

Laurel Bridge Exodus User Manual



Laurel Bridge Software, Inc.
302-453-0222
www.laurelbridge.com

Document Version: 4.2.4
Last Saved: 2/24/2023 6:58:00 PM

Table of Contents

| | | |
|-------|---|----|
| 1 | What is Exodus? | 2 |
| 1.1 | Exodus Applications | 2 |
| 1.2 | Migration Basics | 2 |
| 1.2.1 | Workflow | 3 |
| 1.3 | Process Overview | 4 |
| 1.3.1 | Query the Source Archive | 4 |
| 1.3.2 | Migrate the Exams | 4 |
| 1.3.3 | Validate the Exams | 4 |
| 1.3.4 | Report Migration Status | 4 |
| 2 | Installation | 5 |
| 2.1 | Minimum System Specifications | 5 |
| 2.2 | Install Software Prerequisites | 5 |
| 2.3 | Main Exodus Install | 5 |
| 2.4 | Archive Configuration Prerequisites | 6 |
| 2.5 | SQL Install Examples | 6 |
| 2.5.1 | Installing SQL Server 2012 Express x64 with Tools | 6 |
| 2.5.2 | Installing SQL Server 2014 Express x64 with Tools | 7 |
| 3 | Getting Started | 8 |
| 3.1 | Overview of the Process | 8 |
| 3.2 | Migration Types | 8 |
| 3.3 | Exodus Application Configurations & License Options | 8 |
| 3.3.1 | Exodus Migration Controller | 8 |
| 3.3.2 | Exodus Assessment | 9 |
| 3.3.3 | License Mode: Charge By | 9 |
| 3.3.4 | License Mode: Segment Types | 9 |
| 3.4 | Installing a License | 10 |
| 3.4.1 | Activating a license over the internet | 11 |
| 3.4.2 | Installing a downloaded license file | 11 |
| 3.5 | Creating a Migration | 12 |
| 3.6 | Running a Migration | 12 |
| 4 | User Interface Details | 13 |
| 4.1 | Sessions / Segments Panel | 13 |
| 4.2 | Session Configuration Dialog | 14 |
| 4.2.1 | File-based Migration Settings on Network Tab | 15 |
| 4.3 | Segment Configuration Dialog | 16 |
| 4.4 | Exam Viewer | 17 |

- 4.5 Query Viewer 18
- 4.6 History Dialog..... 18
- 4.7 Containing Segments Dialog 18
- 4.8 Activity Monitor 19
- 4.9 Operation Progress Panel 19
- 4.10 Tool Bar Functions 19
- 4.11 Menu Bar Options..... 20
- 4.12 Audit Logging Configuration 21
- 5 Segments..... 21
 - 5.1 Overview 21
 - 5.2 Normal Segment 22
 - 5.3 Specific Segment..... 22
 - 5.4 List Segment..... 22
 - 5.5 Priors Segment..... 22
 - 5.6 Fixed Segment..... 23
- 6 Operations 25
 - 6.1 Query..... 25
 - 6.2 Move 25
 - 6.3 Validate 26
- 7 Filters..... 28
 - 7.1 Conditions 28
 - 7.2 Original Attribute Sequence 28
 - 7.3 Standard Actions..... 29
 - 7.4 Mapping List Action 31
 - 7.5 Composer Action Examples 36
 - 7.5.1 Working with DICOM sequences 37
 - 7.5.2 Modifying private tags 38
 - 7.6 De-Identify Filter Action..... 39
 - 7.6.1 DICOM Attribute Tags to Anonymize..... 39
 - 7.6.2 Miscellaneous Settings..... 40
 - 7.6.3 Pixel Regions to Anonymize 41
 - 7.6.4 One-Way Anonymization 44
 - 7.6.5 Additional Notes 44
 - 7.7 Re-Identify Filter Action 45
 - 7.7.1 Additional Notes 46
- Appendix A Exodus Privacy and Security Statement..... 47
- 1 Management of Private Data..... 47
 - 1.1 Types of PHI Maintained..... 47

- 1.2 Persistence of Private Data 47
- 1.3 Transmission of Private Data 48
- 1.4 Payment Card Industry (PCI) Data Security Standard 48
- 2 Security Capabilities 49
 - 2.1 Automatic Logoff..... 49
 - 2.2 Audit Controls 49
 - 2.3 User Authorization 49
 - 2.4 Security Configuration 49
 - 2.5 Security Updates 50
 - 2.6 De-Identification of PHI..... 50
 - 2.7 Backup and Restore 50
 - 2.8 Emergency Access..... 50
 - 2.9 Data Integrity and Authenticity 50
 - 2.10 Malware Protection 50
 - 2.11 Node Authentication..... 51
 - 2.12 Person Authentication 51
 - 2.13 Physical Locks..... 51
 - 2.14 Device Life Cycle Roadmap 51
 - 2.15 System and Application Hardening..... 51
 - 2.16 Security Guidance 52
 - 2.17 Data Storage Confidentiality 52
 - 2.18 Data Transmission Confidentiality 52
 - 2.19 Data Transmission Integrity 52
 - 2.20 Other Security Considerations 52
- 3 GDPR Notes..... 53
- Appendix B How to not migrate (i.e., drop) non-image DICOM objects..... 54
- Appendix C Creating a Priors Fetch Segment for Exodus Priors Fetcher 57
- Appendix D Exodus FAQ 58
 - 1. How do I know how many Exams I am allowed to move? 58
 - 2. What do the numbers at the bottom of the Main Dialog mean?..... 58
 - 3. What does the number under the Query Progress Bar mean if I do not have a Date Range Segment? 59
 - 4. Exodus seems to be running slow, what can I do? 59
 - 5. Should I Query and Move and Validate all at the same time? 60
 - 6. How do I configure Exodus to use Tag Filters for data being moved?..... 60
 - 7. How do I apply a filter to all images (datasets)?..... 61
 - 8. How do I validate a migration where new images are added to the destination? 62
 - 9. I did not migrate SR (Structured Report) objects; how can I perform validation? 62

- 10. After querying the Source, I do not see any exams in the Exam Viewer; what should I do? 65
- 11. Exodus was shut down for several days, will my Priors Segment catch-up?..... 65
- 12. Where can I find more information on Regular Expressions used by Exodus? 65
- 13. My low priority Segments do not get run; what is wrong with my configuration?..... 65
- 14. I want to use a Mapping List Action with private tags..... 65
- 15. How does Exodus respond to C-ECHO messages?..... 66
- 16. How do I know when to use SQL Express, Standard, or better? 66
- 17. How do I add custom prequalification segments? 66
- 18. I have upgraded to version 4.0.12 or newer and I cannot activate my license? 66
- 19. Help I'm stuck! 66
- Appendix E Source of Truth Discovery and Migration Process..... 68
 - 1 Overview of the Typical Migration Process 68
 - 1.1 Query the Source Archive 68
 - 1.2 Migrate the Exams 68
 - 1.3 Validate the Exams..... 68
 - 1.4 Report Migration Status..... 69
 - 2 Source of Truth (SOT) – Process..... 69
 - 2.1 Overview 69
 - 2.2 Technical Discussion 69
 - 2.3 PreprocessSot.exe Application 69
 - 2.4 DirectoryScanner.exe Application 70
 - 2.5 SourceOfTruthQRSCP.exe Application 70

Table of Figures

| | |
|--|----|
| Figure 1 - Typical Migration Workflow | 3 |
| Figure 2 - Migration Workflow with Data Filtering (Tag Morphing) | 3 |
| Figure 3 - Installer Welcome Screen | 5 |
| Figure 4 - Exodus About box & License Info | 10 |
| Figure 5 - Exodus License Activation Dialog..... | 11 |
| Figure 6 - Sessions / Segment Panel | 13 |
| Figure 7 - Sessions / Segment Panel Context Menu | 13 |
| Figure 8 - Session Context Menu | 13 |
| Figure 9 - Segment Context Menu | 14 |
| Figure 10 - Session Configuration Dialog | 14 |
| Figure 11 - Segment Configuration | 16 |
| Figure 12 - Exam Viewer | 17 |
| Figure 13 - Column Display Menu..... | 17 |
| Figure 14 - Exam Menu | 17 |
| Figure 15 - Exam Search Pane..... | 17 |
| Figure 16 - Exam History Dialog | 18 |
| Figure 17 - Exam Containing Segments Dialog | 18 |
| Figure 18 - Association Activity Monitor..... | 19 |
| Figure 19 - Operation Progress Panel | 19 |
| Figure 20 - Exodus Main Tool Bar | 19 |
| Figure 21 - Exodus Main Menu Bar Options | 20 |
| Figure 22 – Audit Logging Configuration | 21 |
| Figure 23 - Session Properties, Filters Dialog..... | 28 |
| Figure 24 - Mapping list action specification | 31 |
| Figure 25 - Creating a Mapping List Filter Action (3 parts shown)..... | 32 |
| Figure 26 - Example of a MapFilter.csv file | 34 |
| Figure 27 - Exodus Main Screen after attempting a migration..... | 34 |
| Figure 28 - Example 1 of the Query Client Tool | 35 |
| Figure 29 - Example 2 of the Query Client Tool | 36 |
| Figure 30 - De-Identify Filter Action window | 39 |
| Figure 31 - Negative Pixel Region Coordinate..... | 42 |
| Figure 32 - Pixel Regions to Anonymize Demo Form | 43 |
| Figure 33 - De-identify pixel region color picker..... | 43 |
| Figure 34 - Re-identify Filter Action Editor window..... | 45 |
| Figure 35 - Create a tag filter to Drop an instance | 54 |
| Figure 36 - Level 2 Validation Rule that Queries the Source PACS at the Image Level..... | 55 |
| Figure 37 - Validate Command Filter to Request SOP Class UID tag from Source PACS | 56 |
| Figure 38 - Validation Command Filter to drop Basic Text SR from responses from the Source PACS..... | 56 |
| Figure 39 - Meaning of numbers on the Main Exodus Dialog Window | 58 |
| Figure 40 - Locking the View of the Query Viewer and Exam Viewer | 59 |
| Figure 41 - Setting Exodus to use Tag Filtering | 60 |
| Figure 42 - Exodus AE-Title & Port..... | 60 |
| Figure 43 - Example: apply a filter to all images (datasets) | 61 |
| Figure 44 - Validating a migration where new images are added to the destination | 62 |
| Figure 45 - Create a Validation Rule to Compare Series Level to Study level..... | 63 |
| Figure 46 - Create a Filter to Drop SR Images | 63 |
| Figure 47 - Create a Validation Command Filter to insert Modality as a return tag | 64 |
| Figure 48 - Create a Validation Command Filter to Drop Responses containing SR in the modality tag ... | 64 |

1 What is Exodus?

Exodus is a collection of software applications that assist in the automation of DICOM migration tasks. Exodus applications focus on reliability, flexibility, and a simplified user experience.

Legacy archives often have features that present challenges for moving DICOM data, including:

- Merger of two or more archives
- Access to historical relevant priors
- Mismatched patient/study information
- Archive vendor proprietary issues
- Private DICOM tag handling
- Non-compliant/inconsistent DICOM data
- Unknown size of the job
- Uncertainty of completeness
- Inability to validate the data moved
- Excessive manual effort
- Inability to pre-fetch relevant priors
- Unresponsive support

Exodus allows the user to **automate the process** of collecting the information needed from a legacy archive to assess, plan, configure, perform, and validate a DICOM Migration. This collected information can be used to **assess exactly what data the user has** and help decide exactly what needs to be moved and what needs to be changed.

Using built-in validation and reporting systems, the user can **determine exactly what has moved** and what has not. Exodus ensures that exams are moved in a timely way and that they are available for use in their entirety on the new archive - all automatically.

From start to finish, the goal of Exodus is to provide a complete and transparent view of the issues related to moving DICOM studies, plus provide options to **automatically control and report the movement of the DICOM data** in a simple, high-level way; freeing the user to concentrate on other tasks.

1.1 Exodus Applications

Exodus comes in two basic application configurations:

Certain features described in this document may not be applicable depending on the application in use:

Exodus Migration Controller

Migrate DICOM studies from one archive to another—the complete archive migration solution.

Exodus Assessment

Assess the contents of the source archive prior to a migration and/or validate the results of a previous migration.

1.2 Migration Basics

A data migration is defined as the process of copying data from one location (the “source”) to another (the “destination”). While the source data is left intact, this process is often referred to as “moving” the data, since primary ownership of the data is normally transferred to the destination upon completion. A determination of which exams should be moved is made, and a list of exams to move is collected and

subsequently moved. To ensure that the exam was processed properly by the destination, a query may be issued to the destination to verify the existence and/or validity of the exams moved.

1.2.1 Workflow

DICOM migrations function through a set of third-party DICOM commands that allow a client to interrogate an archive for content and then selectively issue move requests to transfer that data to another system. In the simplest case, the third-party client (Exodus) never sees the actual data:




Figure 1 - Typical Migration Workflow

If the data needs to be changed in any way (also known as “**filtering**” or “**tag morphing**”) it must pass through a system that can do so. The main difference is that the move destination as far as the source is concerned is the system that will do the filtering. This filtering system is responsible for forwarding the image on to the final destination:




Figure 2 - Migration Workflow with Data Filtering (Tag Morphing)

1.3 Process Overview

Exodus performs the following main functions: assess, plan, perform, validate, and report for DICOM data migration activities for DICOM image archives/PACS.

A variation of this standard process is described in Appendix E Source of Truth Discovery and Migration Process.

Here are the standard steps involved in completing a DICOM data migration:

1.3.1 Query the Source Archive

During this initial phase, Exodus discovers (“learns” via DICOM C-Find requests) the contents of the source archive in preparation to analyze, plan, and perform the DICOM migration. There are two methods for querying the source PACS:

1. Using **DICOM Query/Retrieve (Q/R) services** typically employed by the Source Archive, or
2. Via a **file-import process** where Exodus can scan/traverse a file-store of DICOM datasets searching for DICOM exams.

The advantage of the Q/R method of discovery is that the Source Archive will interact with its internal database of study and patient information, and then during delivery of the exam to Exodus will apply appropriate database updates to the outbound DICOM study. For example, if a “John Doe” patient is scanned while in emergency room, then his images will be stored with the patient name of John Doe. Later, when the patient’s name is identified, the database in the Source Archive is updated to reflect patient’s actual name. Then when a DICOM Q/R is processed, the Source Archive will retrieve the images (still containing name of John Doe) but will deliver the study using the patient’s real name located in the internal database.

Regardless of the exam discovery process above, Exodus will build an internal database of study and patient information that will be used to control and perform the DICOM migration.

1.3.2 Migrate the Exams

Using the Exodus Exam Table generated in the query step above, Exodus will facilitate the migration of DICOM exams from Source Archive to Destination Archive. While doing so, Exodus can perform several other functions such as tag-morphing, transfer syntax (compression) modification, throughput throttling, and progress/failure notification.

1.3.3 Validate the Exams

Once the source PACS reports that a move has succeeded, Exodus will again query the Source Archive/PACS and the Destination Archive/PACS, and then compare the two responses to ensure all instances have in fact migrated.

1.3.4 Report Migration Status

Exodus provides built-in reporting mechanisms to report summary level, as well as detailed level, migration activities. Alternately, views of the Exodus exam table, which contains study and patient information, migration status (Queried, Moved, Validated), timestamps, etc., can be exported to Excel or CSV format for user-customizable reporting.

2 Installation

2.1 Minimum System Specifications

System configuration should be based on the scope and size of the planned migration and should be selected in consultation with LBS support staff.

The minimum system specifications for Exodus migrations are:

- Intel i7 or Xeon enterprise server processor, 16GB RAM, 512GB HD (minimum). If using a virtual machine, the virtual system performance must be equivalent to the above hardware.
- Windows 10, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019
- Microsoft SQL Server 2012 x64, SQL Server 2014 x64, or SQL Server 2016 x64.
The Express edition should be sufficient for smaller migrations (up to 500,000 exams). However, if using a virtual machine and SQL Server Express, each virtual processor socket must be configured with 4 or more virtual processor cores, for optimal performance.

2.2 Install Software Prerequisites

Laurel Bridge Exodus utilizes several software components that must be installed for it to work properly.

The software prerequisites are:

- Microsoft .NET Framework 4.8 (minimum)
- Microsoft SQL Management Studio
- Microsoft Visual C++ Redistributable for Visual Studio 2017 x64 version (Note: This will be automatically installed if needed)

Microsoft .Net Framework and Microsoft SQL Server should be installed before running the `Exodus.msi` installer.

2.3 Main Exodus Install

After installing the prerequisites, the Exodus application installer (`Exodus.msi`) should be run. For machines with an older Exodus version installed, this installer will upgrade your previous install appropriately, while maintaining your configuration settings.

Once the application installation is complete, you will also have to install the application license as described below in Section 3.4, Installing a License.

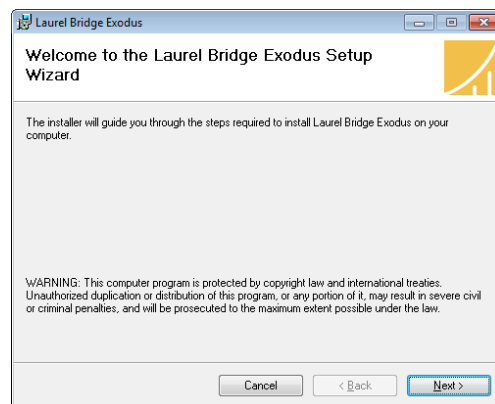


Figure 3 - Installer Welcome Screen

2.4 Archive Configuration Prerequisites

Both the Source Archive/PACS and the Destination Archive/PACS must be configured to communicate with each other and with Exodus. This may include configuring a new AE Title, hostname/port configuration on the Source and/or Destination PACS. The typical configuration changes are summarized below:

- For any Migration:
 - 1) The Source PACS should recognize Exodus as a DICOM Query/Retrieve SCU device.
 - 2) The Destination PACS must recognize Exodus as a DICOM Query/Retrieve SCU device.
- For Migrations where exams **will be routed or filtered** by Exodus:
 - 1) The Source PACS must be configured with Exodus as a DICOM C-MOVE destination (C-STORE SCP).
 - 2) The Destination PACS must be configured with Exodus as a DICOM C-STORE client (C-STORE SCU).
- For Migrations where exams **will not be routed or filtered** by Exodus:
 - 1) The Source PACS must be configured with the Destination PACS as a DICOM C-MOVE destination (C-STORE SCP).
 - 2) The Destination PACS must be configured with the Source PACS as a DICOM C-STORE client (C-STORE SCU).
- Exodus itself can be configured to use any AE Title and any available port may be used as its DICOM port. Section 4.2, Session Configuration Dialog, of the manual describes the configuration of Exodus itself.

2.5 SQL Install Examples

Install the SQL version appropriate to your environment as noted in Section 2.2 above.

2.5.1 Installing SQL Server 2012 Express x64 with Tools

These are instructions for installing SQL Server Express in its most basic configuration for use by Exodus. These instructions are valid for Windows 10, Windows Server 2012, and Windows Server 2016. The installation procedure may differ if an older version of SQL Server is installed, if the full version of SQL Server is preferred, or if SQL Server authentication mode must be enabled.

1. Log in to Windows as a user with administrative privileges
2. Run the [SQL Server 2012 Express x64 with tools \(SQLEXPRTW_x64_ENU.exe\)](#) installer
3. On the [Setup](#) screen select [New installation or add features to an existing installation](#)
4. On the [License Terms](#) screen Accept the license, click the [Next>](#) button
5. On the [Product Updates](#) screen choose whether to check for updates based on your corporate policy.
6. On the [Feature Selection](#) screen make sure all the checkboxes are checked for all of the [Instance Features](#). Make sure that the [Management Tools](#) checkboxes are checked, click the [Next>](#) Button
7. On the [Instance Configuration](#) screen, the defaults should be correct. The named instance should be [SQLExpress](#). Allow it to install in the default location, which should be C:\Program Files\Microsoft SQL Server\

8. On the **Server Configuration** screen, the defaults should be fine for the **Service Accounts** tab and the **Collation** tab defaults.
9. On the **Database Engine Configuration** screen on the **Account Provisioning** tab, select **Windows Authentication Mode**.
The Current user (who must have Administrative Privileges) should be in the list under **Specify SQL Server Administrators**. If it is not, click the button to **Add Current User**.
Leave the defaults on the other two tabs.
10. Installation should complete in several minutes

2.5.2 Installing SQL Server 2014 Express x64 with Tools

These are instructions for installing SQL Server Express in its most basic configuration for use by Exodus. These instructions are valid for Windows 8.1, Windows 10, Windows Server 2012, and Windows Server 2016. The installation procedure may differ if an older version of SQL Server is installed, if the full version of SQL Server is preferred, or if SQL Server authentication mode must be enabled.

1. Log in to Windows as a user with administrative privileges
2. Run the **SQL Server 2014 Express x64 with tools (SQLEXPRTW_x64_ENU.exe)** installer
3. On the **Setup** screen select **New installation or add features to an existing installation**
4. On the **License Terms** screen Accept the license, click the **Next>** button
5. On the **Microsoft Update** screen choose whether to check for updates based on your policy.
6. On the **Feature Selection** screen make sure all of the checkboxes are checked for all of the **Instance Features**. Make sure that the **Management Tools** checkboxes are checked, click the **Next>** Button
7. On the **Instance Configuration** screen, the defaults should be correct.
The named instance should be **SQLExpress**. Allow it to install in the default location, which should be C:\Program Files\Microsoft SQL Server\
8. On the **Server Configuration** screen, the defaults should be fine for the **Service Accounts** tab and the **Collation** tab defaults
9. On the **Database Engine Configuration** screen on the **Account Provisioning** tab, select **Windows Authentication Mode**.
The Current user (who must have Administrative Privileges) should be in the list under **Specify SQL Server Administrators**. If it is not, click the button to **Add Current User**.
Leave the defaults on the other two tabs.
10. Installation should complete in several minutes

3 Getting Started

3.1 Overview of the Process

Exodus manages migrations by creating a database known as a **session** that contains all the configuration information as well as all the state and tag information for every exam it knows about. Groups of exams are viewed and operated upon by creating configurable views known as **segments** within a session.

