this schema cannot be linked to the SEER*DMS production database. Registry staff are responsible for all activities related to the maintenance and use of this schema. This includes the creation of tables, populating tables with values from SEER*DMS data warehouse tables or external sources, deletion of tables, updates

of values within tables, and the quality control of all scripts and programs involved in these tasks.

SEER*DMS Technical Support

Critical issues that inhibit registry operations should be reported immediately to IMS staff. Non-critical issues should first be triaged by the registry's local technical support team.

The registry must appoint a local technical support team who can respond to routine issues. This team should include representatives from IT, editing, and management. The registry team will triage routine technical support issues and perform an initial investigation. If an issue cannot be resolved on site or requires a system change, provide a description of the problem and the results of the investigation to IMS. Unresolved issues and requests for new features should be submitted via the Technical Support Squish project (https://www.squishlist.com/seerdms/support).

System Tasks

Requires system permission: *system_administration* enables access to all tasks; *view_management_tasks* enables access to tasks that do not change any fields in the database

The Tasks section of the SEER*DMS System menu provides access to utilities or "tasks" that enable you to implement processes in batch mode. For example, the Build DCO system task enables you to auto-build batches of DCO cases from death certificate records as opposed to using the SEER*DMS editor to build each DCO case individually.

The majority of the System Tasks require the *system_administration* permission. Tasks that do not implement a change to any data in the database may only require the *view_management_tasks* permission.

Management and IT staff should develop registry-specific standard operating procedures (SOPs) related to the System Tasks. The SOPs should address these topics: scheduling considerations, procedures that must be completed prior to executing the task to ensure that the task is not executed on data inappropriately (e.g., that all death clearance processes are performed prior to executing the Build DCO system task), and a systematic review of data and relevant reports when the task finishes execution.

The System Tasks enabled in SEER*DMS vary by registry. All system tasks are listed below and are described in subsequent sections of this chapter:

- System Tasks to Build CTCs from Non-Abstract Records:
    - Build CFOs
    - Build DCOs
    - Build SHOs
- Identify Patients for Active Follow-up (AFUP)
- Patient Set Edits
- Purge Death Certificate Records
- Purge Supplemental Records
- Rematch Unlinked Records

System Tasks to Build CTCs from Non-Abstract Records

Requires system permission: *system_administration*

There are three separate, record-specific tasks for building CTCs from reportable, unlinked records. These tasks enable you to submit data to SEER or other sources despite the fact that an abstract is currently unavailable. To reduce the number of CTCs built solely from these records, it is recommended that you follow your registry's Death Clearance and Casefinding procedures prior to executing any of these tasks. Refer to *Chapter 17: Death Clearance* and *Chapter 21: Managing Abstracting Assignments* for further instructions.

- Build CFOs – build CTCs from casefinding, HL7 E-Path, or NAACCR Abstract records with a value of 7 for class of case (NAACCR abstract – class of case 7 is only an option in Hawaii's version of SEER*DMS).
- Build DCOs – build CTCs from death certificate records
- Build SHOs – build CTCs from short health records

If a reportable record is not linked to any Patient Set in the database, a new Patient Set will be created. If the record is linked to a Patient Set but not to a CTC in that Patient Set, a new CTC will be created in that Patient Set. An unassigned Visual Edit Patient Set task will be added to the worklist for each newly created or modified Patient Set.

Alternatively, you may use the record menu in the SEER*DMS editor to create a CTC for each record individually. Instructions are provided in the *Chapter 21: Managing Abstracting Assignments*.

To build CTCs from unlinked, reportable records:

1. Use reports and other means to verify that all appropriate death clearance or casefinding procedures have been completed. Execute system reports and registry-defined external reports to stratify the records by facility or other relevant factors. If necessary, use the appropriate system report to list candidate records and review the records manually. The reports below are available within the system. The following system reports show the number of records that will be considered by each task. Instructions for generating system reports are provided in *Chapter 24: Creating Reports and Extracting Data*.

    a. Frequency Reports – Use these reports if you are interested in the number of records at a particular facility that will be used to build CFO, DCO, or SHO cases. These reports show an estimate that is not as accurate as the count displayed on the screen when you set the build task's options. However, only a total count is shown in the build task.

        i. RPT-067B – Frequencies for the Build CFOs task.

        ii. RPT-068B – Frequencies for the Build DCOs task.

        iii. RPT-089B – Frequencies for the Build SHOs task

    b. Record Listings – You may use these reports if you wish to review the individual records being considered by the build tasks.

        i. RPT-067C – Candidates for the Build CFOs task.

        ii. RPT-089C – Candidates for the Build SHOs task.