A session represents a separate database that stores configuration information about a migration. Information about the source, destination, and validation archives is included, as well as settings for each of these operations (e.g. timeouts, schedules, delays, etc.). In addition to this information, the list and state of every query and exam is contained within a session which represents the work that has been done for a given migration. While multiple sessions can be configured, typically only one is needed for a given migration.

A segment provides the definition of what to query and represents a logical grouping of exams to migrate based on those queries. There are multiple types of segments available allowing for various methods by which to gather listings of exams to work on. A given migration can be broken up into multiple segments or left as one segment representing the entirety of the migration.

Recommendation: For DICOM Migrations of approximately 500,000 studies or more, Laurel Bridge recommends, if possible, that the migration be divided into multiple sessions within Exodus, each limited to no more than 500,000 studies. Performance issues have been noted for sessions larger than 500,000 studies.

3.2 Migration Types

Exodus supports two types of migrations: **File** and **Archive**. The type of migration defines where the source data is located. For file-based migrations, data is collected recursively from a directory; for archive-based migrations data is housed in a DICOM-compatible archive. Both migrations are processed very similarly and have: query phases where a list of exams is collected; move phases where the chosen exams are moved; and a validate phase where the moved exams are verified. In fact, after indexing the information from the raw DICOM image files from disk into a built-in archive, file-based migrations function exactly like an archive-based one (i.e. – both are communicating with a source archive – one is simply built-in).

3.3 Exodus Application Configurations & License Options

Exodus comes in two basic configurations: **Exodus Migration Controller** and **Exodus Assessment**. Information on which version you are licensed for can be found by clicking [Help > About Laurel Bridge Exodus](#). Certain features described in this document may not be applicable depending on the application version you are using:

3.3.1 Exodus Migration Controller

This application configuration is the complete archive migration solution, allowing the user to migrate DICOM studies from one archive to another:

- Provides functionality to assess, plan, perform, validate, and report activities associated with performing a DICOM migration from a single source archive to a single destination archive at a given site.
- Includes advanced dataset filtering, tag morphing, and data element processing in real-time as data migrates, e.g., standardize dataset contents based on a master patient index or other site-specific standards, rules, or workflows.
- Includes all functions available in Exodus Assessment and Exodus Priors Fetcher (described below).
- Provides functionality to monitor a modality **Worklist** (or pseudo **Worklist**) to identify scheduled exams and move patient prior studies from a single DICOM source archive to a single DICOM destination.

3.3.2 Exodus Assessment

This application configuration allows the user to assess the contents of the source archive prior to a migration or validate the destination data following a migration:

- Provides functionality to query source DICOM archive to build Exodus database of study demographics, including ability to validate a previously performed DICOM migration.
- Does NOT provide migration functionality, i.e., does not move studies from a source to a destination archive.
- Provides ability to run pre-qualification algorithms against the imported DICOM demographic data.
Pre-qualification criteria have been developed based upon typical problems encountered during a DICOM Migration, e.g., blank accession numbers, blank patient names, duplicate MRNs for same patient, etc.
Results of this pre-qualification, or custom pre-qualification filters developed by customer, provides valuable information for developing migration strategy, including problem identification and mitigation.
- Provides ability to validate a previous migration by comparing of source and destination using Exodus Validation rules and levels; includes the ability to generate comprehensive reports off what actually did migrate and what did not.

3.3.3 License Mode: Charge By

Exodus comes in two modes: **Charge by Exam** and **Charge by Move**. Information on which mode you are licensed for can be found by clicking [Help > About Laurel Bridge Exodus](#).

- **Charge by Exam:** Licensed by number of exams (i.e. - a given exam can move multiple times and will only count against the total once).
- **Charge by Move:** Licensed by number of moves (i.e. – a given exam can move multiple times but will be counted against the total once for each move).

3.3.4 License Mode: Segment Types

Exodus can be configured to support any or all the available segment types. Information on which segment types you are licensed for can be found by clicking [Help > About Laurel Bridge Exodus](#).

3.4 Installing a License

The first step after launching Exodus is to install a license. When the **About Laurel Bridge Exodus** dialog appears, there are two options to install a license. Information on the status of your license, including when it expires, can be found on this dialog available at any time by clicking **Help > About Laurel Bridge Exodus**.

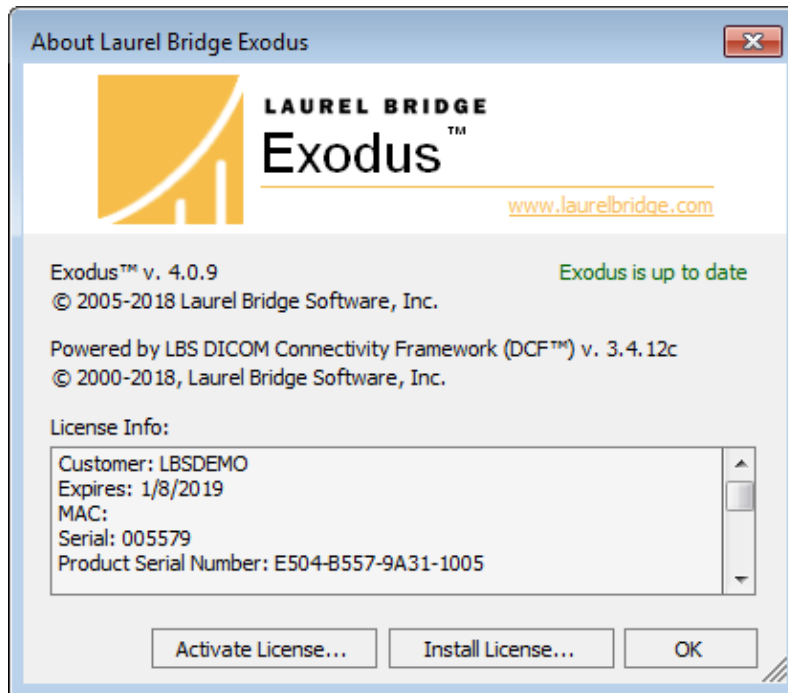


Figure 4 - Exodus About box & License Info

3.4.1 Activating a license over the internet

Clicking the **Activate License...** button will open the **License Activation** form. You will need to provide a **16-digit Product Serial Number** to activate your license. All fields on the form must be filled out. If you have Internet access and press the **Lookup** button, then the form will be prepopulated with the currently known license information, this is useful for upgrades.

The screenshot shows a 'License Activation' dialog box with the following fields and values:

- Product Section:**
 - Name: Exodus
 - Ver.: 4.0.9
 - Product S/N: E504-1234-5678-0123
 - ARC: 064C-31F0-FA2E-7237
 - MAC: (empty)
- User Section:**
 - Site: My Site Name
 - Contact: My Name
 - Email: My E-mail Address
 - Host Name: My Hostname
- Maintenance Section:**
 - Contact: (empty)
 - Email: (empty)
 - Phone: (empty)

Buttons at the bottom: 'Lookup' and 'Activate'.

Figure 5 - Exodus License Activation Dialog

Special case: If the Exodus server does not have internet access you will need to record the **16-digit Activation Request Code (ARC)** from the **License Activation** form. You can then proceed to the manual web activation form at:

https://www.laurelbridge.com/product_activation.php

and enter your **16-digit Product Serial Number** and the **16-digit Activation Request Code (ARC)**. After completing the steps on the manual web activation form, you can download a fully activated Exodus license and copy it to the system where you are installing Exodus. You can then use the **Install License...** button on the **About Exodus** dialog to install your downloaded and activated license.


3.4.2 Installing a downloaded license file


Clicking the **Install License...** button will open a file searcher dialog, allowing you to browse to a fully-activated license file. You will need to use this button to install a license if a fully-activated license was provided to you via email (for example, a MAC-based license), if you downloaded the license file from www.laurelbridge.com, or if you performed a web activation as described in the "Special Case" above.

3.5 Creating a Migration

To begin using Exodus to control a migration, you need to create a **session** to define the source and destination archives and at least one **segment** to define what queries you want to issue and the resulting exams you want to work on.

To create a session, click **File > New Session** (or right-click in the **Sessions / Segments** panel and select **New Session**). Enter the desired session name, taking care to avoid the following reserved names: Beacon, Compass, Compass2, Lighthouse, Navigator, NavigatorNet, Waypoint, and `__INTERNAL__`.

After you enter the desired session name a configuration dialog will appear that contains network settings for all the archives involved in the migration, alert settings for email notifications, and advanced settings for timeouts, rates, schedules, delays, try counts, etc. Context help is available by clicking the  icon present next to each field in the dialog.

To create a segment, select the new session and right-click **New Segment** in the **Sessions / Segments** panel. After you enter the desired segment name a configuration dialog will appear that contains settings for the type, priority, and type-specific selection information. Context help is available by clicking the  icon present next to each field in the dialog.

3.6 Running a Migration

After creating a session and at least one segment you can begin a migration. To do so, check mark the desired segments in the **Sessions / Segments** panel, check mark the appropriate operations at the bottom of the main dialog (**Query**, **Move**, and **Validate**), and click the **Run** button on the toolbar. The progress bars on the bottom of the screen will reflect how far along each of the selected operations are. Once there are no more operations to complete, the migration will automatically stop.


While the operations can be run concurrently, they are still sequential. At least one successful query must complete before there can be any exams eligible to move. Similarly, at least one successful move must complete before there can be any exams eligible to validate.

4 User Interface Details

To configure Exodus to control a migration you must first create a **session** to define the source and destination archives and then you must create at least one **segment** to define what queries to issue and, as a consequence, the resulting exams to work on.

4.1 Sessions / Segments Panel


The **Sessions / Segments** panel (shown in Figure 6) is available on the left-hand side of the main window. It is a two-level tree control that contains and shows the available sessions (top level) and each session's available segments (bottom level).



Only one session can be active at a time and is represented on the display by bold text (Figure 6 shows **example session 1** as the currently active session and **example session 2** as an inactive session). Only the active session displays its segments, each of which can be checked to select it for viewing and running.

Figure 6 - Sessions / Segment Panel

When the migration is running, the checkboxes are replaced by icons representing the state of the segment. A green icon means the segment is in progress. A white icon means the segment is not currently loaded. A red icon means some unexpected error has occurred (dynamic segments only). A black icon means that the segment is currently loaded and complete.

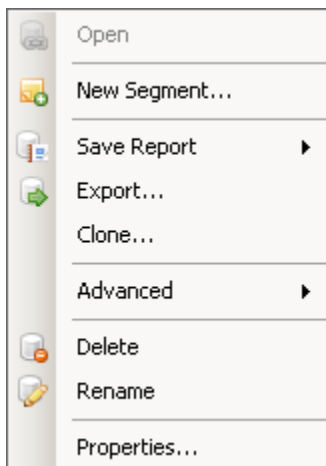


The session panel itself has a context menu (shown in Figure) available by right-clicking any blank space within in to access two options:

Figure 7 - Sessions / Segment Panel Context Menu

- **New Session:** Creates a new session
- **Import Session:** Imports a session from a file

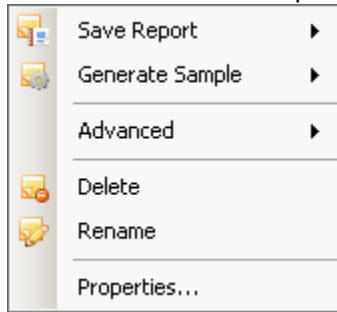
Each session entry also has a context menu (shown in Figure) available by right-click. Note: some context menu options are disabled depending on whether the particular session is active or not and whether multiple sessions are selected.



- **Open:** Activates a session
- **New Segment:** Creates a new segment for a session
- **Save Report:** Generates an import report (file-based migrations only) or a progress report about all the contained segments and their composite exams contained within a session
- **Export:** Exports a session, typically for backup
- **Clone:** Creates a new session with the same settings and segments as the selected session
- **Advanced:** Perform advanced options including optimizing the database, resetting query/move/validate errors, and resetting all query/move/validate progress (to redo all work)
- **Delete:** Deletes a session
- **Rename:** Renames a session
- **Properties:** Displays the session configuration dialog

Figure 8 - Session Context Menu


Each segment also has a context menu (shown in Figure 9) available by right-click. Note: certain options are disabled when multiple segments are selected.



- **Save Report:** Generates a progress report about the exams contained within a segment
- **Generate Sample:** Creates a sample of representative data from the given segment (selected by date and modality)
- **Advanced:** Performs advanced options including resetting query/move/validate errors and resetting all query/move/validate progress (e.g., to redo all work)
- **Delete:** Deletes a segment
- **Rename:** Renames a segment
- **Properties:** Displays the segment configuration dialog

Figure 9 - Segment Context Menu

4.2 Session Configuration Dialog

The session configuration dialog (shown in Figure 10 -) contains all the connectivity and alert information about a migration. Details about each entry field are available by clicking the  icon next to a particular field. The contents of the various menu tabs are described below.

- **Network:** The network tab contains the connectivity information for the source, destination, and validation archives. The network tab shown in the figure is specific to archive migrations. The validation archive is the system that responds to query requests after an exam is moved and is typically the same physical machine as the destination archive (unless routing occurs, for example). Click the **advanced** button to specify validation archive information if it is different from the destination archive. Support for TLS (by version) can be configured for connections both to/from the archives.
- **General:** The general tab contains the migration type (archive or file-based), migration level (study or series), reset time/action, and the alert information including email server configuration, alert recipient, and alert type information. The migration type and level cannot be changed after the session is created. Email alerts are sent out to the specified email addresses depending on what events occur and what alert types are registered for. When a stop trigger (certain number of successive failed query, move, or validate operations) is reached alerts are sent out to email addresses registered for **Error Alerts**. When a new day is reached daily reports are sent out to email addresses registered for **Daily Reports**.
- **Query:** The query tab contains specific information about the query operations including timeouts, try counts, delays, partitioning modes, and user defined tags. User defined tags are non-standard tags that aren't normally requested during a query. Any DICOM tag can be

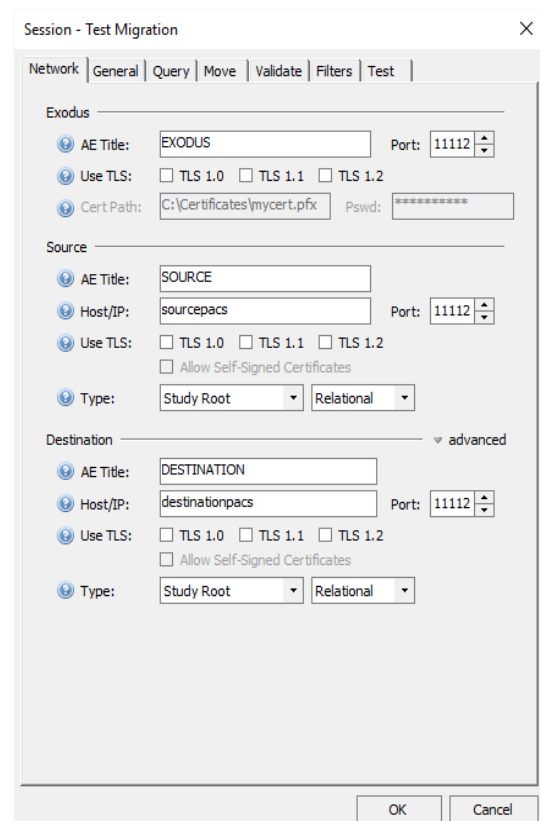


Figure 10 - Session Configuration Dialog


queried, but there is no guarantee that an archive will return data for a particular tag. See 6.1 for more details on the query operation including user defined tags.


- **Move:** The move tab contains specific information about the query operations including timeouts, move order, pre-checks, try counts, delays, and schedule. The move schedule is defined on a per hour basis with three basic modes: **Normal**, **Off**, and **Rated**. During **Rated** mode, the configurable rate (given as a percentage) is used. Right-clicking the move schedule will bring up a presets menu allowing you to select from popular schedule configurations. See 6.2 for more details on the move operation.
- **Validate:** The validate tab contains specific information about the query operations including timeouts, try counts, delays, and rules. Each validation rule consists of a set of options including the tags to use to find the exam on both the source and destination archives, level from which to obtain image count or list information for comparison, and comparison logic to use to compare the collected data (e.g. – equal to, greater than, etc.). See 6.3 for more details on the validate operation.
- **Filters:** The filters tab contains the configuration window used to create image (dataset) filters that are used to change image information during the migration. If image filtering is enabled all images are routed through Exodus (via C-STORE from the source) before being forwarded on the destination (via C-STORE to the destination). The filtering mechanism can be configured to allow data transfer as uncompressed, compressed, or both (via forcing the source to negotiate a particular transfer syntax).
- **Test:** Selecting the test tab allows the user to initiate a series of connectivity and compatibility tests for the entered information. It is recommended that any errors revealed by testing be addressed as they are likely the result of misconfiguration that will impede the migration or limit the information returned from the given archives.


4.2.1 File-based Migration Settings on Network Tab

When a file-based migration is selected for a new session on the general tab, some file-specific Source settings are available. Other settings on the network tab are the same as for an archive migration.

Source

Path: 

Extensions: Include files with no extension 


Normalize: Coerce Patient IDs to UPPERCASE 

- **Path:** The path to the files to index. Clicking the file icon on the right enables a directory path to be selected.
- **Extensions:** A comma separated list of file extensions (without the dots) used to select the files to import. This is used for importing files with specific, known file extensions. All characters, including spaces and dots, that are included in the list of extensions will be used exactly as entered. Because extensions are separated by commas, there is no way to include a comma in an extension. Here are some examples:
 - “dcm” matches files ending with “.dcm”
 - “dcm,txt” matches files ending with “.dcm” or “.txt”
 - “ xyz” matches files ending with “. xyz” (including spaces)
 - Leaving the field empty matches files that have any extension

- **Include files with no extension:** Checking this option (default) will also include files with no extension. This is used when importing files that do not have dots anywhere in the filename. To list **only** files that **lack** an extension, check **Include files without extensions** and enter an extension that does not exist among any of the files to import (e.g., “mybogusext”).
- **Coerce Patient IDs to UPPERCASE:** Checking this option causes the Patient IDs of imported files to be normalized in the Exodus database by coercing them to all UPPERCASE. DICOM data sent to the destination is migrated in its original state, so Patient IDs will not be normalized.

For each of these file-based Source settings, once the file import has completed successfully their values cannot be changed.

4.3 Segment Configuration Dialog

The segment configuration dialog (shown in Figure 11 -) contains the information that defines a segment and the queries and exams the segment contains. Each segment type has a different set of fields available that apply to that particular segment type. Details about each entry field are available by clicking the  icon next to a particular field.

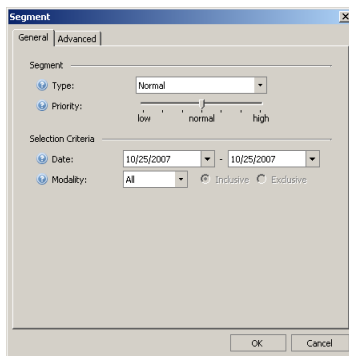


Figure 11 - Segment Configuration

- **General:** The general tab contains the segment type and general selection criteria for the segment. The available types include **Normal**, **Specific**, **List**, **Priors**, and **Fixed**. See section 5 for more information on each segment type.
 - ❖ **Normal:** Provides a date range to query, one day at a time, backwards through time
 - ❖ **Specific:** Provides specific patient data to be queried in a single query
 - ❖ **List:** Provides a list of exams to query and move, provided by a file or command
 - ❖ **Priors:** Provides a set of patient matching requirements, used after querying for a list of patients from a modality worklist server or command-line output to then use to match against the exams queried from a source archive
 - ❖ **Fixed:** Used as a simple fixed list of studies, usually created by right-clicking an exam or set of exams and selecting “Add To Segment”.
- **Advanced:** The advanced tab contains specialized selection criteria used to refine the segment definition explicitly. Also available is a move schedule that can be optionally enabled which will override the global session move schedule. The move schedule is defined on a per hour basis with three basic modes: **Normal**, **Off**, and **Rated**. During **Rated** mode, the configurable rate (given as a percentage) is used. Right-clicking the move schedule will bring up a presets menu allowing you to select from popular schedule configurations.

4.4 Exam Viewer

The exam viewer (shown in Figure 12) is available on the right-hand side of the main window (if activated by clicking **View > Exams**) and contains a list of all the exams that are contained within any loaded segment (i.e., segments with a checkmark next to them).

| Uid | Exam Date | Patient Name | Patient Id |
|---|------------|------------------|---------------|
| 1.2.840.114089.1.0.1.3229814446.126529... | 06/23/2005 | BLACK^PATRICIA | PID000005FA22 |
| 1.2.840.114089.1.0.1.3229814446.126529... | 03/26/2001 | HOWARD^NYDIA | PID000007A68C |
| 1.2.840.114089.1.0.1.3229814446.126529... | 08/09/2006 | WOODS^JERALD | PID0000052CBK |
| 1.2.840.114089.1.0.1.3229814446.126529... | 05/09/2001 | STRECKER^KENNETH | PID000007B592 |
| 1.2.840.114089.1.0.1.3229814446.126529... | 07/19/2000 | SKELDON^MARIA | PID0000047DCI |
| 1.2.840.114089.1.0.1.3229814446.126529... | 06/25/2001 | SKELDON^MARIA | PID0000047DCI |
| 1.2.840.114089.1.0.1.3229814446.126529... | 09/16/2005 | SCOTT^THERESE | PID000002F70C |

Figure 12 - Exam Viewer

Columns can be reordered by dragging the column headers left or right within the viewer. Columns can be sorted on by clicking on a particular column header. The direction of the sort (i.e., ascending or descending) is toggled on each subsequent click. Columns can be hidden by right-clicking on any column header and using the menu that appears (shown by Figure 13 -) to toggle the visibility of a particular column.



Figure 13 - Column Display Menu

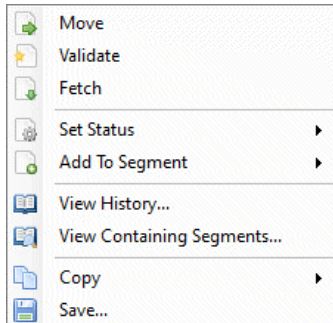


Figure 14 - Exam Menu

Each exam listed in the viewer has a context menu (shown in Figure 14) available by right click. Note: certain options are disabled if configuration information about specific archives (e.g., source, destination, validate) is missing.

- **Move (or Query):** Move selected exams (or Query with selected queries) after setting status to **Waiting**
- **Validate:** Validate selected exams after setting to **Moved**
- **Fetch:** Move selected exams (or queries) to local machine
- **Set Status:** Override the selected exams (or queries) status. Setting an exam to a status of **Validated** will set the status to be **Validated (0)** to indicate that it was validated manually.
- **Add to Segment:** Add the selected exams to a new or existing list segment
- **View History:** View the move/validate history for selected exams (or queries)
- **View Containing Segments:** View the containing segments for selected exams (or queries)
- **Copy:** Copy the selected rows/cell text to the clipboard
- **Save:** Save information for the selected exams or queries to a file

The exam search pane (shown in Figure 15 -) is available above the exam viewer (if activated by clicking **View > Exams Search** or the small magnifying glass icon in the viewer title bar) and allows the list of exams displayed in the exam viewer to be limited by a set of match requirements. This can be used to find a logical subset of all the loaded exams without defining a new segment. Every column in the exam viewer can be used in the exam filter control with a number of match modes (e.g. **Equal To**, **Greater Than**, etc.). The exam filter control only limits the exams displayed. It does not, however, change the loaded segment definitions. Thus, if an exam is contained within a loaded segment it is still eligible to be moved and/or validated regardless of whether it is displayed in the viewer.



Figure 15 - Exam Search Pane

To increase performance (especially when doing a large number of queries) the exam viewer can be locked (i.e., updates to the viewer are disabled) by clicking the small lock icon in the viewer title bar.

4.5 Query Viewer

The query viewer is available on the right-hand side of the main window (if activated by clicking **View > Queries**) and displays a list of all the queries that are contained within any loaded segment (i.e., segments with a checkmark next to them). Column display, context menu options, searching, and locking are similar to the exam viewer dialog options. See section 4.4 for more details.

4.6 History Dialog

The history dialog (shown in Figure 16) is available by right-clicking on any number of exams or queries and selecting **View History**. The entire history for each item is listed with the time and result along with any comments (typically error comments).

History entries for each of the selected items can be stepped thru using the navigation arrows at the bottom or with the Home, Left, Right, and End keys.

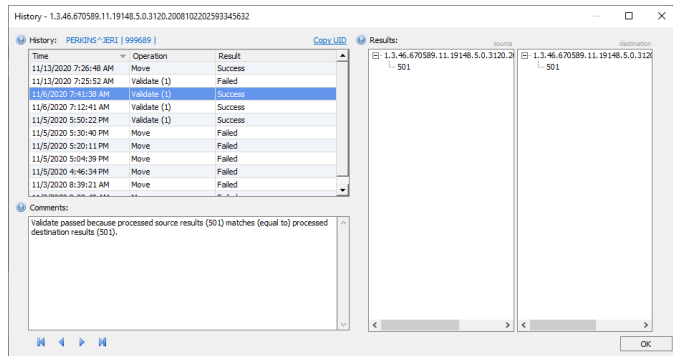


Figure 16 - Exam History Dialog

Any of the exam or query values presented in **blue text** can be copied to the system clipboard by clicking on them.

For validate operations each rule for a given score attempted is listed as a separate operation. For each try the information that was returned from the source and validation archives is listed to the right in the **Results** panel. A short explanation as to why a validation operation failed or passed is displayed in the **Comments** box.

4.7 Containing Segments Dialog

The containing segments dialog (shown in Figure 17) is available by right-clicking on any number of exams or queries and selecting **View Containing Segments**. The list of segments in which each selected exam or query is contained is displayed in the **Containing Segments** box. Entries for each of the selected items can be stepped thru using the navigation arrows at the bottom or with the Home, Left, Right, and End keys.

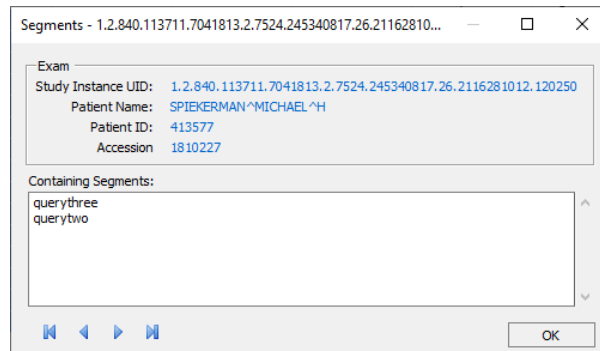
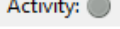
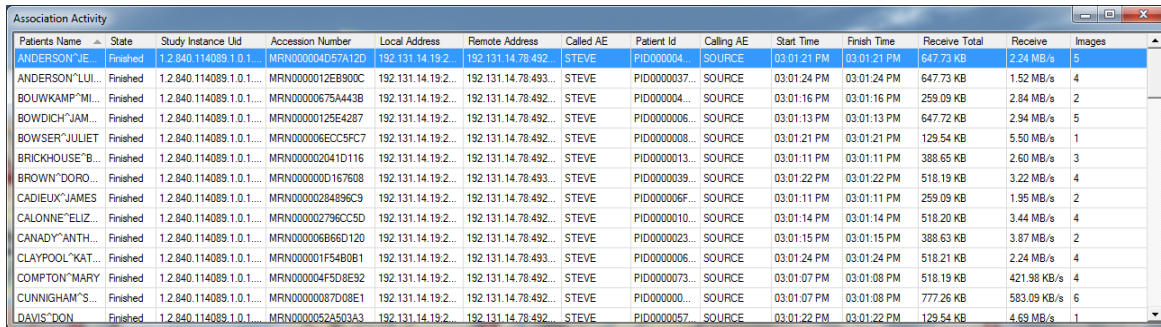


Figure 17 - Exam Containing Segments Dialog

Any of the exam or query values presented in **blue text** can be copied to the system clipboard by clicking on them.

4.8 Activity Monitor

The Activity Monitor can be opened by clicking the  **Activity Button** on the **Tool Bar**. The **Activity Monitor** displays DICOM Associations that have been routed by Exodus. **Association Information** is not persisted forever. Association information is persisted based on the configuration parameters set on the Exodus **Options Form**. You can right-click on the **Activity Monitor** at any time to open a pop-up dialog that will allow you to remove all **Association Information**.



| Patients Name | State | Study Instance Uid | Accession Number | Local Address | Remote Address | Called AE | Patient Id | Calling AE | Start Time | Finish Time | Receive Total | Receive | Images |
|-----------------|----------|--------------------------|------------------|--------------------|---------------------|-----------|---------------|------------|-------------|-------------|---------------|-------------|--------|
| ANDERSON"JE... | Finished | 1.2.840.114089.1.0.1.... | MRN000004057A12D | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID000004... | SOURCE | 03:01:21 PM | 03:01:21 PM | 647.73 KB | 2.24 MB/s | 5 |
| ANDERSON"LUI... | Finished | 1.2.840.114089.1.0.1.... | MRN0000012EB900C | 192.131.14.19.2... | 192.131.14.78.493.. | STEVE | PID0000037... | SOURCE | 03:01:24 PM | 03:01:24 PM | 647.73 KB | 1.52 MB/s | 4 |
| BOUWKAMP"MI... | Finished | 1.2.840.114089.1.0.1.... | MRN00000675A443B | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID000004... | SOURCE | 03:01:16 PM | 03:01:16 PM | 259.09 KB | 2.84 MB/s | 2 |
| BOWDICH"JAM... | Finished | 1.2.840.114089.1.0.1.... | MRN00000125E4287 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000006... | SOURCE | 03:01:13 PM | 03:01:13 PM | 647.72 KB | 2.94 MB/s | 5 |
| BOWSER"JULIET | Finished | 1.2.840.114089.1.0.1.... | MRN000006ECC5FC7 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000008... | SOURCE | 03:01:21 PM | 03:01:21 PM | 129.54 KB | 5.50 MB/s | 1 |
| BRICKHOUSE"B... | Finished | 1.2.840.114089.1.0.1.... | MRN000002041D116 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000013... | SOURCE | 03:01:11 PM | 03:01:11 PM | 388.65 KB | 2.60 MB/s | 3 |
| BROWN"DORO... | Finished | 1.2.840.114089.1.0.1.... | MRN0000000167608 | 192.131.14.19.2... | 192.131.14.78.493.. | STEVE | PID0000039... | SOURCE | 03:01:22 PM | 03:01:22 PM | 518.19 KB | 3.22 MB/s | 4 |
| CADIEUX"JAMES | Finished | 1.2.840.114089.1.0.1.... | MRN00000284896C9 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID000006F... | SOURCE | 03:01:11 PM | 03:01:11 PM | 259.09 KB | 1.95 MB/s | 2 |
| CALONNE"ELIZ... | Finished | 1.2.840.114089.1.0.1.... | MRN000002796CC5D | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000010... | SOURCE | 03:01:14 PM | 03:01:14 PM | 518.20 KB | 3.44 MB/s | 4 |
| CANADY"ANTH... | Finished | 1.2.840.114089.1.0.1.... | MRN000006866D120 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000023... | SOURCE | 03:01:15 PM | 03:01:15 PM | 388.63 KB | 3.87 MB/s | 2 |
| CLAYPOOL"KAT... | Finished | 1.2.840.114089.1.0.1.... | MRN000001F5480B1 | 192.131.14.19.2... | 192.131.14.78.493.. | STEVE | PID0000006... | SOURCE | 03:01:24 PM | 03:01:24 PM | 518.21 KB | 2.24 MB/s | 4 |
| COMPTON"MARY | Finished | 1.2.840.114089.1.0.1.... | MRN000004F508E92 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000073... | SOURCE | 03:01:07 PM | 03:01:08 PM | 518.19 KB | 421.98 KB/s | 4 |
| CUNNINGHAM"S... | Finished | 1.2.840.114089.1.0.1.... | MRN00000087D08E1 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID000000... | SOURCE | 03:01:07 PM | 03:01:08 PM | 777.26 KB | 583.09 KB/s | 6 |
| DAVIS"DON | Finished | 1.2.840.114089.1.0.1.... | MRN0000052A503A3 | 192.131.14.19.2... | 192.131.14.78.492.. | STEVE | PID0000057... | SOURCE | 03:01:22 PM | 03:01:22 PM | 129.54 KB | 4.69 MB/s | 1 |

Figure 18 - Association Activity Monitor

4.9 Operation Progress Panel




Figure 19 - Operation Progress Panel

The operation progress panel (shown in Figure 19 -) is available at the bottom of the main window and is used to display the individual progress of the **Query**, **Move**, and **Validate** operations for the loaded segments. Each section of the operation panel has a checkbox to control whether that operation will run when a migration is run. When running, these checkboxes are replaced by a *spinning* animation to indicate that a specific instance of a given operation (e.g., a single move) is occurring or a *pulsing* animation to indicate that a given operation is waiting for more work to do. Progress is represented through a progress bar (blue for success and red for failure) as well as through numeric statistics below each progress bar. The estimated amount of time remaining for a given operation is also displayed.

4.10 Tool Bar Functions






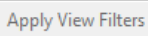


Figure 20 - Exodus Main Tool Bar

The tool bar (shown in Figure 20) is shown at the top of the Exodus main control window and contains multiple buttons for controlling, configuring, and viewing a migration.

- **Run:** Starts a migration involving the selected operations on the loaded segments
- **Stop:** Stops a migration (prompts if operation currently in progress)

-  **[Session]**: Displays the session context menu
-  **[Segment]**: Displays the segment context menu
-  **[Record Log]**: Starts full-detail logging to user-defined file for simplified debugging purposes
- **Activity:**  Flashes if there are exams moving through the Image Routing component. Click to open the Activity Monitor.
-  **Apply View Filters** If clicked, then View Filters will be applied to the Exam Viewer and Query Viewer

In addition to these controls, an exam viewer row selector is available on the far right which displays the current selected row and the total number of rows available for display in the viewer.

4.11 Menu Bar Options




Figure 21 - Exodus Main Menu Bar Options

The menu bar is available at the top of the main window. Menu drop-down options are:

- **File > New Session:** Creates a new session
- **File > Import Session:** Imports a session from a file
- **File > Record Log:** Starts full-detail logging to user-defined file
- **File > Exit:** Exits the program (after stopping migration)

- **View > Queries:** Toggles the display of the query viewer
- **View > Exams:** Toggles the display of the exam viewer
- **View > Queries Search:** Toggles the activation/display of the query search pane
- **View > Exams Search:** Toggles the activation/display of the exam search pane
- **View > Sessions / Segments:** Toggles the activation/display of the session/segments viewer

- **Tools > Log Viewer:** Displays a log viewer dialog with display filtering options
- **Tools > Current Application Log:** Opens the current log file for Exodus
- **Tools > Application Log Directory:** Opens a file explore to the current log directory for Exodus
- **Tools > Apply View Filters F5:** Applies the set filters to the current log exam view
- **Tools > Activity Monitor:** Opens a pop-up to the display the current DICOM association history
- **Tools > Query Client:** Displays a query dialog that performs C-FIND requests
- **Tools > Options:** Displays an options dialog to set program options (e.g. – logging)

- **Help > Getting Started:** Displays a guide on how get started with the program
- **Help > User Manual:** Opens this user manual
- **Help > Release Notes:** Opens the Release Notes version and change history
- **Help > About Laurel Bridge Exodus:** Displays an about dialog containing program and license information; provides access to License Activation and Installation options

4.12 Audit Logging Configuration

To configure Exodus audit logging capabilities, select the **Audit Log Configuration** edit link in the **Tools > Options** dialog. The Audit Logging dialog will be displayed as shown to the right. This dialog allows the Exodus audit logging capabilities (an implementation of DICOM PS 3.15 Appendix A.5 “Audit Trail Using Syslog” functionality) to be enabled or disabled. Exodus audit log messages are typically sent to a remote (secured) syslog server, but copies of all audit log messages will also be stored locally in the Exodus logs.

If audit logging is enabled, the syslog server **Host/IP** address, **Transport** protocol, and **Port** can be configured. (Note that if the **Host/IP** address is left blank, audit logging will still be enabled locally.) The **Site ID** field allows configuration of the audit enterprise site ID, which is used to uniquely identify this Exodus instance within the enterprise.

If the **UDP** transport option is selected, the port defaults to the well-known port 514. If the **TCP** transport option is selected, the port defaults to the well-known port 601. If **TLS** transport is selected, the port defaults to the well-known port 6514. In addition, the supported version(s) of TLS can be configured, as well as whether or not to ignore certificate name mismatch errors and whether or not to allow self-signed certificates. Per RFC 5424, use of the **TLS** transport option is strongly recommended unless the underlying network is secure.

Finally, the types of messages to be audit logged can be configured. Enabling **Application Start/Stop** messages will log a message each time the Exodus application is started or stopped. Enabling **Software Configuration** messages will log a message each time an Exodus configuration (global, session, or segment) is changed or the Exodus migration is started/stopped. Enabling **User/Security Alerts** will log a message whenever a user performs a restricted function, such as creating/renaming/deleting sessions, creating/renaming/deleting segments, exporting sessions, performing manual exam operations, and generating manual reports. And finally, enabling **Queries/Retrieves** will log a message each time a query or retrieve (move) is performed. Caution should be used when enabling this option, as it could generate a very large number of syslog messages.

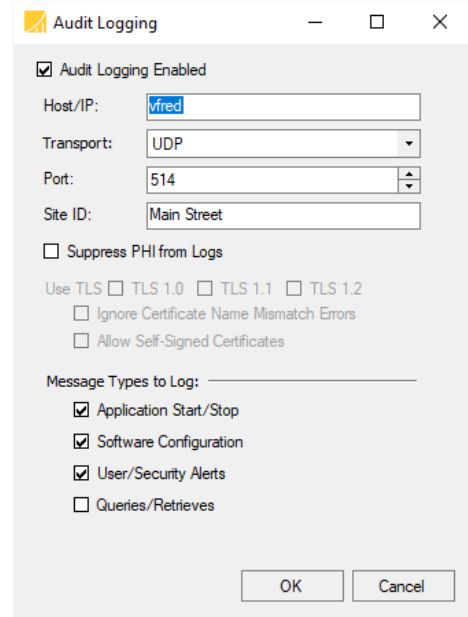


Figure 22 – Audit Logging Configuration

5 Segments

5.1 Overview

A segment provides the definition of what to query and represents a logical grouping of exams to migrate based on those queries. There are multiple types of segments available allowing for a variety of ways to identify and gather collections of exams. The set of queries done to gather exams may collect a larger list of exams into the database than will be shown in the segment display. In other words, every exam returned by a given query won't necessarily be part of the segment that initiated the query. This is because of the limited query functionality available in DICOM and especially in legacy archives. A given migration (session) can be broken up into multiple segments or left as one segment representing the entirety of the migration.

All segment types have segment-specific criteria as well as shared criteria including a **Name**, **Priority**, **Custom SQL Selection String**, and **Move Schedule Override** (except in **Assessment** versions of Exodus).

Note that each session can contain no more than 100 segments (at which point the **New Segment** menu item will be disabled).

5.2 Normal Segment

Normal segments are the most typical segment type for a historical (date-based) migration. This segment type defines a **Date** range to query: one day at a time with the most recent date queried first. A **Modality** selector is also available with the option for **Inclusive** (i.e., that modality is included in list of modalities) or **Exclusive** (i.e., that modality is not included in list of modalities) logic.

5.3 Specific Segment

Specific segments allow for selective information to be gathered from a single query. To limit the query scope a selective piece of information that is to be queried for must be included; typically, these are: **UID**, **Accession #**, **Patient Id**, and/or **Patient Name**. In addition, **Date** and **Modality** options are available to further restrict the selection criteria.

5.4 List Segment

List segments allow for a list of exam UIDs to be provided to Exodus. Each UID in the list will be queried to gather additional information about the exam. Lists can be provided to Exodus as either a **File** or as the output of a **Command**. This list can be repopulated on a configurable **Interval**.

5.5 Priors Segment

Priors segments allow a list of relevant prior exams to be gathered from a particular **Date** range. Exams are considered to be priors when they match based on a configurable set of criteria. The list of exams used to trigger a search for priors may come from a list of upcoming exams gathered from a **Worklist** server or from the standard output of a **Command** (see the format below). The set of match tags that can be used include **Patient Name**, **Patient Id**, **Patient Dob***, **Patient Sex***, **Modality**, and **Body Part*** (* indicates a tag that is available when it has been added to the session as a user-defined tag). The matching modes between the upcoming exam list and potential prior exams include **Exact** (matches values that match exactly), **Smart** (matches values in a criteria specific manner to avoid minor inconsistencies common to that specific tag), **Map** (matches values after translating upcoming exam values through a mapping given by a file), and **Custom** (matches values using a custom SQL string). Prior exams are no longer considered priors after the upcoming exam they match has been off of the upcoming exam list, gathered at a given **Interval**, for more than the time given by the **Fall Off** time. Each new exam can match a given number of match criteria specified by **Max Matches**. The reset time and action (**Charge by Move** mode only) allows for exams to be reset to the **Waiting** state (with **Reset** action) or to be deleted (with the **Delete** state) after a configurable amount of time for use with transient/temporary destinations or large capacity sources.

Smart priors matching follows these guidelines for comparing upcoming exam values with potential prior exam values (values are not case sensitive):

- **Patient Name:** Matches when either name (after being stripped of spaces and carets) is a subset of the other (e.g. – SMITH^BOB^J matches SMITH^BOB)
- **Patient Id:** Matches by ignoring leading zeros (e.g. – 0001234 matches 1234)
- **Patient Dob:** Matches when dates differ by one day or less (e.g. – 19900101 matches 19900102)
- **Patient Sex:** Matches when first letter matches (e.g. – Male matches M)
- **Modality:** Matches when upcoming exam value in an item in the list of values in the prior exam (e.g. – MR matches MR\CT)
- **Body Part:** Matches when upcoming exam value in an item in the list of values in the prior exam (e.g. – CHEST matches CHEST\HEAD)

Map priors matching functions the same for all tags. Matches occur when the upcoming exam value is contained within a list of *keys* in a map file and the associated *value* for that key is contained within the list of values in the prior exam. For example, if the upcoming exam has a value for **Patient Id** of 123456 and the map file has a mapping for 123456 to 56789 then this will match any potential prior exam that has a **Patient Id** of 56789 and *not* one that has a value of 123456. As another example, if the upcoming exam has a value for **Body Part** of HEAD and the map file has a mapping for HEAD to BRAIN and HEAD to HEAD then this will match any potential prior exam that has a **Body Part** of HEAD or HEAD\CHEST or BRAIN or BRAIN\ARM, etc.

The format of the standard output to be provided when using the **Command** mode to supply the list of exams follows an XML format, structured as follows:

```
<Exams>
<Exam
PatientName="<name>"
PatientId="<id>"
PatientDob="<yyyymmdd>"
PatientSex="<sex>"
Modality="<modality>"
BodyPart="<body part>" />
    . . .
<Exams>
```

Notes:

- Attributes are not required.
- The order is not significant.
- Newlines aren't significant.

Custom priors matching uses a custom SQL 'where' clause. For example, if you only wanted to match Exams where the modality is MG you would provide a string like this `Exam.Modality LIKE '%MG%'`

5.6 Fixed Segment

Fixed segments allow for a simple user-defined list of studies that are added through the user interface (as opposed to a file of UIDs, for example). These segments are usually created and managed by right-clicking on an exam or set of exams and selecting "Add To Segment" and "Remove From Segment". A

fixed segment might be used to manually and pre-emptively move selected exams for a particular patient, for instance.

6 Operations

6.1 Query

Every segment definition results in a set of queries that are performed against the source archive one at a time. Each query returns a set of exams that match the given query. To protect against potential DICOM query limitations, some queries are purposefully made more generic than might be expected. The result of this is that not all the exams that return necessarily match the exact segment definition. However, regardless of the precision of the query, the resulting exams are post-processed to mark only the exams that match the segment definition. The end result is that each high-level segment definition contains the appropriate, intelligent queries which will, in turn, provide the appropriate exam entries to the Exodus data base.