2. To continue with the Build CTC task, click **System > Tasks**.

3. Based on the record-type you wish to use in the build, click the link for the appropriate Build CTC system task. Because of the demand on resources, the task links for the Build tasks will be disabled if any of the Build tasks are currently executing.

4. Set the parameters for the task:

   a. Enter a value for **Year**. The year parameter is required for all Build CTC system tasks. Registry-specific requirements determine which database field is used for year (e.g., date of diagnosis or screening year for CFO records, date of death for DCO records, date of diagnosis for SHO records). Please refer to the documentation section of the screen to determine the year used for each record type in your registry's configuration.

   b. If you would like to specify a particular Facility Source, enter the appropriate ID in the **Facility** field. This parameter is not available for the Build DCOs system task.

   c. You may specify the type of record considered by the Build CFO system task. The specific types of record considered by this task vary by registry. In all registries, casefinding and HL7 E-Path records may be used. NAACCR abstract records with a value of 7 for class of case may also be used in some registries. Consult registry management to determine which types of records should be used in the build.

5. Click **Recalculate** to view the number of records that will be considered by the task.

6. You may enter text related to this task in the **Comment** field. The comment for the last build of each type is stored in the database (*utility_history* table).

7. Click **Start**. SEER*DMS will begin executing the appropriate automated Build CTC task and return you to the System Tasks page.

8. After the task completes, follow the instructions below to review a report showing the results of the task.

SEER*DMS creates an unassigned Visual Edit Patient Set worklist task for each Patient Set created or modified by the build. To find the editing tasks, enter "DCO Build", "CFO Build", or "SHO Build" in the worklist's **Information** filter. To complete these editing tasks, follow your registry's guidelines for editing missing CTC data fields and resolving errors in patient data. The build task also generates a report listing each Patient Set created or modified by the build.

To review a Build Report for any of the Build CTCs System Tasks:

1. Select View > Home.

2. Click **Report Output** in the **My Tasks** section of your worklist summary. If this link is not shown, then the build may have failed or may still be running.

3. If you have a large number of Report Output tasks, use the filter to search for the results of the Build CTC System Task:

   a. Enter "Build" in the **Information** filter to search for a worklist task related to any of the Build CTC System Tasks. Or you may indicate the specific task by specifying "Build CFOs", "Build DCOs", or "Build SHOs" in the filter.

   b. Click **Apply**.

4. Click on the **Task ID** to open the Report Output task.

5. Click **View**. Depending on your browser settings, the report may open automatically or you may need to click **Open**. The PDF will open in an Adobe Acrobat window.

6. Use the Adobe controls to print or save this report.

Patient Set Edits Task

All edits are executed each time a patient set is opened, validated, or saved in the SEER*DMS editor. The Patient Set Edits system task enables you to re-execute the edits on patient sets in the database. You may run the edits on all patient sets or on a cohort defined by year of diagnosis. Use the Patient Set Edits task to ensure that new or modified edits are evaluated.

A polisher is a system utility that derives, calculates, or assigns data field values. For example, polishers are used to derive collaborative stage variables; assign census tract based on address; and calculate the age at diagnosis based on date of birth and date of diagnosis. When a patient set is opened, saved, or validated, a polisher will be executed if the value of a related data item changes.

Polisher classes are defined in the registry configuration files. These "classes" are categories which are used to control their execution sequence during system processes, and to determine whether a polisher is available in the Patient Set Edits system task. In the Patient Set Edits task, you have the option of executing the polishers in the "Pre-Edits" and "Post-Edits" classes (the polishers in these classes are listed on the screen when you open the Patient Set Edits system task). Typically, these should not be run when this task is executed and should only be considered during the initial transition of data to SEER*DMS. You also have the option of executing one additional polisher in the "Standard" and "Extra-Edits" polisher classes. In SEER*DMS, the polisher class definitions can be viewed by selecting Configuration from the System Administration page (the *system_administration* permission is required). Additional information related to polishers can be found in the *SEER*DMS Technical Reference*.

To re-execute the edits for some or all patient sets in the database:

1. Click **System** > **Tasks**.
2. Click the **Patient Set Edits** link.
3. To limit the edits to data by year of diagnosis, enter the **Start Year**. Patient sets with a diagnosis date during or after this year will be considered.
4. If you would like to define an end range, enter an **End Year**. If the end date is not specified, today's date will be used by default.
5. To include data with unknown year of diagnosis, set **Include Unknown Year** to *Yes*.
6. If you wish to execute polishers in the Pre-Edit and Post-Edit classes, set **Run Edit Polishers** to *Yes*. It is recommended that this option be set to *No* unless there is a specific need related to the transitioning of data into SEER*DMS.
7. To execute a single polisher from the "Standard" or "Extra-Edits" classes, select a polisher from the **Extra Polisher** drop-down list.
8. You may enter text related to this task in the **Comment** field. The comment for the last execution of the task is stored in the database (utility_history table).
9. Click **Start**.