Normally, the set of queries is defined at segment creation/modification time. For dynamic segments like **Priors** and **List**, the query list changes periodically. Thus, dynamic segments will never finish as far as the application is concerned. Human assessment is required to determine when the dynamic segment is no longer needed. For example, the **Priors** segment can fetch prior exams for any patient on the worklist server. Since new patients can be added at any time, the segment is never truly complete. However, once all the exams on the source archive have been migrated, the segment is effectively complete. This determination is not automatically made, but the human assessment can be aided by examining the completeness of a secondary segment that contains a list of all the exams on the source archive.

Queries are tried a configurable number of times (until success) with a configurable timeout. Once a query is marked as failed it must be manually reset (available in the segment and session context menus). The delay in between queries is also configurable. In addition to the standard tags displayed in the exam viewer, any other tag, known as a user-defined tag, can be requested during all queries. These settings are found in the session configuration dialog.

When migrating at the study level any user defined tag can have its query level overridden to the series level forcing a single query at the series level for each study returned in the base query. The overridden tag will have its information coalesced from the information at the series level and grouped together by backslashes. So, for example, if you wanted to fetch Body Part Examined, but you knew it was only available at the series level. You would add a user defined tag for it (in the session configuration dialog) and override it by right-clicking the word "Study" next to it and then selecting "Series." In this example, the final returned study entries, that are stored in the internal database, would contain the Body Part Examined values from the series level.

6.2 Move

Each segment is made up of exams that match the given segment definition. Each exam contained in a segment is eligible for move provided its status is **Waiting**.

Each move is tried a configurable number of times (until success) with a configurable timeout. Moves are skipped (and marked as **Moved**) if the move pre-check is configured and the score passes (the pre-check process is identical to the validation process and can be thought of as pre-move validation where a certain score must be met to skip the move). The number of simultaneous moves is also configurable and may be limited due to configuration settings on the source archive. The move schedule that defines the delay in between moves and when moves should not take place is also configurable on a per hour

basis. Delays between moves are available in **Rated** mode and are a function of the percentage of the duration of the most recent move. These settings are found in the session configuration dialog.

Exams with failed move status can be reset individually by the exam viewer context menu or at the segment/session level through the segment and session context menus. This is typically done to cause the exam move to be retried.

6.3 Validate

Each exam contained in a segment is eligible for validation provided its status is **Moved** (when using the Exodus Assessment application, the necessary status is **Waiting**). The purpose of validation is to ensure that exams were successfully received and processed by the destination archive. The validation archive can be, and generally is, the same archive as the destination unless the destination is a router or gateway.

Validation consists of applying a set of rules, with a given rank, where each rule is processed (starting with rank 1) until a successful match is found. The resulting validation rank is the highest (i.e., lowest numbered) rank for which a match occurred. Thus, more stringent requirements should be used for higher ranks.

The first step in validation is to locate each exam or set of exams on the source and destination archives. Each rule consists of a configurable set of tags specified by the **Tags** list that will be used as the match tags in the queries to the source and destination. The matching values used are the values collected previously during the initial query phase of the migration. Depending on which tags are used and the content of each archive, one or more exams may be returned from each archive. In the ideal situation, validation would have one rule, matching solely on the exam UID, to return one result from each archive. However, since UIDs can potentially be changed during a migration, either intentionally or otherwise, this is not always possible. Choosing a different set of match tags can provide an alternate way of locating exams as well as a way to ensure that certain tags were not changed. However, if one or more of those tag values are changed during migration that exam will not be located on the destination. Thus, care should be taken to choose the appropriate match tags that will ensure that exams will be located correctly.

The next step is to determine from which level data will be obtained; this is designated by the **Level** selection. The options include **Study** which is a count of the number of images in a particular study (Number of Study Related Instances), **Series** which is a count of the number of images in each of the series in a particular study (Number of Series Related Instances), and **Image** which is a list of all the images in a particular study (SOP Instance UID). This decision is generally determined by the capabilities of the archives involved, as some systems do not support certain options. **Image** level validation is stricter as it requires that not only the image counts match, but that the exact list of images is the same.

The next step is to determine how the data collected is compared; this is designated by the **Comparison** selection. When retrieving count information (**Study** or **Series**), the options for comparison are numerically based (e.g. – **Equal To**, **Greater Than**, **Less Than or Equal To**, etc.) or list based (e.g. – **Equal To**, **Contains**, etc.) when retrieving list information (**Image** level). Note, that if either the source or the destination is set to return count information, then numerical based comparisons are used, as the list data is converted to count data for comparison (e.g. – a list of 5 images will be converted to “5”). This allows for flexibility in terms of the comparison such that, for example, any exam that has at least as

many images for a particular study on the validation archive as on the source archive can be considered a match. The options should be read as:

“results from source archive are <comparison option> the results from the validation archive”.

The last step is to determine how duplicates and multiple responses are processed, if allowed, as designated by the **Duplicates** and **Multiples** selections. When retrieving count information (e.g. – **Study Image Count**), the options are numerical based (e.g. – **Sum**, **Max**, **Min**, etc.) and list based (e.g. – **Unique**) when retrieving list information. Duplicate responses for count-based selections (**Study** or **Series**) occur when more than a single DICOM C-FIND response for a given UID is returned. For list-based selection (**Image**), duplicates occur when a duplicate image UID is returned. Duplicates are symptomatic of some DICOM incompatibility. Multiple responses occur when more than single exam UID is returned and would typically only be allowed when validation is expected to find multiple exams on either the source or the destination (e.g. – two exams are intentionally merged during the migration). These settings represent advanced capabilities and should only be used for specific cases. As such, the default is to **Disallow** these types of responses, resulting in a failed validation step.

7 Filters

Filters allow image (or dataset) header elements to be changed as Exodus receives them from the Source archive. Filters can also be used to modify DICOM C-FIND and C-MOVE messages as they are sent/ received to/from the Source and Destination archives.

A new Filter can be created by selecting the **Filters** tab on the **Session Properties** Dialog (Figure 23). The special C-FIND and C-MOVE filters can be created from the **Query**, **Move** and **Validate** tabs.

To add a dataset filter, press the green plus button located under the Filters list. This will create a new Filter with the name “New Filter n” where “n” is the next available number, starting at 1. Filter names are customizable and can be modified at any time by clicking on the name. The list

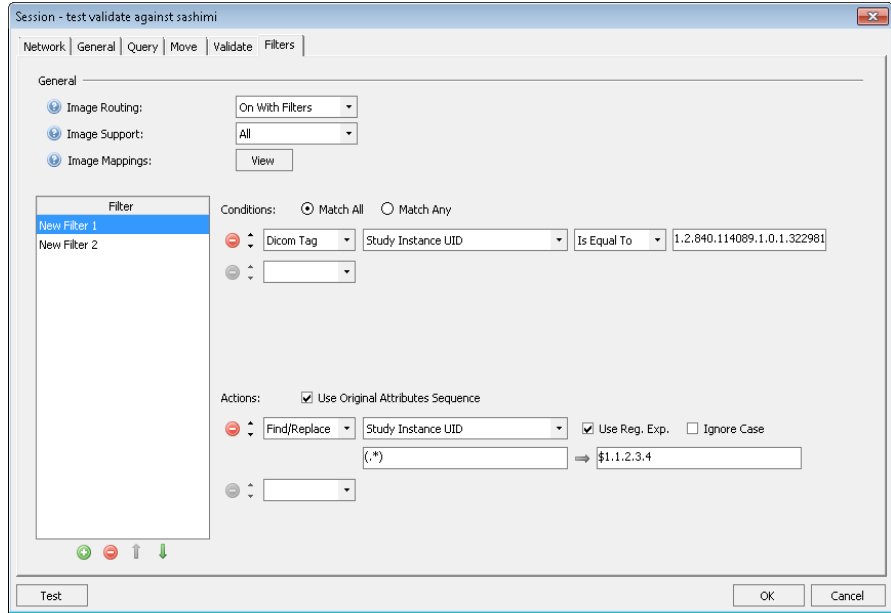



Figure 23 - Session Properties, Filters Dialog

of Filters may be reordered by selecting a Filter and then pressing the up or down arrows under the list display box. Filters are applied in order, top to bottom, so *ordering is important*.

Once a Filter place holder has been created the next step is to configure its Conditions and Actions.

7.1 Conditions

A Filter’s **Conditions** determine whether or not its Actions will be applied to the image being processed. A Filter may have multiple conditions; therefore, two options are provided that all Conditions must match (“Match All”) or any Conditions can match in order (“Match Any”) for the Filter to be applied. The other configuration is to have no Conditions, thereby causing the filter to be always applied with no conditional matching.



For each Condition select the **DICOM Tag** condition type (additional types may be added in the future). Configure the rest of the Condition based upon the desired test. Conditions may be removed by pressing the red minus sign located to the left of each condition and may be sorted via the up and down arrow buttons.

7.2 Original Attribute Sequence

The original tag value of a filtered tag can optionally be stored in the DICOM original attribute sequence by checking the **Create Original Attribute Sequence** checkbox. If left unchecked, the original values modified by the filter will not be recorded in the original attribute sequence.



7.3 Standard Actions

If the Conditions testing is satisfied, then any Actions defined by a filter will be applied to the image / dataset.

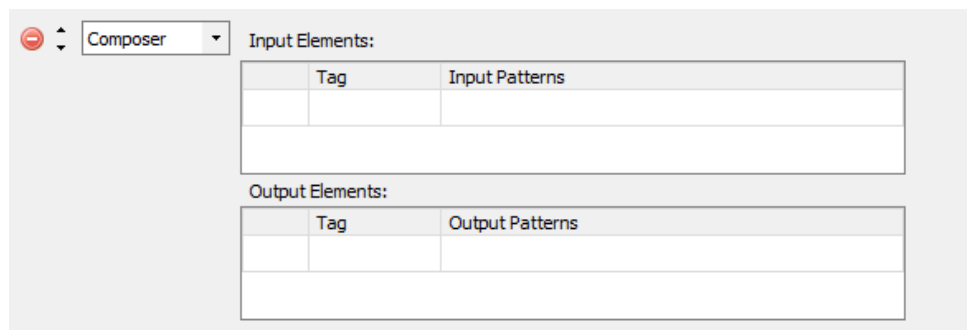
- The **Change Case** action changes the case of the selected tag to either upper-case or lower-case.



- The **Copy** action copies the tag specified in the From field to the tag specified in the To field.



- The **Composer** action supports the merging and splitting of elements. It can also be used to manipulate DICOM sequence elements with ease. It uses regular expressions to parse values from DICOM tags and combine the values into other DICOM tags. For example, it can take parts from two different tags and combine them to make a new value in a third tag. Regular expressions are specified to parse each input tag and substitution patterns determine how the output tags are constructed.



In the “Input Elements” table, specify the tags and how the regular expressions should parse each value into groups – the regular expressions go into the “Input Patterns” column. The groups that are produced are specified in the Output Patterns column of the “Output Elements” table. The first match in the first input element is referred to as “\${1.1}”, the second match in the first input element is “\${1.2}”, and the third match in the second input element would be “\${2.3}”, and so on.

The Output Elements specify what parts of the input elements to combine and how to combine them; note that the parts can be used multiple times and also combine them with plain text.

- The **Execute** action executes a piece of specialty custom code that is stored in a local file and that adheres to a particular coding interface. For example:

```
public class DateComparisonFilter : IFilterAction
{
    public void ApplyAction(CFGGroup config, DicomSessionSettings dss,
        RelevantTagMarker tagMarker, ref DicomDataSet dds)
    {
        // custom code
    }
}
```

Note: Contact Laurel Bridge Software for more information on creating custom Execute filters.



- The **Find/Replace** action changes the value of the selected tag via hardcoded values or .Net style regex values.



- The **Insert/Overwrite** action creates the specified Tag (if it didn't already exist) or overwrites the specified tag (if it did already exist) with the specified Value.



- The **Move** action moves the tag from its original location to its final location. Once moved, the tag will not be present in its original location. Note that this is not the same as the copy action.



- The **New UID** action replaces the Unique Identification (UID) value of the specified UID type tag in the Tag field with a new, auto-generated UID value.



- The **Pad** action pads the specified tag with the specified character. The **Left** and **Right** radio buttons specify whether to pad the left or right side of the value. The new total length of the tag value is specified by the **Length** field. If the length of the original value is equal to or greater than the Length field, the value will not be changed.



- The **Prepend/Append** action prepends or appends the specified text to the specified tag's value.



- The **Regex Substitution** action performs a regular expression match on the specified tag's value using the given **Pattern**. It then replaces the value with the **Replace** string which may contain regex variables from the match. The **Test** box can be filled with an example tag value and clicking the green arrow processes the test data.



- The **Remove** action removes the specified tag. If the **Only if Empty** checkbox is checked, the tag will be removed only if the specified tag's value is empty.



- The **Trim** action removes the specified characters from the left, right, or both sides of the specified tag's value.



- The **Drop Image** action drops any images (datasets) that match the conditions and then returns either a successful C-STORE Response message or a failure C-STORE response message to the Source archive.



7.4 Mapping List Action

Overview:

The **Mapping List** action allows the user to define and specify an external file that lists match tags and replace tags. When creating a Mapping List action in Exodus, the user specifies the format of the mapping list file to be used. The example below shows that the mapping list file will contain three comma-separated fields, in this case: an accession number, a patient ID, and a patient's name. If the match tags on any given row match the current dataset, the values of the replace tags are substituted into the dataset. Below is an example of a mapping list file that matches the format specified by the mapping list action below:

Mapping list file contents:

```
PIKR0004, PAT124, SMITH^JOSEPH  
ABC123, PAT45, MINER^STEVEN
```

Mapping list action specification:

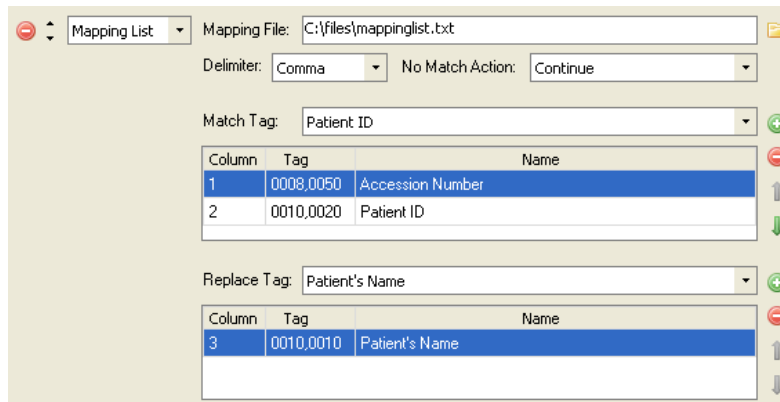


Figure 24 - Mapping list action specification

Example: Creating a Mapping List filter

Create your Mapping List Filter under the **Filters** tab of the **Session Properties Dialog**:

The **Actions** section has limited screen space and you have to scroll to see it in pieces. You can increase the size of that area of the dialog box by hovering slightly above the **Actions** label and getting a handle to that portion of the screen. Examples of scrolling the Actions are shown in the three screenshots below.

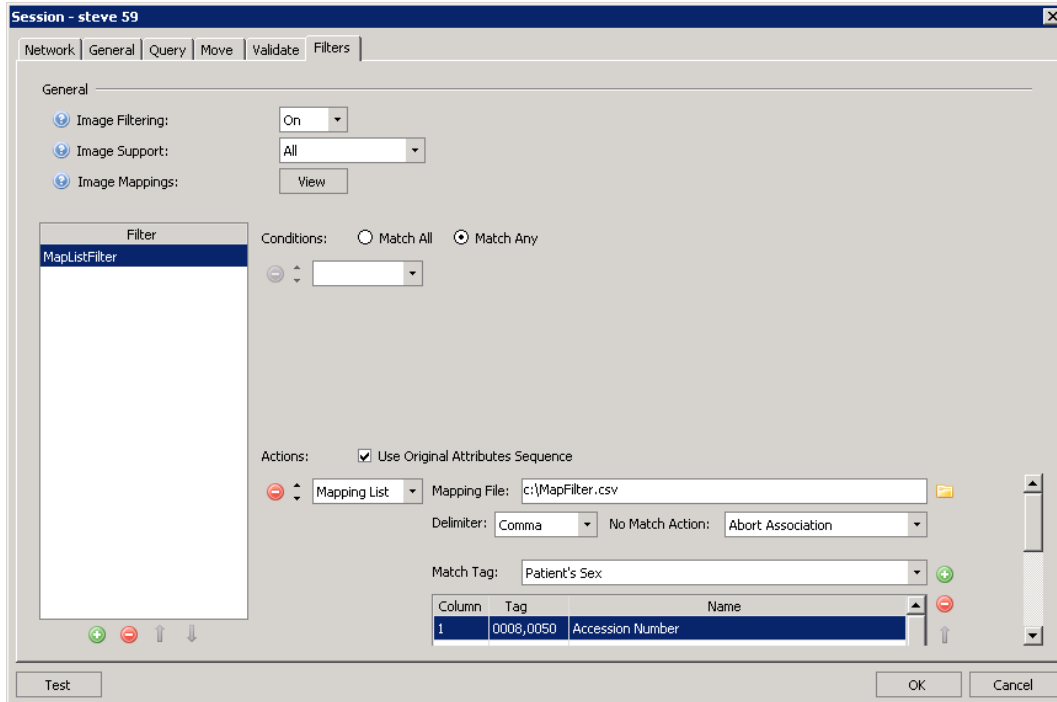
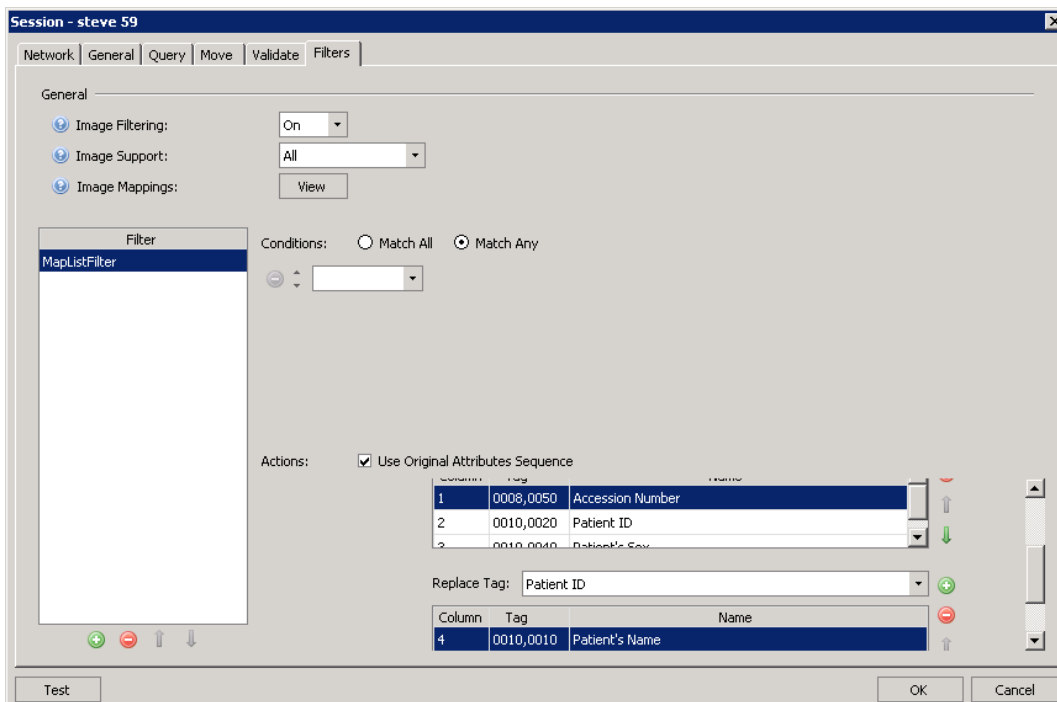
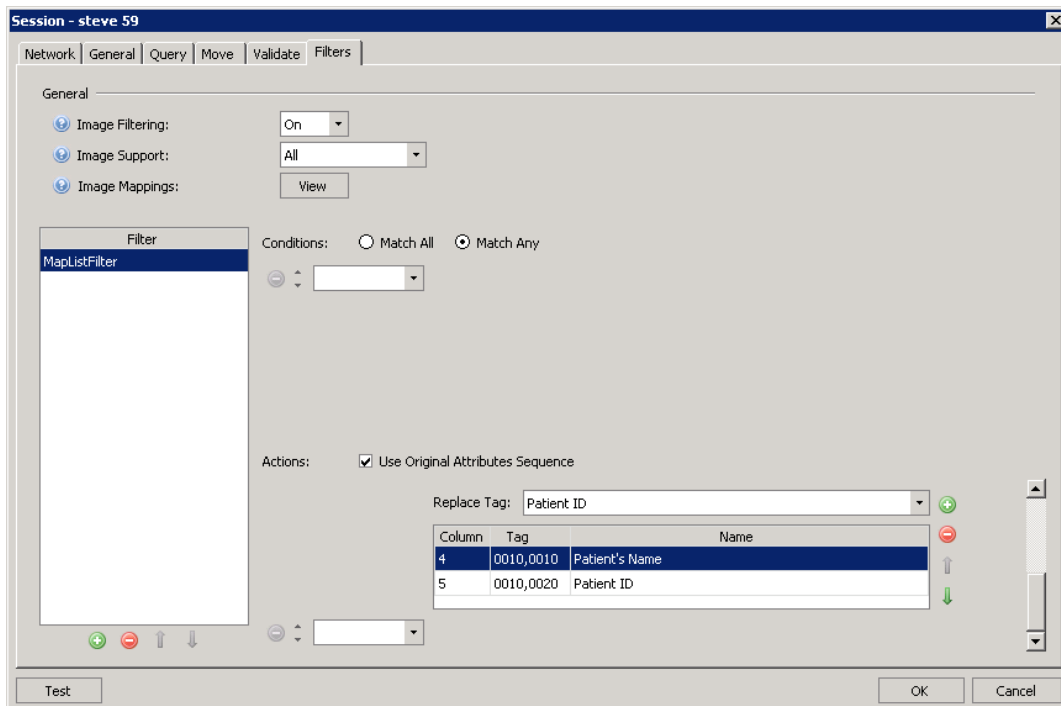


Figure 25 - Creating a Mapping List Filter Action (3 parts shown)





If you look at all three of the previous example screen shots in succession you'll see selected a mapping list filter that uses a comma delimited file called c:\MapFilter.csv and that should contain three match tags (Accession number, Patient ID and Patient's Sex were chosen in this example, although any tags may be chosen); and two replacement tags (Patient's Name and Patient ID were chosen in this example). Note that in screen 1 the "No Match" action is set to "Abort Association", this means that only datasets that match an entry will be moved, others will not be moved. If "Abort Association" is chosen and an exam contains an instance for which no match is found, then the association for that instance is aborted. If "Continue" is chosen, the instance will be sent on to the destination without a mapping being applied. Exodus issues Move requests to the source PACS and has no control over the association boundaries chosen by the source to fulfill those requests. Exodus will however continue to issue Move requests for subsequent Exams in your migration.

To the left of each tag on the screen is the **Column** number – which indicates which column in the file contains that data element.

Here's an example of what a MapFilter.csv file might look like for the sample definition:

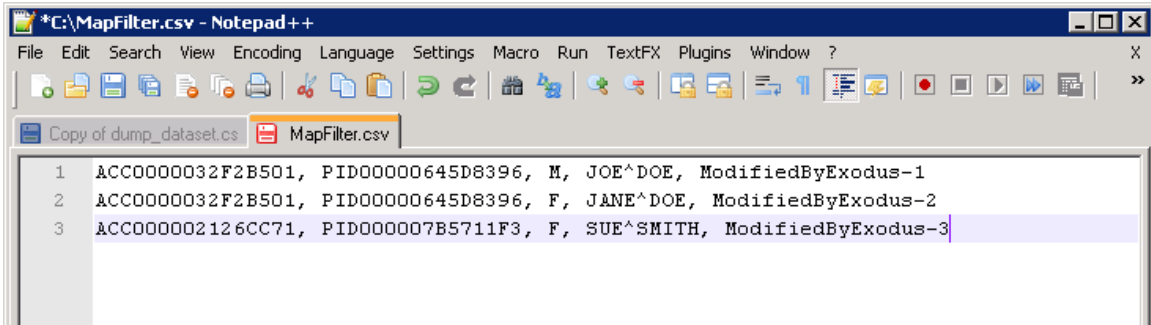


Figure 26 - Example of a MapFilter.csv file

Please note that this is a contrived example; in practice, a mapping list file like this should not ever be created—the first two lines show two different patients with the same Patient ID. This example is used to make a point about matching rules below.

Here’s the Exodus Main Screen after attempting a migration. Since the MapFilter.csv was so small, only the first two exams visible in the exam viewer below partly matched the mapping list criteria.

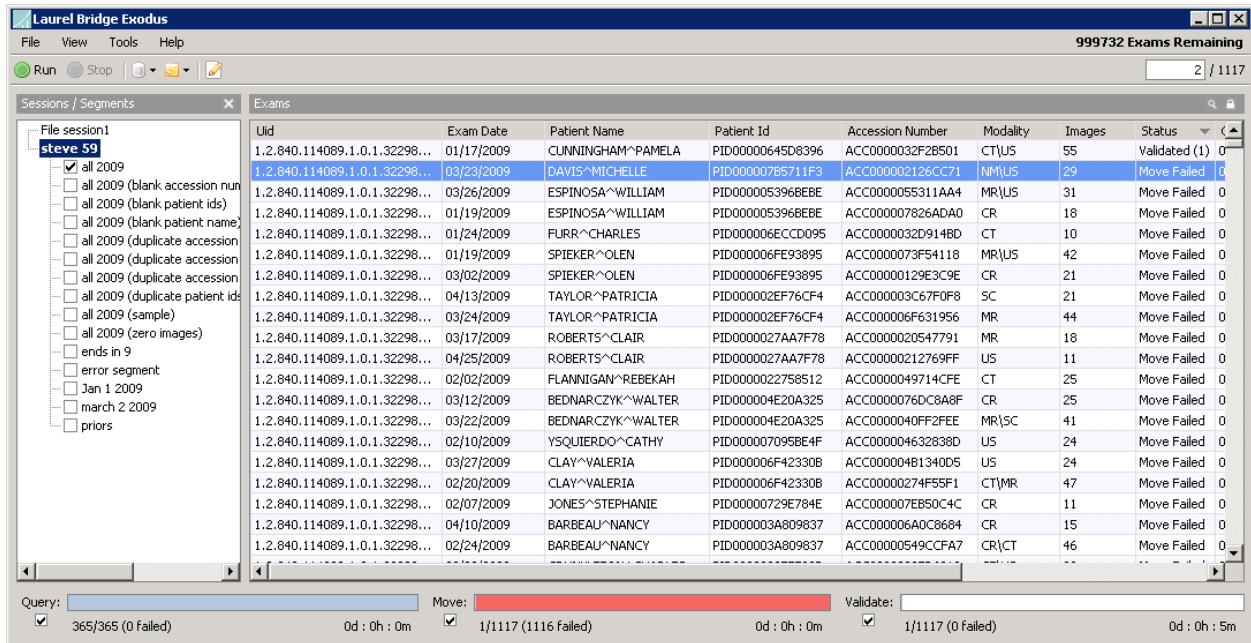


Figure 27 - Exodus Main Screen after attempting a migration

The Query tool shows that the (first and only) moved study contains data from line 2 of the mapping list CSV file; note that this study didn’t match line 1 since the “sex” field did not match.

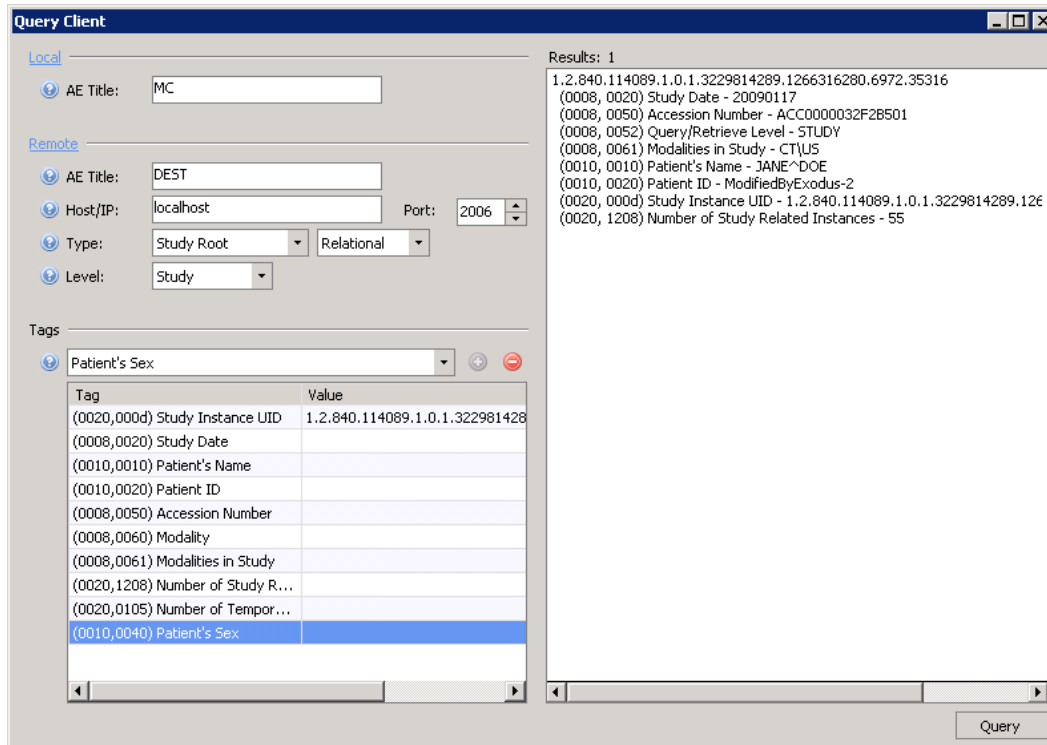


Figure 28 - Example 1 of the Query Client Tool

The second study listed on the Exodus main screen didn't move – Query Tool shows that on the source archive for that study UID, patient Michelle Davis was marked as having Gender “M” so it wasn't a match. The “No match action” was set to “Abort Association,” which caused the Move to fail. If the “No match action” had been set to “Continue,” then the exam would have moved, but its header data would not have been modified.

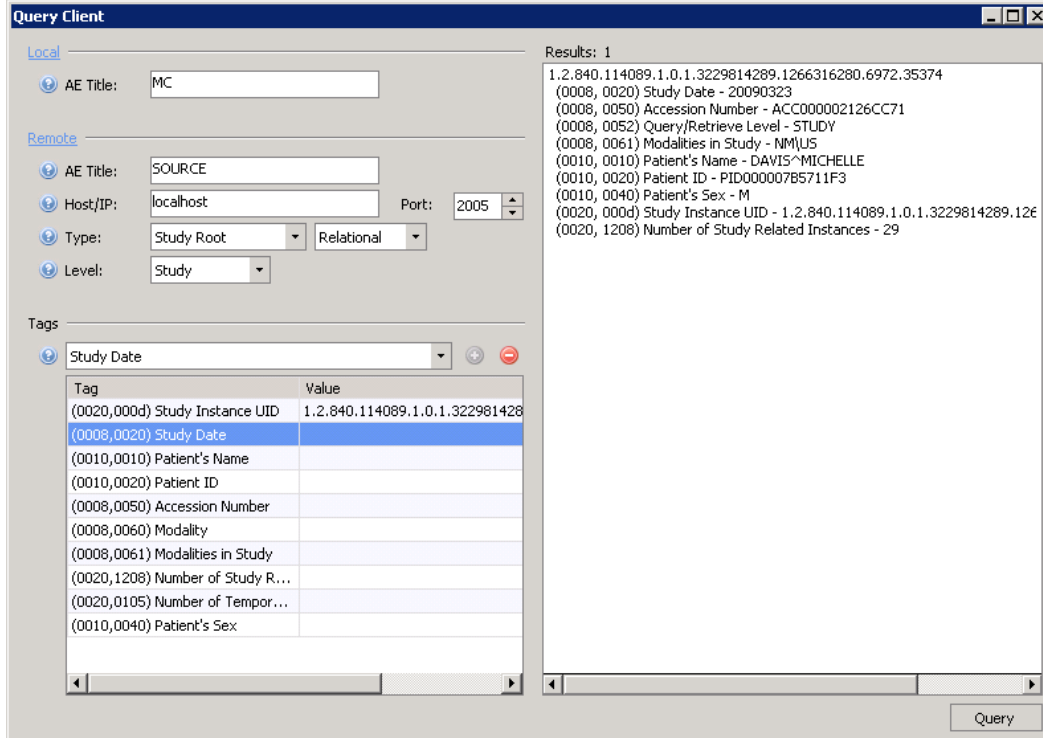


Figure 29 - Example 2 of the Query Client Tool

7.5 Composer Action Examples

The Composer action uses .NET regular expressions to parse an element’s value and combine the parts into new elements.

- 1) **Swap two tags** – Swap the Patient’s Name and the Patient ID to be in each other’s place you would specify the following inputs:

| Tag | Input Pattern |
|-----------|---------------|
| 0010,0010 | (.*) |
| 0010,0020 | (.*) |

The regexes shown here mean that the entire value should be one capturing group, then the outputs would look like this:

| Tag | Output Pattern |
|-----------|----------------|
| 0010,0010 | \${2.1} |
| 0010,0020 | \${1.1} |

This results in the first pattern $\${1.1}$ from the first input tag $\${1.1}$ being put into the second output tag (0010,0020), and the first pattern $\${2.1}$ from the second input tag $\${2.1}$ being put into the first output tag (0010,0010). (In this case, the first pattern is also the entire value.) So, if you started with “John Doe” and “1.2.3.4.5” in Name and ID respectively, your result would be a Patient ID of “John Doe” and a Patient’s Name of “1.2.3.4.5”.

- 2) **Split one tag into two tags** –To take the Accession Number (0008,0050) and keep only the first 10 characters in it and put the rest of it into the Requested Procedure ID (0040,1001). In this case, the regex for the input pattern has to specify how to split the Accession Number.

| Tag | Input Pattern |
|-----------|---------------|
| 0008,0050 | (^{10})(.*) |

This regex means the first 10 characters will be the first capturing group and everything else will be the second capturing group, then the outputs would look like this:

| Tag | Output Pattern |
|-----------|----------------|
| 0080,0050 | \${1.1} |
| 0040,1001 | \${1.2} |

This means that the first capturing group – the first 10 characters – will go into the Accession Number; everything else from the Accession Number will go into the Requested Procedure ID. If the initial Accession Number was “ABCDEF1234567890”, then you would have “ABCDEF1234” as the Accession Number and “567890” as the Requested Procedure ID. (Note that the output tag does not necessarily have to be parsed as an input.)

- 3) **Combine two tags** – To take parts of the Accession Number and parts of the Requested Procedure ID and “mix and match” them.

| Tag | Input Pattern |
|-----------|---------------|
| 0008,0050 | (^{10})(.*) |
| 0040,1001 | (^{6})(.4) |

These regular expressions mean to split the first tag into two capturing groups – the first one having 10 characters and the second one having whatever is left – and to split the second tag into two capturing groups, one of the first 6 characters and the second of the following 4 characters, then the outputs might look like this:

| Tag | Output Pattern |
|-----------|-----------------------------|
| 0080,0050 | \${1.1}---\${2.2}---\${2.1} |
| 0040,1001 | \${2.1}\${1.2} |

If the initial Accession Number was “ABCDEF1234567890” and the initial Requested Procedure ID was “1.2.3.4.5.6.7.8.9.0”, then the resulting Accession Number would be “ABCDEF1234---4.5.---1.2.3.”; the resulting Requested Procedure ID would be “1.2.3.567890”.

Note that the patterns can be used multiple times and also combined with plain text.

7.5.1 Working with DICOM sequences

A sequence may be entered as a tag by appending it to a numeric tag (the traditional group-element pair) with a period (“.”). You may also indicate an item in the sequence with “#” and the sequence item ID, followed by the tag indicating the sequence. There may be multiple sequences and sequence IDs as part of one “tag”. Examples are shown below:

- Simple tag - 0010,0010
- Tag within sequence - 0080,0100.0008,0060
- Tag within specific sequence item - 0080,0100.#0.0008,0060
- Tag within nested sequence with sequence items
0080,0100.#1.0080,0100.#0.0008,0060

If no item number is specified, the first item (#0) is assumed. Specify the last element in a sequence by “#L” (upper-case is important!) if the number of sequence items is unknown. Specify the next item in the sequence via “#N” (again, case is important) to append to the sequence. For example:
0080,0100.#L.0010,0010.#N.0008,0060

Please notice that:

- The sequence IDs (e.g., #1) and the tag-value pairs for the sequences are all separated by periods (“.”).
- The tags for the sequences are simple group-element pairs themselves.

7.5.2 Modifying private tags

Private tags can be modified using the **Composer Filter Action**. But first, you must add a private data dictionary to Exodus, so Exodus can know the VR (Value Representation) and VM (Value Multiplicity) of the private tags.

In the C:\ProgramData\Laurel Bridge Software\Exodus\cfg directory create a new directory called **dicom**. In that directory create a file called **ext_data_dictionary**. The **ext_data_dictionary** file must not have an extension. Its full pathname should be:

C:\ProgramData\Laurel Bridge Software\Exodus\cfg\dicom\ext_data_dictionary

The contents of the file shall look like this:

```
#
# The following is an example extended data dictionary file.
# If this file is named "ext_data_dictionary",
#
[ elements ]
0029,1020 = CS,1,Example Private Attribute 1
0039,1020 = US,1,Another Example Private Attribute 2
0049,1001 = DS,1,Private DS attribute 3
0049,1002 = UL,1,Private UL attribute 4
0049,1003 = SL,1,Private SL attribute 5
0049,1004 = UI,1,Private UI attribute 6
```

Note: A line beginning with a pound sign “#” is a comment.

The line containing “[elements]” must be present and can exist only once in the file. All other lines define private DICOM elements and should be specified in the format:

gggg,eeee (where gggg is the group number and eeee is the element number) followed by an equals sign, followed by three comma-delimited fields: VR, VM, and the long text description.

For example, to add private tag 0029,1020 to the data dictionary with a VR of CS and a VM of 1, the **ext_data_dictionary** file should have the following line added under the [elements] line:

0029,1020 = CS,1,Example private attribute text description

Note: You must restart Exodus after creating or modifying your private dictionary.

7.6 De-Identify Filter Action

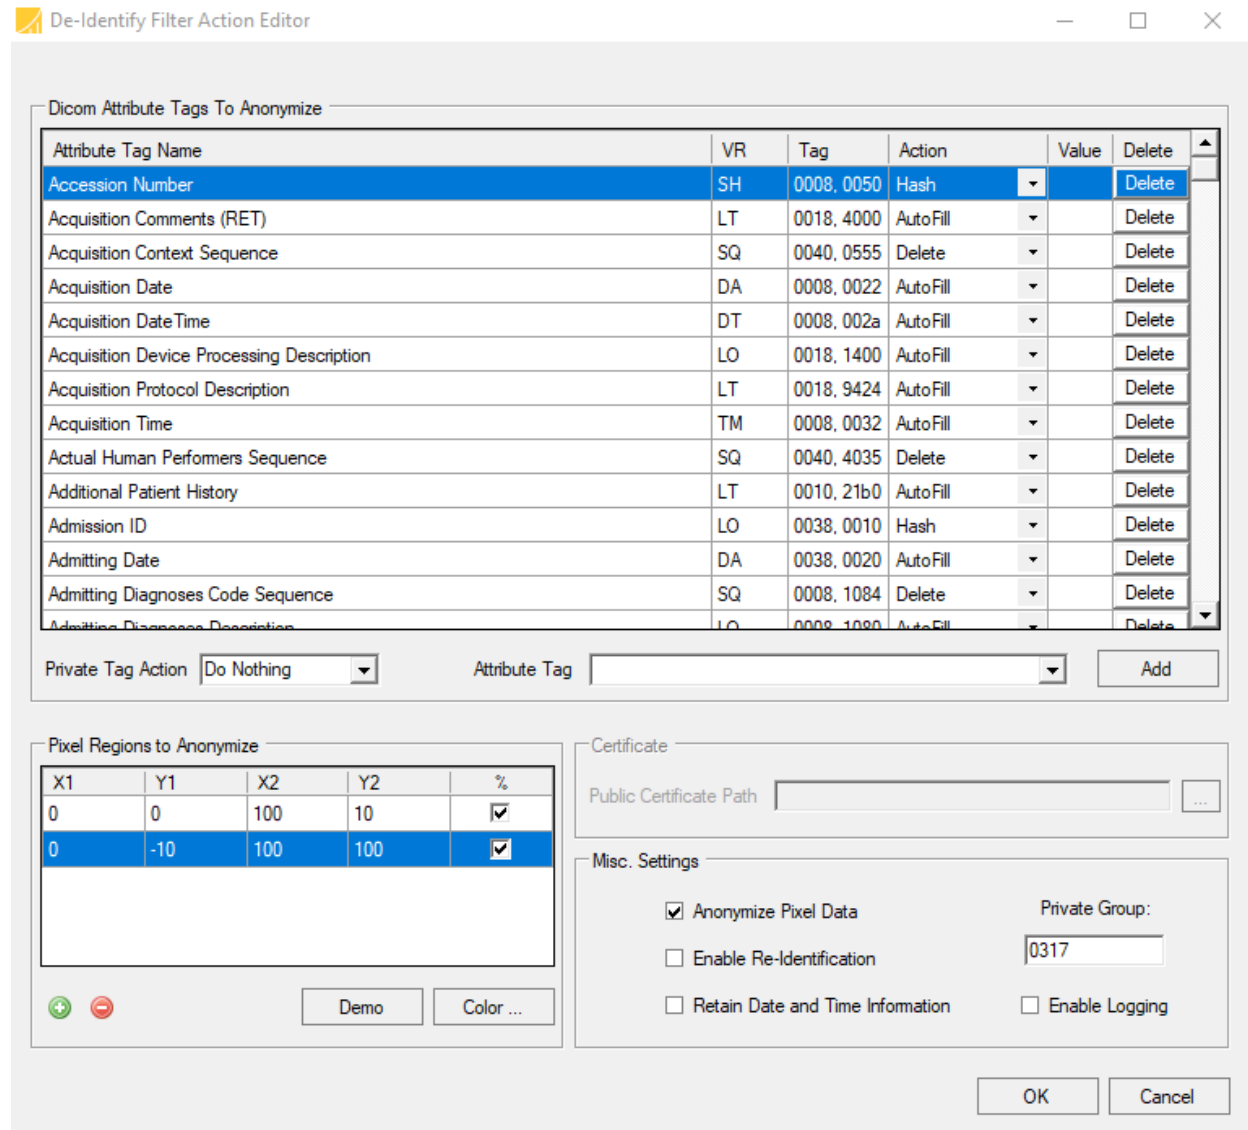


Figure 30 - De-Identify Filter Action window

The **De-Identify Filter Action** provides a mechanism to remove and replace certain attributes within a DICOM dataset that may lead to patient identification.