The edits will be re-evaluated for each patient set in the cohort. In order to avoid creating an inordinate number of worklist tasks, a Resolve Patient Set Errors task will *not* be created for each patient set with an edit error. If the logic of a new or modified edit is implemented incorrectly, it could erroneously create an edit error for a large number of patient sets. Therefore, you must use reports rather than tasks to identify the patient sets with errors and to evaluate the error levels in

the patient set data. Two system reports are available for identifying the edit errors that were triggered and the patient sets that are involved:

- *RPT-064A: Frequency of Edit Errors in the Patient Set Data*- Run this report to evaluate the error levels in the patient data. In order to verify that modifications to edits, polishers or data did not have unexpected results, set the parameters to generate frequencies of errors for all edits in all patient set data.
- *RPT-064B: Patient Sets with Edit Errors* - Run this report to obtain a listing of Patient Sets with an error related to a particular edit.

Instructions for creating reports are provided in *Chapter 24: Creating Reports and Extracting Data*.

Identify Patients for Active Follow-up (AFUP)

Requires system permission: *system_administration*

Registry-specific algorithms are used to identify patient sets for active follow-up (AFUP). These criteria are typically based on vital status, date of last contact, and whether the patient's follow-up status is monitored by the registry (e.g., non-reportable cases may not be followed by some registries).

The "Identify Patients for AFUP" task creates an open 'AFUP Need' for each patient set requiring active follow-up. You may use the AFUP Manager to facilitate and track communications to the patients, their physicians or other contacts. The AFUP Manager also enables you to enter follow-up data as it is received. An AFUP Need is closed when the patient's vital status and date of last contact indicate that the need has been fulfilled. Instructions for using the AFUP Manager are provided in *Chapter 16: Follow-up*.

To populate the AFUP table:

1. Click **System > Tasks**.
2. Click the **Identify Patients for AFUP** link.
3. Enter a date in DOLC Cutoff that defines whether a patient set requires active follow-up. A patient set will be considered if the patient set's value for vital status is alive and the date of last contact is prior to this date.
4. Click **Recalculate** to view the number of patient sets that meet the criteria. If the number of patient sets is unacceptably high based on registry policy, adjust the DOLC Cutoff and recalculate.
5. You may enter text related to this task in the **Comment** field. The comment from the last execution of the task is stored in the database (utility_history table).
6. Click **Start**. SEER*DMS will create an AFUP Need for each patient set which matches the registry's active follow-up criteria. Once the task completes, the follow-up staff may use the AFUP Manager to initiate follow-up procedures for these patients.

Task to Purge Death Certificate Records

Requires system permission: *system_administration*

The purpose of the Purge Death Certificate Records system task is to permanently remove unused death certificate records from the system. Typically, this task would be used to delete old, non-reportable records. Consult with registry management and write a query or report to review the

records prior to executing this task. Purge Death Certificate Records will delete records which meet these criteria:

- The record is not reportable.
- The record is unlinked. Linked records may have contributed to patient set data and, therefore, cannot be deleted.
- The record is not the focus of a workflow task. If you need to delete a batch of records which are involved in workflow tasks then you must terminate the tasks prior to using the Purge Death Certificate Records task.
- The record's type and import date meet the specifications that you set in the task's options.

To permanently remove unlinked Death Certificate records from the database:

1. Click **System > Tasks**.
2. Click the **Purge Death Certificate Records** link.
3. Specify the **Date of Death Prior To**. If an Import ID is not specified, the task will delete records with a date of death (*record.date_of_last_contact*) prior to this date or a null value for *record.date_of_last_contact_yyyy*.
4. If an **Import ID** is specified, all records in the import will be considered by the task (the date parameter will be ignored). Records in the import that are linked or in the workflow will not be deleted.
5. You may enter text related to this task in the **Comment** field. The comment from the last execution of the task is stored in the database (utility_history table).
6. Click **Recalculate** to view the number of records that will be considered by the task. Determine whether this value is reasonable.
7. Click **Start**.

Task to Purge Supplemental Records

Requires system permission: *system_administration*

The purpose of the Purge Supplemental Records system task is to permanently remove unused records from the system. Typically, this task would be used to delete old records from non-medical sources (DMV, SSA, etc). Consult with registry management and write a query or report to review the records prior to executing this task. Purge Supplemental Records will delete records which meet these criteria:

- The record is unlinked. Linked records may have contributed to patient set data and, therefore, cannot be deleted.
- The record is not the focus of a workflow task. If you need to delete a batch of records which are involved in workflow tasks then you must terminate the tasks prior to using the Purge Records task.
- The record's DOLC and import date meet the specifications that you set in the task's options.