7.6.1 DICOM Attribute Tags to Anonymize

By default, this list is populated with attributes suggested in **PS 3.15 Annex E Attribute Confidentiality Profile** in the 2011 DICOM specification. Attribute tags may be added to the list of DICOM tags to

anonymize by using the 'Attribute Tag' drop-down combo box, which contains an enumeration of the public tags found in the 2011 DICOM dictionary. Public tags not found in the dictionary, nor private tags, may be added to the list of attribute tags to anonymize using this control. See section 7.6.5 for additional information regarding how the filter De-Identifies private tags. The following describes each aspect of an entry in the data grid shown above under [DICOM Attribute Tags to Anonymize](#):

- Attribute Tag Name – Specifies the name of the attribute tag to be anonymized if present in a given DICOM dataset's header data.
- VR – Indicates the attribute's Value Representation. This becomes useful when determining an appropriate replacement value for de-identification when "AutoFill" is selected.
- Tag – Indicates the attribute tag of the attribute to be anonymized.
- Action – Specifies the replacement action to use for a given attribute when determining the replacement value. The available replacement actions are determined based on the attribute's VR. The following list contains all available replacement actions:
 - Hash (SH, LO, PN elements) – For SH, LO, and PN elements, the replacement value will be a trimmed hash string (using SHA256) of the original value.
 - GenerateUID – For a given UI element, a new UID will be generated for the replacement value.
 - TransformUID – For a UI element, the existing UID will be transformed in a consistent manner to allow for repeatability. For instance, this option is typically selected for SOP Instance UIDs to allow a given UID to map to the same new UID across multiple instances of the [De-Identify Filter Action](#).
 - AutoFill – Replace the current DICOM attribute value with an appropriate replacement value based on the attribute's VR. The replacement value is predetermined and is not configurable.
 - UserDefined – Replace the current attribute value with the given replacement string defined in the "Value" column.
 - Default – Replace the current DICOM attribute value with a user-specified value. For example, specifying the default of 'M' for Patient's Sex will set all instances of Patient's Sex to 'M' by default.
 - Delete – Specifying this option will remove the given attribute from the de-identified dataset. If "Delete" is selected for a given attribute and the attribute is present in a sequence (SQ), that sequence will be removed from the de-identified dataset.
- Value – If "UserDefined" is selected as a replacement value action, the string specified in the Value column will be used as a replacement value for the given attribute. This value will appear as is. Care should be taken when entering a value for non-standard string VR types ("CS" for instance).
- Delete – Press the delete button to remove the given attribute from the list of attributes to de-identify.

7.6.2 Miscellaneous Settings

These settings allow for finer-grained control over common options needed during de-identification.

- Anonymize Pixel Data – If selected, the list of rectangles defined in the [Pixel Regions to Anonymize](#) grid will be removed from the DICOM dataset's pixel data. This includes pixel data in (7FE0, 0010), and not any additional overlays or private tags. This option will not work correctly if no valid pixel data is present in the DICOM dataset.

- Enable Re-Identification – If checked, the original values for each attribute anonymized will be persisted in the resulting de-identified dataset using the specified public certificate. More information about this process is specified later in this section.
- Retain Date and Time Information – If this option is checked, all date and time information (DT and TM elements) will not be anonymized. This option may be useful when performing clinical trials, where the DICOM dataset’s date and time information may need to be retained after de-identification.
- Private Group – The Private Group box specifies which group will be used during the de-identification process. After De-Identification, this private group will contain information necessary in determining what version of the DCF library was used during the De-Identification process.
- Enable Logging – Checking this option will enable logging for the de-identification process.

7.6.3 Pixel Regions to Anonymize

This grid specifies which rectangles to anonymize in a DICOM dataset’s pixel data (7FE0,0010). Each region is defined by providing the two corner points of a rectangle, (x1,y1) and (x2,y2), as either a percentage or absolute coordinates. The following section will outline the various options, as well as provide simple examples to demonstrate the various configuration settings.

The indicated regions to anonymize are only anonymized if the ‘Anonymize Pixel Data’ option is checked from the Misc. Settings section of the De-Identify Filter Action Editor. These same pixel regions are restored during re-identification if the ‘Enable Re-Identification’ option is selected and a valid Public Certificate is provided.

7.6.3.1 Percentage

Both points (x1,y1) and (x2,y2) may be given as a percentage of the image’s rows (y) and columns (x). This option is the most flexible, as it does not require prior knowledge of the image size of the DICOM image to anonymize.

For instance, the following images demonstrate the sample configuration for blanking out the top and bottom 10% of an image:

Consider the following percentage calculations for an 800 x 600 image:

Image Size: 800w (columns) x 600h (rows)

Region 1: x1, y1 = (0%, 0%) x2, y2 = (100%, 10%)

x1, y1 = (0%x800=0, 0%x600=0)

x2, y2 = (100%x800=800, 10%x600=60)

Region 2: x1, y1 = (0%, -10%) x2, y2 = (100%, 100%)

x1, y1 = (0%x800=0, -10%x600=540)

x2, y2 = (100%x800=800, 100%x600=600)

7.6.3.1.1 Negative Pixel Region Coordinates

Note how the above region 2 was calculated with a y1 percentage of -10%. Negative percentages correlate to percentages starting from the other end of image. In other words, negative percentages for width mean starting from the right side of the image opposed to the left. See the illustration below for more information about negative coordinate values. Note that negative values may be applied to both percentages and absolute coordinates and are treated the same.




Figure 31 - Negative Pixel Region Coordinate

7.6.3.2 Absolute Coordinates

Both coordinates (x1, y1) and (x2, y2) may also be given in absolute coordinates in reference to the actual number of rows and columns for a particular image. This mode is only recommended if the type of DICOM image remains constant for this filter, and the number of rows and columns are known at this point. This mode tends to be more precise than percentages when the exact location of the region (or regions) to anonymize are already known.

7.6.3.3 Demo Button

A demonstration button has been added to the Pixel Regions to Anonymize group to assist in the configuration of the list of regions to anonymize. After configuring the data grid with one or more regions to anonymize (again by specifying (x1, y1) and (x2, y2) coordinates), press the 'Demo' button to view a thumbnail bitmap of the regions specified to anonymize. (Please refer to the PixelRegionsToAnonymizeDemoForm image below for an example). The Pixel Replacement Color selected in the previous form will be used to render the rectangles seen drawn on the thumbnail bitmap.

NOTE: At the point the demonstration form is launched, the actual image rows and columns are not known. This will lead to inaccurate renderings when using absolute coordinates until an actual image is loaded via the '...' button located by the image path text box.




Figure 32 - Pixel Regions to Anonymize Demo Form

7.6.3.4 Color Button

Clicking **Color...** allows the user to define an appropriate replacement pixel value for each type of supported photometric interpretation via the following dialog box:




Figure 33 - De-identify pixel region color picker

The supported photometric interpretations for de-identification are as follows:

- MONOCHROME 1
- MONOCHROME 2
- RGB/YBR
- PALETTE COLOR

See the conformance statement for Anonymization for a complete breakdown of the supported image types for pixel data De-Identification.

Note: If you a color is picked that is not grayscale and photometric interpretation is MONOCHROME1 or MONOCHROME2 then grayscale representation for the selected pixel replacement color will be used.

7.6.4 One-Way Anonymization

For one-way anonymization of the patient demographic data and/or pixel data (re-identification cannot be performed), the following feature must be unchecked in the **Misc Settings** group box:

- **Enable Re-Identification**

NOTE: If the **Enable Re-Identification** setting is enabled without a valid **Public Certificate**, the de-identification will fail with the following error:

```
'LaurelBridge.DCS.DCSException: Failed to apply filter action --->  
System.Security.Cryptography.CryptographicException: Cannot find the requested object.'
```

The **Public Certificate Path** text box will be disabled when performing one-way anonymization. A valid public certificate is only necessary when **Enable Re-Identification** is checked. A valid certificate should either be purchased or generated by the user.

The expected certificate file formats are:

- the public certificate (export format PEM): <filename>.crt
- the private certificate (export format PKCS #12): <filename>.p12

7.6.5 Additional Notes

Pixel data and header data re-identification are only available with a valid public and private certificate.

Disable **Use Original Attributes Sequence** for all de-identification filters if you are going to anonymize SOP instances. Having this option enabled will defeat the purpose of anonymization by embedding the list of attribute tags and their original values that changed after applying the de-identification filter action.

The **De-Identify Filter Action** also supports removing all private tags from the DICOM dataset's header data. The current options and their descriptions are as follows:

- Do Nothing – Leave the private tags as they are. By default, this option is selected.
- Remove All – Remove all private tags from the DICOM dataset's header data. This does not include any private tags added during the de-identification process. This removal is unidirectional and cannot be undone, even during re-identification.
- Encrypt All – Remove all private tags from the DICOM dataset's header data. This does not include any private tags added during the de-identification process. This option can only be enabled if **Enable Re-Identification** is selected and a valid public certificate is provided.

Sequences are De-Identified per the following rules:

- If a given sequence (SQ element) contains an attribute tag that is marked for de-identification, the entire SQ element is anonymized. That is, each of the DICOM elements contained within that sequence are anonymized per either the specified tag replacement action, or AutoFill by default when determining the appropriate replacement value.
- If a sequence is marked as 'Delete', then the entire sequence is removed. A sequence will not be only partially altered. If any of the contents of a sequence are marked for de-identification, regardless of the action, then the entire sequence will be de-identified.

7.7 Re-Identify Filter Action

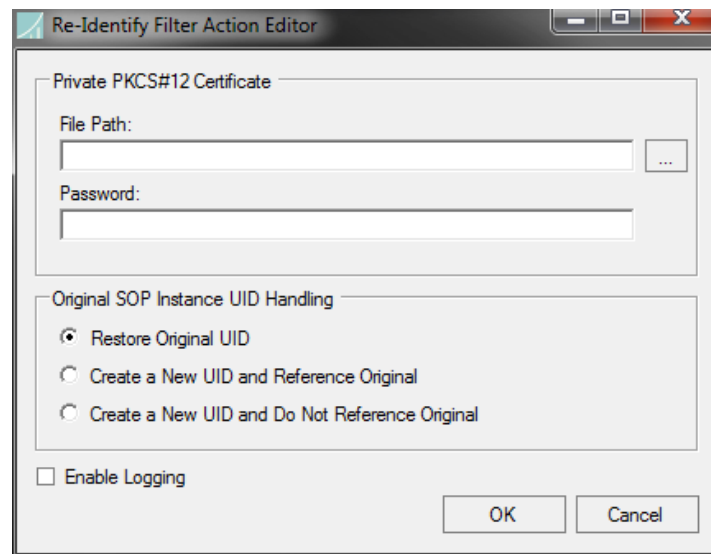


Figure 34 - Re-identify Filter Action Editor window

The **Re-Identify Filter Action** allows restoration of previously anonymized patient demographic information and pixel data for a given DICOM dataset.

The patient demographic information and pixel data restoration are only available if a valid private certificate (and password) is available. This private certificate must pair with the public certificate used to de-identify the original DICOM dataset.

Select one of the following actions when re-identifying a dataset to dictate how to handle the original SOP Instance UID:

- Restore Original UID – Select this option to restore the original SOP Instance UID. This option will not create a Referenced Image Sequence (0008, 1140).
- Create a New UID and Reference Original – Select this option in order to create a new SOP Instance UID and create a Referenced Image Sequence (0008, 1140) that references the original SOP Instance UID that was de-identified.
- Create a New UID and Do Not Reference Original – This option will create a new SOP Instance UID for the re-identified dataset without creating a Referenced Image Sequence (0008, 1140).

7.7.1 Additional Notes

Enable Logging – This option enables logging for the re-identification process.

Patient Identity Removed (0012, 0062) Attribute Tag will be added to the re-identified dataset in all cases, including selecting “Restore Original UID”.

As per [PS 3.15 Annex E Attribute Confidentiality Profile](#), the following Attribute Tags are added when re-identifying a given DICOM dataset, unless restoring the original SOP Instance UID:

- Patient Identity Removed (0012, 0062)
- Purpose of Referenced Code Sequence (0040, A170)
- Referenced Image Sequence (0008, 1140)

Appendix A Exodus Privacy and Security Statement

Because the Laurel Bridge Exodus application is installed on hardware that is provided, configured, and controlled by the Exodus customer, Laurel Bridge Software (LBS) cannot make assertions about the privacy and security of a particular Exodus installation. It is up to the customer to ensure that the host Windows system onto which Exodus is installed has been adequately secured and locked down. However, LBS does provide technology, tools, and guidance to assist customers in locking down their Exodus installations. In the context of this appendix, the term “Exodus customer” refers to the administrators for the host hardware system and for the Exodus application.

Section 3 - GDPR Notes found below contains comments regarding the European Union’s (EU’s) GDPR - General Data Protection Regulation.

An overview of the Exodus application privacy and security features is given in the sections below, roughly following the format given in the HIMSS/NEMA Standard HN 1-2013, “Manufacturer Disclosure Statement for Medical Device Security”, or MDS2 for short. For more details about this form or to download it, see <https://www.nema.org/Standards/Pages/Manufacturer-Disclosure-Statement-for-Medical-Device-Security.aspx>

The headers in the following sections map directly to the headers in the MDS2 document. The Exodus MDS2 document is included with the installation files and is available for a particular release version upon request from LBS.

1 Management of Private Data

The Laurel Bridge Exodus application migrates DICOM instances, which may contain protected health information (PHI), from one host system (usually an archive) to another. Exodus can move these instances from one or more sources to one or more destinations. To do this, Exodus may ingest, store, display, and transmit PHI. However, since the PHI only resides in Exodus temporarily, Exodus is not considered a primary repository of electronic health record (EHR) or electronic medical record (EMR) data, and thus is not maintaining part of the designated record set (as defined by HIPAA). Also, the Exodus application and the data it stores and manages is entirely resident within the customer premises (i.e., no part of the application or its data is cloud-hosted or hosted by LBS).

1.1 Types of PHI Maintained

While migrating DICOM instances, Exodus potentially transports and caches the following types of PHI:

- Patient demographic information
- Patient medical record information
- Patient diagnostic and therapeutic information (including diagnostic images)
- Patient financial information

1.2 Persistence of Private Data

Exodus maintains PHI both temporarily in memory (while running) and on disk (persistent storage). PHI may be found in data transmitted or cached by the application, and in log files generated during use of the application. Data can be imported from other medical systems via network-mounted hot folders or removable media, but these features must be explicitly configured by the Exodus customer, and their use remains under the full control of the customer. Available security features to protect PHI when at rest are described below and, in more detail, elsewhere in this Exodus User Manual.

Note: Due to the sensitive nature of the PHI that Exodus handles, the only non-destructive and completely safe way to decommission a (non-virtual) computer system on which a production Exodus application has been running is to wipe the hard drive clean using a suitable hard drive wiping application. For self-encrypting drives, changing or overwriting the encryption key(s) may be sufficient.

1.3 Transmission of Private Data

PHI can be transmitted or received over the network via DICOM or via files. The ability to configure and control the behavior of this functionality is under the full control of the Exodus customer, and the use of these features remains under the full control of the customer. Available security features to protect PHI when in transit are described below and, in more detail, elsewhere in this Exodus User Manual.

1.4 Payment Card Industry (PCI) Data Security Standard

Because Exodus does not process any patient billing transactions, it is not subject to the requirements of the Payment Card Industry (PCI) Data Security Standard.

2 Security Capabilities

The Laurel Bridge Exodus application is comprised of a single, monolithic Windows Forms application. The following sections briefly describe available security features of the Exodus application. For more details, see the Exodus User Manual.

2.1 Automatic Logoff

The Exodus application does not automatically time out or log a Windows user off, so LBS recommends that the Windows password-protected, inactivity-activated screen lockout functionality be enabled for all users with Exodus privileges on the host Windows system, to help avoid unintended administrative access by unauthorized users.

2.2 Audit Controls

Exodus can be configured to send DICOM PS3.15 Appendix A.5 (“Audit Trail Message Format Profile”) audit messages to a syslog server (such as **syslog-ng** or **nsyslog**). Messages can be sent via the TLS (recommended), UDP, or TCP protocols, and all messages include the user ID of the user performing the action as well as a date/time stamp.

The following types of audit trail messages can be enabled/disabled independently:

- **Application Start/Stop** – Logs when an application is started/stopped.
- **Software Configuration** – Logs when changes are made to the software configuration.
- **User/Security Alerts** – Logs when web user or security alerts occur.
These include events such as web user logon/logoff, web user addition/removal, web user password/role changes, and manual modifications of DICOM or HL7 jobs.
- **Queries/Retrieves** – Logs when instances are queried or moved.

The following DICOM PS3.15 Appendix A.5 audit trail message types are supported by Exodus:

- **Application Activity**
 - Application Start
 - Application Stop
- **Security Alert**
 - Software Configuration
 - Use of Restricted Function
- **Queries**
 - DICOM queries
 - DICOM retrieves

2.3 User Authorization

Windows Authentication (whether locally-administered or domain-based) is used to control access to the Exodus application. It is up to the customer to lock down administrative control of Exodus (e.g., by using a custom Windows user-level group).

2.4 Security Configuration

The Exodus customer has full control over and responsibility for the security of Exodus, both through the ability to lock down the Windows system on which Exodus is installed, as well as through the ability to

configure the security features built into the Exodus application. Extensive information about how to do this is found in this Exodus User Manual.

2.5 Security Updates

The Exodus customer has full control over the installation of Windows security updates, as well as over the installation of any Exodus application updates.

2.6 De-Identification of PHI

Exodus does not support the de-identification of PHI. When sending PHI over the public internet to / from satellite locations, the use of TLS encryption or an encrypted VLAN is the way recommended by the DICOM Standard to protect the confidentiality of PHI in transit.

2.7 Backup and Restore

The Exodus customer has full responsibility to both install and maintain the SQL Server database which provides the backing store for the Exodus data. As such, the customer is also responsible for providing backup and restore capabilities for the SQL Server database. Microsoft provides an extensive set of SQL Server backup, restore, and replication technologies.

2.8 Emergency Access

Since the Exodus customer has full control over the installation and configuration of both the host system and the Exodus application itself, it is up to the customer to provide a means of emergency access (“break-glass” feature) by maintaining alternate access to administrative credentials for the systems involved, if necessary.

2.9 Data Integrity and Authenticity

Since the primary function of Exodus is to query and move DICOM instances between other systems, it is simply not practical to implement a mechanism whereby alteration of data can be detected. However, the following techniques can be used to control access to the data:

- Use Audit Trail logging to record any access to or modification of data.
- Use Windows Authentication to ensure that unauthorized Windows users cannot access the host Windows system on which Exodus is installed.
- If desired, use TLS encryption on the network connections used by the system to ensure privacy, node authentication, and protection against man-in-the-middle (MITM) attacks.

Exodus does not currently use explicit error detection on data at rest, but rather depends on the built-in ECC error detection and correction technology provided by modern hard drives (as supported by Windows). If data redundancy is desired, LBS recommends the use of RAID data storage technology for both the SQL Server database repository and for the migration data.

2.10 Malware Protection

Since the Exodus customer has full control over the installation and configuration of both the host Windows system and the Exodus application itself, it is up to the customer to install and maintain malware protection technology. Exodus itself should be unaffected by the use of such technology (beyond the

obvious potential impact to system performance that can occur when using anti-virus software). For better performance, it is generally recommended that antivirus checking be turned off for the SQL Server data directories used by Exodus.

2.11 Node Authentication

Node authentication (the ability to confirm the identity of both the sender and receiver of DICOM and HL7 data) can be implemented using TLS protocols on all network connections. Exodus supports TLS versions 1.0, 1.1, and 1.2 as both client and server. TLS must be enabled separately for each DICOM input and output. More details about how to do this and further security details can be found elsewhere in this Exodus User Manual.

2.12 Person Authentication

As mentioned earlier, user authentication for the host Windows system can be controlled locally, using a domain with single sign-on technology such as LDAP / Active Directory. User authentication for web interface users can also be controlled either locally or using LDAP/AD.

2.13 Physical Locks

Since the Exodus customer owns and has full control over the host Windows system on which Exodus is installed, it is up to the customer to maintain the physical security of the host system.

2.14 Device Life Cycle Roadmap

The Exodus application currently supports the following Windows operating systems:

- Windows 10
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019

LBS intends to support each of these operating systems up until their respective end-of-extended-support dates.

In addition, the Exodus application has the following software dependencies:

- SQL Server (can be SQL Server 2012 x64 or newer)
- SQL Server Management Studio
- .NET Framework 4.8 (or later)

2.15 System and Application Hardening

Since the Exodus customer provides, configures, owns, and has full control over the host system on which Exodus is installed, it is up to the customer to perform system hardening, as well as to configure the Exodus application for the desired level of application hardening. More details about hardening of the host Windows system and the Exodus application can be found elsewhere in this Exodus User Manual.

Some specific application hardening techniques that are supported by and/or implemented in Exodus include:

- Use of Authenticode digital signatures (currently SHA256) for all LBS executables and DLLs

- Support for TLS encryption for data in transit
- Provision of instructions for how to lock down the TLS protocols and ciphers, which affects any DICOM connections which are configured to use TLS encryption

The implementation of the following lockdown techniques on the host Windows system is the responsibility of the Exodus customer:

- Disabling of unnecessary Windows accounts
- Disabling of unnecessary open network ports (e.g., telnet, ftp, etc.)
- Removal of any unnecessary off-the-shelf applications
- Enabling of Windows password-protected, inactivity-activated screen lockout
- Disabling of the ability to boot from removable media (if physical access to the host Windows system cannot be controlled)
- Enabling of BitLocker or other at-rest, full-disk encryption technologies (if desired)
- Enabling of SQL Server encryption (especially if the database resides on a different, unencrypted system)

2.16 Security Guidance

The security-related features of the Exodus application are described in detail in this Exodus User Manual.

2.17 Data Storage Confidentiality

Exodus does not encrypt data while at rest on the hard drive(s). PHI is stored both in the SQL Server database, as well as possibly in cached data files (depending on how Exodus is configured). If at-rest encryption of PHI is deemed necessary (e.g., if physical access to the host Windows system cannot be controlled), we recommend the use of a full disk encryption technology such as BitLocker or the use of self-encrypting drives. SQL Server at-rest encryption technologies such as Transparent Data Encryption (TDE) may also be necessary if the SQL Server database is resident on a different (unencrypted) system. Exodus does support encrypted SQL Server connections, and their use is highly recommended in the case of SQL Server instances accessed over a network.

2.18 Data Transmission Confidentiality

Exodus can be configured to encrypt data in transit (using TLS), which will protect the data against interception by unauthorized parties. It can also be configured to use reversible de-identification, which will encrypt PHI in such a manner that it can be decrypted by the authorized receiver. And as mentioned above, Exodus supports encrypted SQL Server connections, and LBS highly recommends using them in the case of SQL Server instances accessed over a network.

2.19 Data Transmission Integrity

TLS encryption also protects the data against any attempt to modify the data during transmission (i.e., via man-in-the-middle attacks). Exodus will only transmit data to destinations that have been explicitly configured within the application by the customer.

2.20 Other Security Considerations

Exodus can be serviced remotely by LBS only with the express permission of the Exodus customer, as access to the host system onto which Exodus is installed is completely controlled by the customer. Exodus

does not contain any service backdoors, nor does it contain any secret service accounts. All LBS access to an installed Exodus application must be explicitly enabled/allowed by the customer using standard Windows secure remote access technologies.

By default, Exodus does not listen on any network ports.

3 GDPR Notes

The European Union's (EU) General Data Protection Regulation (GDPR) is a refresh of Europe's data-protection laws that harmonizes statutes across the 28 EU member states; it became effective May 25, 2018. GDPR is a law that applies to any organization doing business in the EU or with EU-based clients. It is up to the Laurel Bridge application customer to ensure that they manage the Exodus application and the medical imaging data processed by it in a way that is conformant to their GDPR policies and practices. Relative to the GDPR some key points to remember are:

- The Laurel Bridge Exodus application is installed on virtual or physical systems that are provided, configured, and controlled by the customer, therefore Laurel Bridge Software (LBS) cannot make assertions about the privacy and security of a particular installation.
- It is up to the customer to ensure that the customer's host systems on which the application components are installed have been adequately secured.
- By virtue of using this application, Laurel Bridge Software receives no private data from the customer or the customer's clients; data remains with and under the control of the customer.
- The application does not maintain a designated record set and is not a primary repository of electronic health record (EHR) or electronic medical record (EMR) data. Data processed and tracked by the application is typically transient and purged after the completion of a migration.
- Log files and generated reports may possibly contain private data associated with the medical imaging data being processed. Such files should be handled in a way that is compliant with the customer's data retention and privacy policies.
- The application's management database will contain private data associated with the medical imaging data being processed by the migration activities. This database should be purged after the completion of a migration or dispositioned in a way that is compliant with the customer's data retention and privacy policies, including GDPR.

Appendix B How to not migrate (i.e., drop) non-image DICOM objects

If a user does not want to migrate and instead wants to drop a non-image type of object, such as a structured report (SR), he can create a filter with a **Drop Action** set to **Drop** the particular type of instance. However, one must be aware that this action could result in later **Validation** failures because there will be missing instances on the **Destination PACS** when it comes time to compare (i.e., validate) source and destination files.

One way to work around this is to create a second validation rule that will query at the Image level of the source system and drop any C-FIND Response DIMSE messages (C-FIND-RSP) that contain a response for the given non-image object type. For more information on performing a validation in this context, see Appendix D Exodus FAQ, item number 9 – “I did not migrate SR (Structured Report) objects; how can I perform validation?”

Example—Do not migrate Structured Report objects:

Let us assume we need to drop all Basic Text Structured Reports (DICOM SOP Class: 1.2.840.10008.5.1.4.1.1.88.11)

Step 1: Create a tag filter to **Drop** any Instance of DICOM Basic Text SR and return a C-STORE-RSP of **success**. Displayed below is the **Filters Tab** of the Session Properties Dialog showing this configuration.

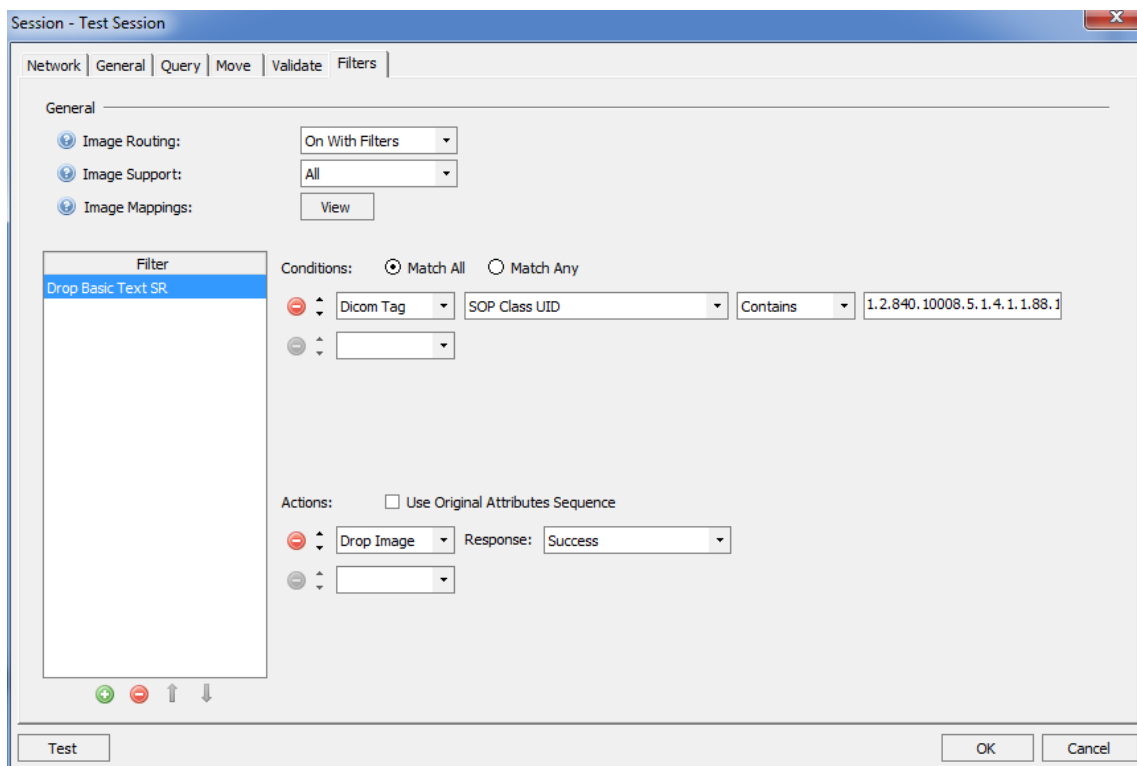


Figure 35 - Create a tag filter to Drop an instance

Step 2: Create a **Level 2 Validation Rule** that Queries the Source PACS at the Image Level. Displayed below is the **Validate Tab** of the Session Properties Dialog showing this configuration.

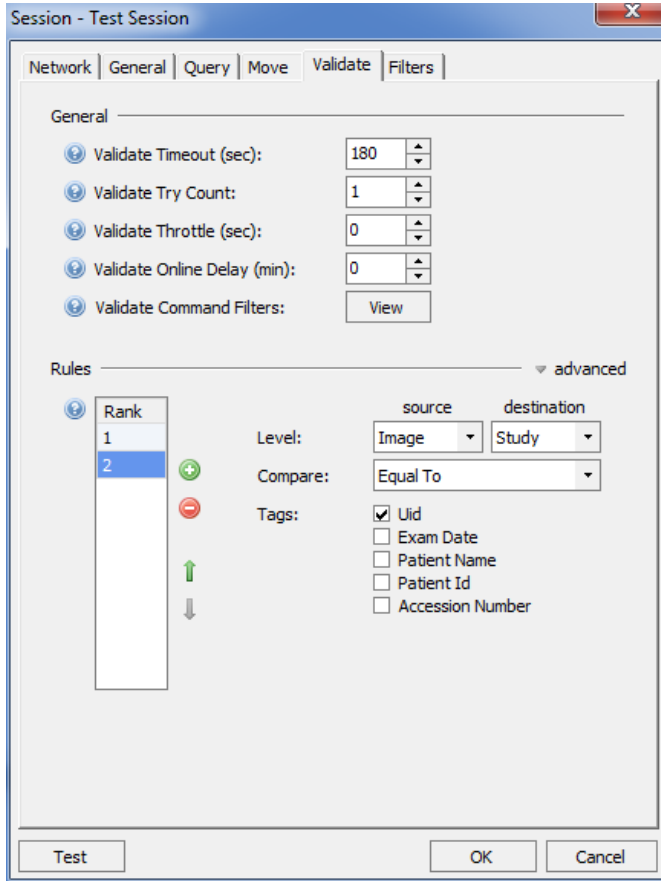


Figure 36 - Level 2 Validation Rule that Queries the Source PACS at the Image Level

Step 3: Create a **Validate Command Filter** to Request SOP Class UID tag from Source PACS. Displayed below is the Validate Command Filters Dialog. This dialog can be reached by clicking the “View” button on the **Validate Tab** of the Session Properties Dialog.

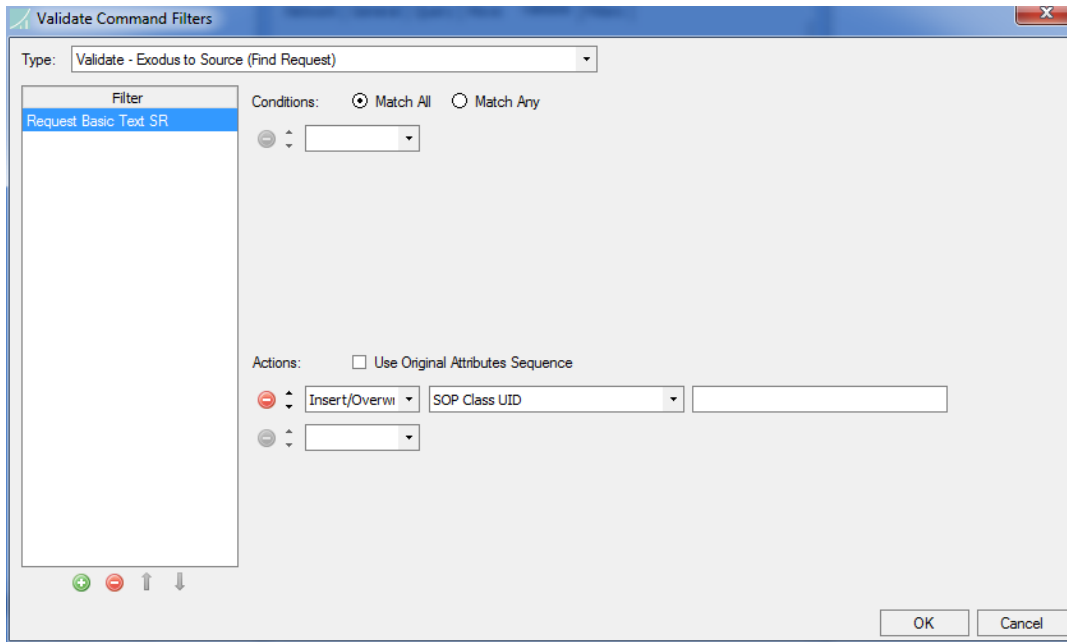


Figure 37 - Validate Command Filter to Request SOP Class UID tag from Source PACS

Step 4: Create a **Validation Command Filter** to drop Basic Text SR from responses from the Source PACS. Displayed below is the Validate Command Filters Dialog. This dialog can be reached by clicking the “View” button on the **Validate Tab** of the Session Properties Dialog.

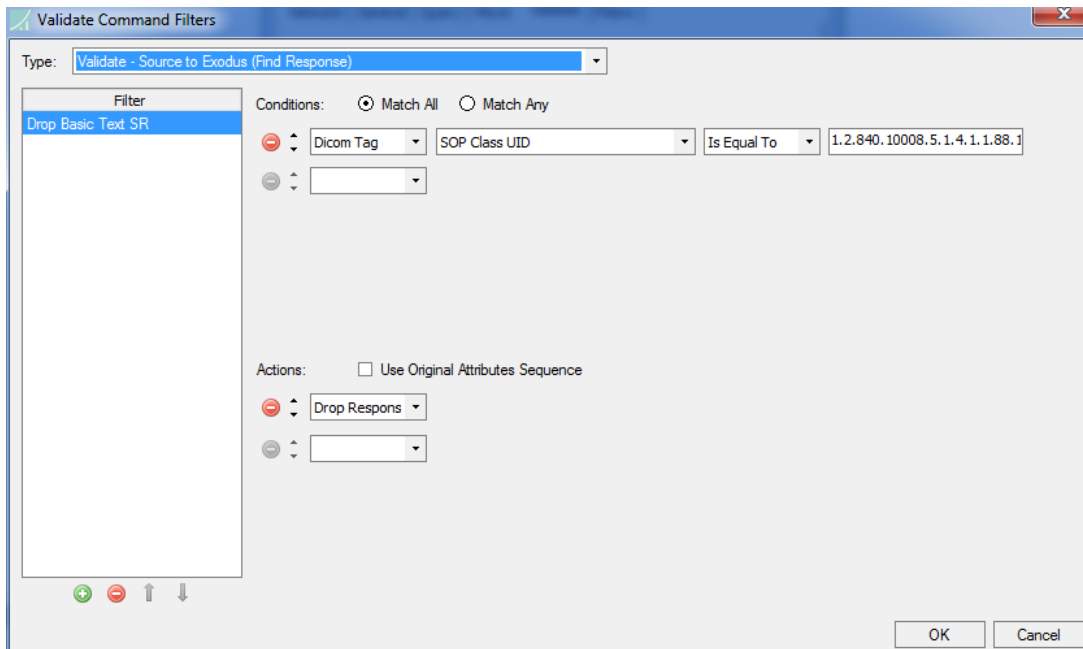


Figure 38 - Validation Command Filter to drop Basic Text SR from responses from the Source PACS

Appendix C Creating a Priors Fetch Segment for Exodus Priors Fetcher

With the former Exodus Priors Fetcher application version, you could only create **Priors Segments**. This optional product functionality has been retired.

Priors fetching for the source archive can still be performed as part of a migration by using the priors segment. More comprehensive priors fetching may be done by the use of a different Laurel Bridge workflow solution, typically Navigator.

Appendix D Exodus FAQ

Listed below are a collection of Frequently Asked Questions about Exodus and its use in migration related scenarios. Additional questions may be submitted to support@laurelbridge.com .

1. How do I know how many Exams I am allowed to move?

In the upper right-hand corner of the **Main Dialog**, the number of Remaining Exams on your currently installed license is displayed (see below). This number is just that, how many more **Moves** you can perform. You can still query and validate once this number reaches 0 (zero), but you will not be able to migrate any more exams without a new or upgraded product license.

2. What do the numbers at the bottom of the Main Dialog mean?

If you are doing a typical migration with Date Range **Segments**, then the **Query** number under the **Query** progress bar is the number of days that have been successfully queried and the total number of days remaining to query. The number of failed queries is shown in parentheses. (See also FAQ 3 below.) Under the **Move** progress bar in the bottom middle of the **Main Dialog**, you see the number of Exams successfully moved and the total number of Exams that have been found so far. The number of failed **Moves** is shown in parentheses. These move numbers are representative of the currently selected **Segments** in the left panel. Under the **Validate** progress bar, the number of Exams that **Validated** to one of currently defined **Validation Criteria** is displayed, next to the total number of exams to validate. The number of **Exams** with a state of **Validate Failed** is displayed in parentheses. These numbers are representative of the currently selected Segments.

| Uid | Exam Date | Patient Name | Patient Id | Accession Number | Modality | Images | Status | Query Time | Move |
|--------------------------------------|------------|-----------------|--------------|-------------------|----------|--------|--------------|-----------------|-------|
| 1.2.840.114089.1.0.1.3229814446.1... | 12/27/2007 | LOWE^CYNTHIA | PID000003... | MRN000007D808506 | SC | 3 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/24/2007 | HILL^WILLA | PID000007... | MRN0000060E79188 | MR | 5 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/29/2007 | GONZELAS^MA... | PID000004... | MRN000004AD98383 | CR MR | 7 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/23/2007 | PROSPERI^WI... | PID000000... | MRN00000015C2F97E | US | 3 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/23/2007 | TAYLOR^PAUL | PID000002... | MRN000000CFE606F | CR US | 4 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 9/9/2007 | SILVERNAIL^L... | PID000005... | MRN0000001F327C38 | CT US | 7 | Moved | 9/9/2011 1:4... | 9/9/2 |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/22/2007 | COX^CRAIG | PID000002... | MRN00000058043C8D | SC US | 2 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/24/2007 | PETERSON^PE... | PID000005... | MRN00000036896F8B | MR NM | 5 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 9/9/2007 | TRINIDAD^KIR... | PID000001... | MRN0000005E11A43E | CR US | 3 | Moved | 9/9/2011 1:4... | 9/9/2 |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/23/2007 | DOVERSPIKE^... | PID000000... | MRN0000005BF1E7E9 | NM US | 5 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/20/2007 | BANK^WALTER | PID000005... | MRN00000018188C62 | US | 4 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/24/2007 | MUELLER^PAM... | PID000003... | MRN00000023E61D0F | CT NM | 4 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 9/9/2007 | ROSENSTEIN^... | PID000003... | MRN00000076CABDD8 | NM | 5 | Moved | 9/9/2011 1:4... | 9/9/2 |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/20/2007 | BIHM^RAUL | PID000005... | MRN0000007E4AC079 | CT SC | 6 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/22/2007 | TILL^NATASHA | PID000005... | MRN0000007532D6C | CT SC | 8 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/30/2007 | GILMER^JOSEPH | PID000005... | MRN000000274C6482 | CT | 2 | Validated... | 8/25/2011 1:... | 8/29/ |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/23/2007 | WETZEL^ROBERT | PID000005... | MRN00000025476155 | CR | 1 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/28/2007 | HAMRICK^GALE | PID000001... | MRN00000050EB0A49 | CR US | 6 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/20/2007 | HOLLAUBAUGH^... | PID000004... | MRN0000006DA4886F | CT | 4 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/24/2007 | RIVERA^NICOLE | PID000007... | MRN0000000AC42C6C | MR US | 7 | Waiting | 8/25/2011 1:... | <n/a> |
| 1.2.840.114089.1.0.1.3229814446.1... | 12/27/2007 | WOODBURY^E... | PID000000... | MRN0000001D927378 | MR US | 6 | Waiting | 8/25/2011 1:... | <n/a> |

Query: 16/365 (0 failed) Move: 244/1897 (0 failed) Validate: 13/1897 (6 failed)

Days Queried/Total Days Exams Moved/Total Exams Exams Validated/Total Exams

Figure 39 - Meaning of numbers on the Main Exodus Dialog Window

3. What does the number under the Query Progress Bar mean if I do not have a Date Range Segment?

If you create a List Segment from a file or a command, then the number of Patient IDs or number of Study UIDs will be displayed.

If you create a Fixed Segment by selecting exams from another segment, then 0 (zero) Queries will be displayed.

4. Exodus seems to be running slow, what can I do?

Locking the View of the Query Viewer and Exam Viewer will reduce some significant display-related computation performed by Exodus. Doing this is especially useful if you are not actively watching the migration and are just letting it run. You should click the Padlock icon on both viewers.

Another technique to improve apparent performance is to query the entire source first, and then do the Moves and Validations as a separate step. Querying is typically quite fast and may take less than a day for even very large archives.




Figure 40 - Locking the View of the Query Viewer and Exam Viewer

You should also verify that you are not logging more information than you need to. To turn off logging except for errors, under the Tools->Options menu, you should have the log level set to 'Normal' and uncheck 'DICOM' logging. Higher log settings will provide useful debugging information but at a performance penalty.

5. Should I Query and Move and Validate all at the same time?

For typical PACS to PACS migrations, we recommend querying the entire possible range of dates of the Source PACS before migrating any Exams. You want to know a count of how many exams you have to move before starting. You will most likely move a sample of the Source PACS' data to the Destination during the assessment phase of the migration to check for image quality, etc. You will need to prequalify the Source PACS data and create any tag morphing filters where appropriate. In order to have a good representative sample and prequalification, you will need to query for everything first. Also, Querying and Moving simultaneously may reduce performance of the Exodus tool.

6. How do I configure Exodus to use Tag Filters for data being moved?

Open the **Session Properties Dialog**. Select the **Filters Tab**. Set the **Image Routing** option to **On With Filters**.

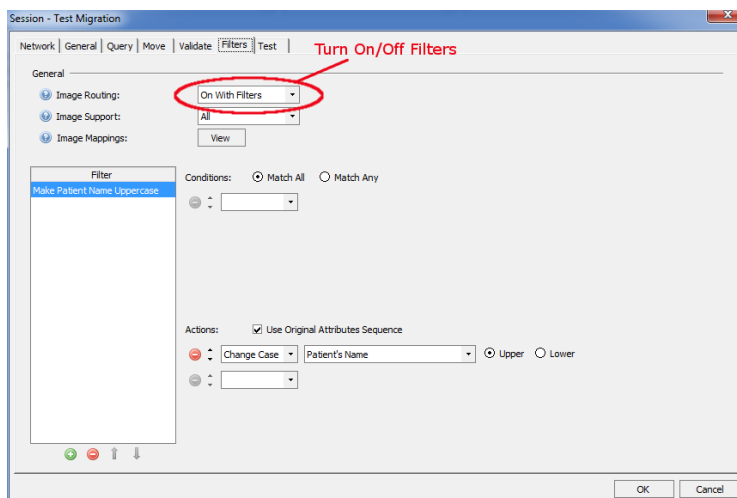


Figure 41 - Setting Exodus to use Tag Filtering

Once you have selected **On With Filters**, Exodus will send the Exodus AE Title as the Move Destination to the Source PACS; Exodus will also use the Exodus AE Title as the move destination for **On Without Filters**.

You must have the Exodus AE Title and port number configured in the Source PACS as a Move (or DICOM Storage) destination.

You must have the Exodus AE Title and port number configured in the Destination PACS as a source for images (datasets).

Exodus will process each and every image and DICOM object sent to it and apply Tag Filters in the order that they are presented on the **Filters Tab**; from the top to bottom of the list. The output of one filter is the input to the next filter in the list, allowing you to chain many filters together. After filtering is complete, Exodus will then forward the filtered datasets on to the destination PACS.

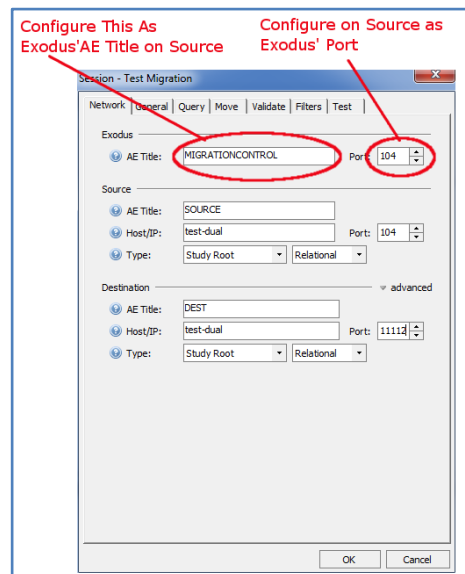


Figure 42 - Exodus AE-Title & Port

7. How do I apply a filter to all images (datasets)?

Leave the Conditions section of the filter definition empty and click the **Match All** radio button.