To permanently remove unlinked Supplemental records from the database:

1. Click **System > Tasks**.
2. Click the **Purge Supplemental Records** link.

3. Specify the **DOLC Prior To**. If an Import ID is not specified, the task will delete records with a date of last contact prior to this date or a null value for *date_of_last_contact_yyyy*.

4. If an **Import ID** is specified, all records in the import will be considered by the task (the date parameter will be ignored). Records in the import that are linked or in the workflow will not be deleted.

5. You may enter text related to this task in the **Comment** field. The comment from the last execution of the task is stored in the database (utility_history table).

6. Click **Recalculate** to view the number of records that will be considered by the task. Determine whether this value is reasonable.

7. Click **Start**.

Task to Rematch Unlinked Records

Requires system permission: *system_administration*

The Rematch Unlinked Records system task reloads records into the workflow. Each record enters the matching task in the workflow branch that appropriate for the record type (main workflow or passive follow-up, for example). The record is not sent through the tasks that precede matching: automatic record recoding, record edits, or screening. The value for reportability assigned when the record was first loaded is retained.

Based on the record type and match result, this task may create manual matching or consolidate FUP tasks in the worklist. Consult with registry management prior to executing this task. This task was originally designed to reload records for the purpose of passive follow-up. For example, unlinked supplemental records that did not match a patient in the database when the record was first loaded. However, the task can be used to rematch any type of record.

Records that meet the following criteria will be loaded into the workflow and rematched:

- The record is unlinked.
- The record is not the focus of a workflow task.
- The record's type, reportability, date loaded, and Import ID correspond to the options set in the task.

To rematch unlinked records:

1. Select **System > Tasks**.

2. Click **Rematch Unlinked Records**.

3. Select the **Record Type**. The task can only be run on one type of record at a time.

4. You have the option of rematching records with a specific value for **Reportability**. Leave this option blank to consider all records regardless of reportability.

5. You have the option of restricting the rematch to records **Loaded on or after** a specific date. Leave this option blank to consider all records regardless of date loaded.

6. If an **Import ID** is specified, all records in the import will be considered by the task (the date parameter will be ignored). Records in the import that are linked or in the workflow will not be considered.

7. You may enter text related to this task in the **Comment** field. The comment for the last Rematch Passive Follow-up Records task is stored in the database (utility_history table).

8. Click **Recalculate** to view the number of records that will be considered by the task. Determine whether this value is reasonable.

9. Click **Start**. The records will be loaded into the workflow and matched against the database. Passive follow-up data will be auto-consolidated, if possible. Match-consolidate, Consolidate FUP, Supplemental Match tasks may be created depending on the record type and match result.

Aborting a System Task

Requires system permission: *system_administration*

You may use the worklist to stop the execution of a system task. However, all changes made by the task may not be reversed. Some system tasks that update data perform the updates in batches. Updates that were made before you click Abort Task will not be reversed. You may use reports or query audit logs, the worklist, and other data to determine what was changed. The Patient Set Edits and Build tasks (Build CFOs, Build DCOs, Build SHOs) create a report. For those tasks, the report will still be created and will list the patient sets that were modified.

To abort a System Task:

1. For system tasks, the automatic task will be assigned to you and a link will be shown in the My Tasks section of your home page. Click the Automated task link for the appropriate task type to access the worklist.

2. Click the **Task ID**.

3. Click **Abort Task**.

4. Click **OK** to confirm.

System Administration Page

Requires system permission: *system_administration*

The System Administration page allows you to monitor system logs, review the registry-specific configuration settings and algorithms implemented in SEER*DMS, and access the version history.

To review information provided in the System Administration page:

1. Click **System** > **Administration**. The current server log will be shown when you first enter the page.

2. To view other logs and documents, select a topic from the **Name** drop-down list.

    a. Configuration – variable definitions in the registry-specific configuration files.

    b. Environment – environment variables for the SEER*DMS servers.

    c. Hibernate – technical information for the SEER*DMS development team. This page shows information concerning internal data structures and caching strategies. The Hibernate section may not be displayed in your registry's version of SEER*DMS.

    d. Logs – Boot and server logs. Use these logs to monitor system activity and to identify the cause of system errors.

    e. Memory – memory allocation, current usage, and memory pool status indicators.

f. System Files – the XML files containing the registry-specific implementation of the Edits, Importer, Lookups, Matching, Reports, Screening, Translator, and Workflow. Select a specific file from the Files drop-down list. These files are particularly useful for accessing and reviewing importer algorithms, matching criteria, and report SQL.

g. System Properties – Linux environment settings (a resource for the SEER*DMS development team)

h. Threads – list of concurrent system processes that are executing.