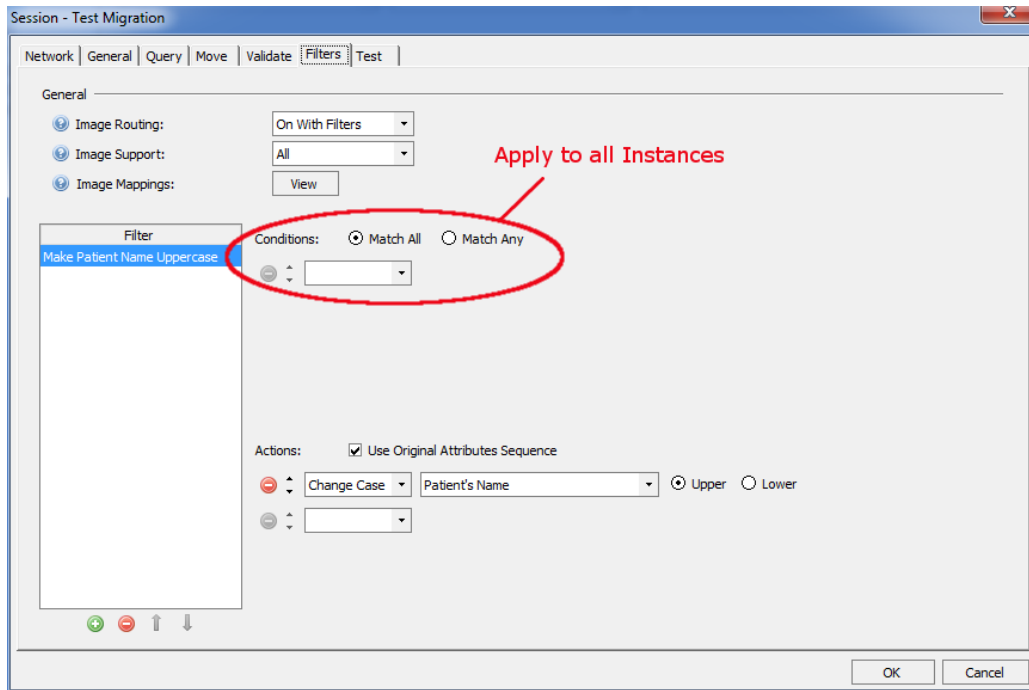


Figure 43 - Example: apply a filter to all images (datasets)

8. How do I validate a migration where new images are added to the destination?

If you are validating a previous migration with the Exodus Assessment tool, new images or reports may be added to the exams on the destination PACS after the original migration has been completed. This change could affect the “Number Of Study Related Instances” returned by a Study-Level query to the Destination PACS, therefore, the default image counting **Validation** rules will not work as originally expected.

To accommodate this change in image counts, you should create a **Validation** rule that queries at the Image Level on both the Source and Destination PACS and set the comparison mode to be **Contained By**. If the list of SOP Instance UID’s returned from the Source PACS is contained in the list of SOP Instance UID’s returned from the Destination PACS, then the **Validation** will succeed.

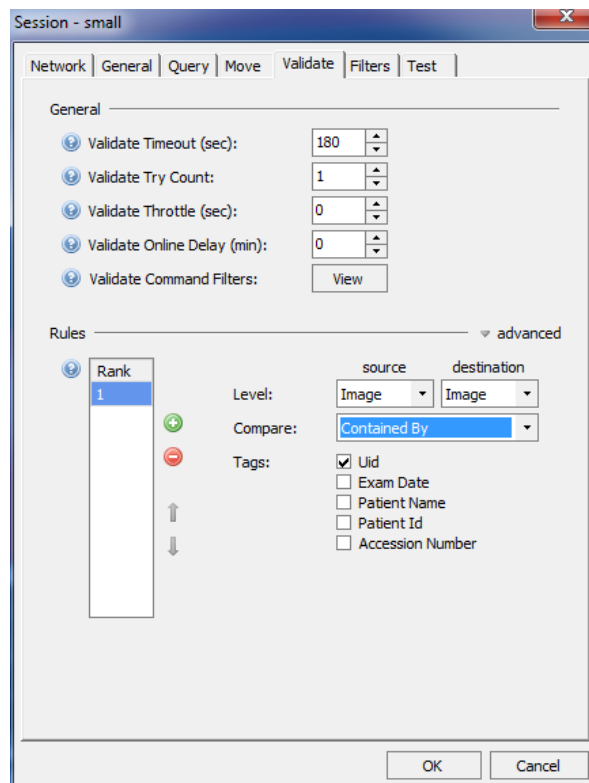


Figure 44 - Validating a migration where new images are added to the destination

9. I did not migrate SR (Structured Report) objects; how can I perform validation?

Sometimes you cannot migrate Structured Reports because a destination will not accept them; this is accomplished with a Drop Image filter. For additional info, see Appendix B - How to not migrate (i.e., drop) non-image DICOM objects.

However, if you do not migrate SR objects, then you have guaranteed that **Validation** will fail because the instance counts on the Source and Destination will be different. In that case we can configure a **Validation** rule to do the Source query at the Series or Image level and ignore the responses that are for SR SOP classes. Doing this will cause Exodus to count the source instances and compare them to the instances on the destination. You could do this same thing for any SOP class, of course; it does not have to be SR.

See the sequence of screens shown below for an example of how to set up this type of validation process:

Figure 45 - Create a Validation Rule to Compare Series Level to Study level

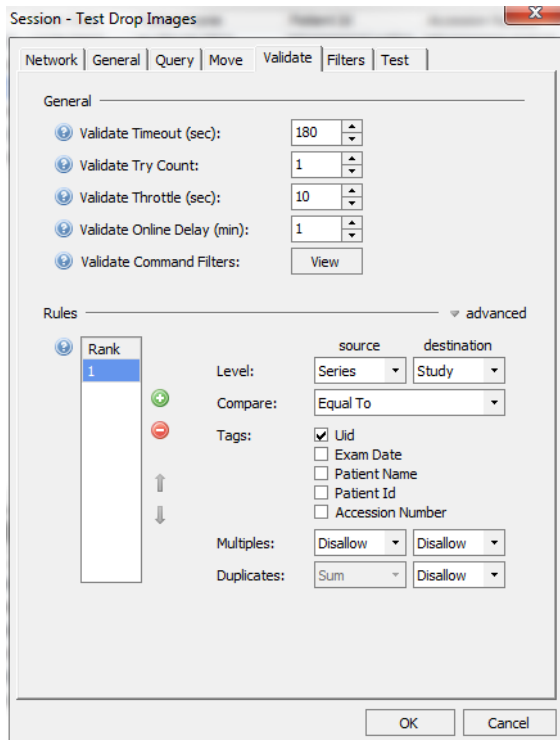


Figure 46 - Create a Filter to Drop SR Images




Figure 47 - Create a Validation Command Filter to insert Modality as a return tag

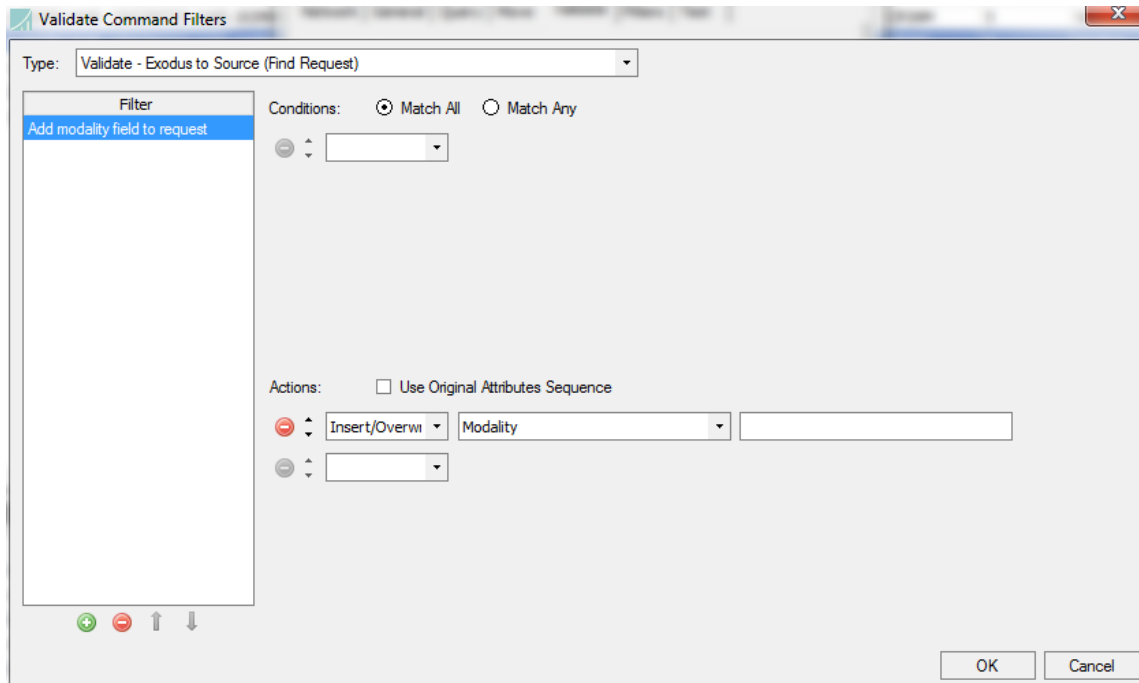
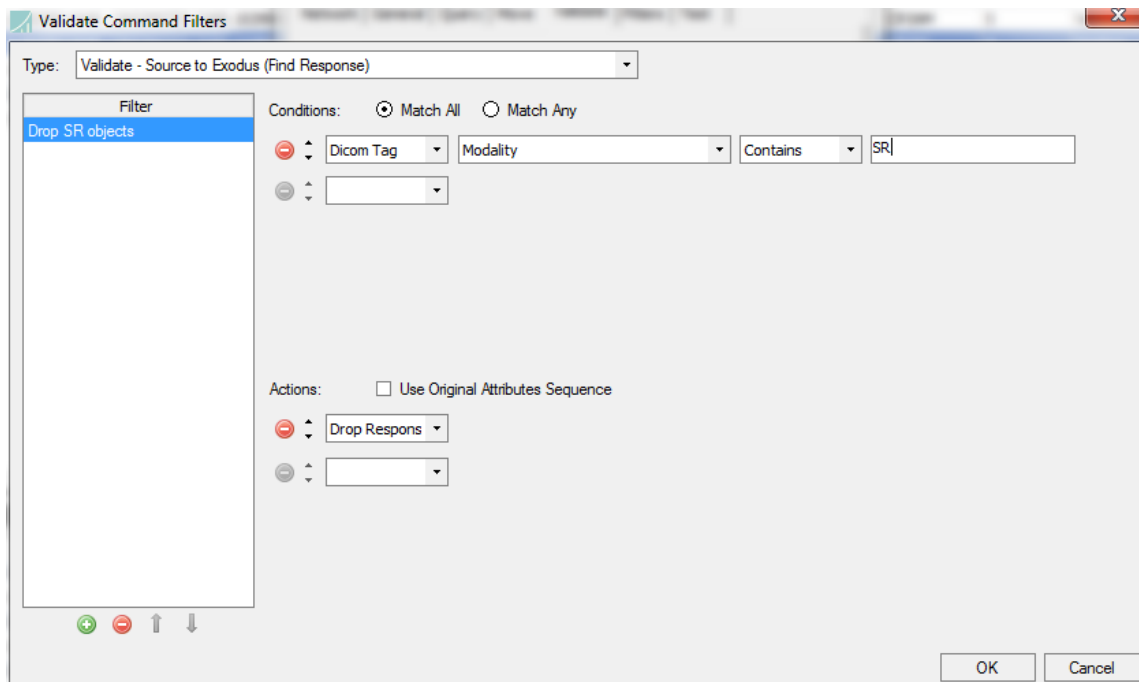


Figure 48 - Create a Validation Command Filter to Drop Responses containing SR in the modality tag



Note: You may need to Drop Responses based on SOP Class UID(s) instead of the Modality tag in some cases.

10. After querying the Source, I do not see any exams in the Exam Viewer; what should I do?

Is your source PACS configured properly to move to Exodus or directly to the Destination?

Are you sure that there are exams in the queried date range on the Source PACS?

Can you verify that the exams are accessible and retrievable on the Source PACS? Can you view them on the Source PACS in its image viewer? Can you push them off the Source PACS via a DICOM C-Store? Do the exams need to be loaded into the PACS cache from long term storage and the system is waiting for you to load a tape or DVD.

Use the **Query Client** (available from the Tools menu) to probe the Source PACS. Exodus, by default, queries per day using a date range of today-today e.g. if you are querying for January 1, 2007, then the query sent will be 20070101-20070101. Some PACS do not support this date range query format. You should try a single day query of 20070101 or a two-day range of 20070101-20070102. You can then reconfigure Exodus to use the preferred format on the Query Tab of the Session Properties.

11. Exodus was shut down for several days, will my Priors Segment catch-up?

Yes, if you have a Priors Segment and Exodus is stopped or shutdown, it will query for the missed days when the migration is restarted. You must have the Query checkbox checked for the extra queries to happen.

12. Where can I find more information on Regular Expressions used by Exodus?

Exodus uses .Net style regular expressions and more information can be found at:

[http://msdn.microsoft.com/en-us/library/az24scfc\(v=vs.90\).aspx](http://msdn.microsoft.com/en-us/library/az24scfc(v=vs.90).aspx)

13. My low priority Segments do not get run; what is wrong with my configuration?

If you have multiple Segment(s), the highest priority segment(s) will run to completion then the next highest priority segment(s). Segments with the same Priority setting will run a round robin fashion. This will allow you to set a high priority **Priors Segment** and guarantee that **Priors** will always be migrated ahead of the normal migration segments.

14. I want to use a Mapping List Action with private tags

To do this you must use a mapping list filter coupled with a Composer Action Filter. Your mapping list filter will have to update standard DICOM tags. These standard tags will be used as temporary place holders. The Composer Action can then copy the values from temporary place holders into the desired private tags.

For example, if you wanted to insert a private Patient ID into a private tag via a mapping file, you could have your mapping file overwrite tag 0010,1000 (Other patient Ids) which is a seldom used standard tag). From there the Composer Action could copy the value from tag 0010,1000 to the private tag. You could then use a Remove Action to delete the temporary tag. See the section on Composer Action for information about working with private tags.

15. How does Exodus respond to C-ECHO messages?

If you are using the **Image Routing** feature on the Filters tab, then Exodus will forward inbound C-ECHO-RQ messages to the Destination PACS configured for the current **Session**. C-ECHO-RSP messages will be relayed from the Destination PACS to the original requester. If **Image Routing** is configured to **On With Filters**, then filters will be applied to the C-ECHO messages, possibly resulting in unexpected results.

If **Image Routing** is not being used, then Exodus will not respond to C-ECHO-RQ messages.

16. How do I know when to use SQL Express, Standard, or better?

It can vary, but if you need to create a single Session with an exam count approaching 1,000,000 exams you will want to consider using a version of SQL more powerful than SQL Express. SQL Express is limited to allocating approximately 1.5 GB of memory. With large numbers of exams SQL Standard may need to allocate 5 or 6 GB of memory. Exodus might also allocate over 4 GB of memory occasionally. In these cases, you will also want to consider having 16 GB of RAM in your machine. If you are doing any compression or transcoding, you may want more memory.

If using a higher performance, paid version of SQL is not feasible, you might want to consider using multiple copies of Exodus on separate computers and then manually divide up the load between the separate migration servers.

17. How do I add custom prequalification segments?

There are several ways to prequalify your data after performing an assessment query. The first is to use the View Filters on the Exam Viewer. You can use Boolean operators on the data queried from the Source Pacs by Exodus. Once the View has been filtered you can add the exams in the View into their own segments or mark them as Ignored.

The second option is to create a Normal Segment encompassing the date range of exam dates you want to prequalify. Then, on the Advanced tab of the Segment dialog, you can enter a SQL where clause to refine the exams in the segment. For example, to only show exams that have an Accession Number that end in 'E9' you would define a selection string like this [AccessionNumber] like '%E9'

18. I have upgraded to version 4.0.12 or newer and I cannot activate my license?

If you are presented with the License Activation Dialog after installing or upgrading and you click the 'Activate' button and Exodus immediately crashes, you may be having an issue with the Microsoft Visual C++ Redistributable for Visual Studio 2017 x64 version. The redistributable should have been installed as part of the Exodus install. However, it may require certain windows updates to be installed. Try installing all available windows updates and that should fix the issue.

19. Help I'm stuck!

Email: support@laurelbridge.com

Or

Submit a request at: http://www.laurelbridge.com/request_support.html

Laurel Bridge Software, Inc.

Tel.: 302-453-0222

www.laurelbridge.com

Appendix E Source of Truth Discovery and Migration Process

This appendix provides a summary overview of the Exodus Migration Controller's typical migration process, as well as a description the **Source of Truth (SOT)** exam discovery and migration process.

In the SOT process a collection of DICOM exams, stored in an accessible location, is migrated to a new archive/PACS. Exodus uses a user or vendor supplied file that contains the known good exam information (i.e., the SOT data) that should be applied to each exam as it is moved.

1 Overview of the Typical Migration Process

Typically, Laurel Bridge Software's Exodus Migration Controller Software is used to assess, plan, perform, validate and report DICOM migration activities, and performs the main functions described in the brief overview included below. For additional information and details, see also Section 1.2, Migration Basics.

Summarized below are the standard steps involved in completing a typical DICOM data migration:

1.1 Query the Source Archive

During this initial phase, Exodus discovers ("learns" via DICOM C-Find requests) the contents of the source archive in preparation to analyze, plan, and perform the DICOM migration. There are two methods for querying the source PACS:

1. Using **DICOM Query/Retrieve (Q/R) services** typically employed by the Source Archive, or
2. Via a **file-import process** where Exodus can scan/traverse a file-store of DICOM datasets searching for DICOM exams.

The advantage of the Q/R method of discovery is that the Source Archive will interact with its internal database of study and patient information, and then during delivery of the exam to Exodus will apply appropriate database updates to the outbound DICOM study. For example, if a "John Doe" patient is scanned while in emergency room, then his images will be stored with the patient name of John Doe. Later, when the patient's name is identified, the database in the Source Archive is updated to reflect patient's actual name. Then when a DICOM Q/R is processed, the Source Archive will retrieve the images (still containing name of John Doe) but will deliver the study using the patient's real name located in the internal database.

Regardless of the exam discovery process above, Exodus will build an internal database of study and patient information that will be used to control and perform the DICOM migration.

1.2 Migrate the Exams

Using the Exodus Exam Table generated in the query step above, Exodus will facilitate the communication and migration of DICOM exams from Source Archive to Destination Archive. While doing so, Exodus can perform several functions, such as tag-morphing, transfer syntax (compression) modification, throttle throughput, and notification.

1.3 Validate the Exams

Once Exodus believes that a study has been successfully migrated, it will again query the Source Archive/PACS and the Destination Archive/PACS, and then compare the two responses to ensure all instances have in fact migrated.

1.4 Report Migration Status

Exodus provides built-in reporting mechanisms to report summary level, as well as detailed level, migration activities. Alternately, views of the Exodus exam table, which contains study and patient information, migration status (Queried, Moved, Validated), timestamps, etc., can be exported to Excel or CSV format for user-customizable reporting.

2 Source of Truth (SOT) – Process

2.1 Overview

A third type of exam discovery is available by using a vendor specified extract file as the “Source of Truth” (SOT) for certain dataset elements. In this case, using the SOT file provided by the vendor, Exodus will build its internal database (exam table) by storing information contained in the SOT table, and it will then use this modified table to facilitate and perform the exam migration.

During the migration, Exodus will:

- Locate the specific DICOM files obtained from the SOT table,
- Ingest the DICOM objects, as if they had been imported or sent (C-Moved) to Exodus;
- Apply study and patient demographic updates based upon data provided in the SOT table or based upon standard Exodus filtering mechanisms that may have been configured; and then
- Send or move the ‘corrected’ DICOM object to the Destination Archive.

These new SOT application features are an extension to the standard Exodus Migration Controller software and provide Query/Retrieve SCP access to the DICOM files found in the Source Archive. Exodus also provides query functionality for date and number of files per study as determined by the DirectoryScanner.exe application described below.

2.2 Technical Discussion

The Exodus Migration Controller processes a vendor provided Source of Truth (SOT) file that is used to identify studies, update demographics, and provide physical DICOM dataset locations. The alternative exam discovery process consists of three applications/processes identified below. These applications are run one after the other to prepare for the SOT migration.

2.3 PreprocessSot.exe Application

This application validates the vendor provided SOT file. The provided file is expected to have nine fields per line, defined as follows (‘|’-delimited):

```
patient name| patient ID| date of birth| study date| study time| sex|  
accession number| study instance UID (SUID)| image path.
```

Each field is validated for specific parameters including existence, length, and format. Some of the fields such as SUID must be unique for all entries. In some cases, field values will be generated if they are NULL. If any of the fields are invalid the validation process exits and the error reported must be resolved before continuing.

2.4 DirectoryScanner.exe Application

This application uses the validated SOT file generated by the PreprocessSot.exe application to scan the vendor's directories for study instance files. The application creates an additional scan result file that contains the image path, link to the SOT, total number of files found, total DICOM files found, total non-DICOM files found, AVI files found, JPG files found, and BMP files found. By default, the application scans each directory and uses the file extensions to determine the file type. If file extensions are not present or are unreliable, there is an advanced function that will attempt to read a DICOM header from each file processed in an attempt to verify that it is a DICOM (.dcm) file.

2.5 SourceOfTruthQRSCP.exe Application

This QRSCP application provides the C-Move and C-Find functionality (service) that is required by the Exodus Migration Controller by using the generated output from the PreprocessSot.exe and DirectoryScanner.exe applications.

The Exodus application controls the ongoing migration process and this QRSCP application provides the services required by Exodus so that the exams are moved as scheduled.

- end -