



Customer Release Notes

eRAD PACS
Version 9.0
Server Build 1

Release Date August 22, 2022

Table of Contents

New Features and Enhancements	4
1. Minimize cache footprint	4
2. Track data/processed/cache storage state and manage dead mounts for study data storage	4
3. Refactor database access in application code to be database-agnostic	4
4. Facilitate runtime server role selection	4
5. Stream server jit not creating raw files even if the format is explicitly requested	4
6. Upgrade java to current stable version	4
7. Upgrade apache/tomcat to the latest stable version	4
8. Upgrade mysql to the latest stable version	5
9. Upgrade DCMTK	5
10. Upgrade gwt to the latest stable version	5
11. Upgrade openssl to the latest stable version	5
12. Deleting study when study resides on multiple mounts in the data repository	5
13. UDI for v9 server	5
14. Provide a warning sign on the WL for studies that reside on multiple mounts	5
15. Forwarding study when study resides on multiple mounts in the data repository	5
16. Editing study when study resides on multiple mounts in the data repository	5
17. Editing/adding report and notes when study resides on multiple mounts in the data repository	6
18. Remove legacy unimplemented jsp-s	6
19. Web Services notification triggered at child did not get sent	6
20. Add Study Update to Web Service Device Message Triggers	6
21. Time based warning message incorrect	6
22. Study with invalid time zone offset value displays empty study date	6
23. Serialize (manage) cw3 thumbnail downloads on tech view page	6
24. Include list name in logs generated by actions	7
25. Warn admin when a study is acquired that might nullify the server license	7

26. Change default media creation engine to local	7
27. Handle lost SQL connections/reconnects from C/C++ more reliably	7
28. Remove configuration options for mandatory v9 features	7
29. Local cache usage support for registration	7
Resolved Issues	7
1. Configured device timeouts did not get picked up properly	7
2. Users can forward or delete studies that aren't on their access list	7
3. MySQL escaping problem in Viewer's previous/next study	8
4. User's forward targets can be ignored	8
5. Global restriction is ignored by Viewer's previous/next buttons	8
6. Users can delete their own .access list	8
7. The forward study process is misspelled in the epws.log	8
8. Page does not refresh after editing/creating DocType	8
9. Viewer log version information parsing wrong	8
10. Extra exception thrown by CacheStateToolbox	8
11. Messaging recipients' limits can be ignored	8
12. User Account page not displayed for support user	9
13. AE title keeps reverting back to PBUILDER	9
14. Curl doesn't exist where it should, actions can't run	9
15. Persistent native database connections not reconnecting after timeout	9
16. Herelod daemons stop processing	9
17. Possible class/jar file missing	9
18. Exportusers.sh won't function - java doesn't exist where it should	9
19. Taskd crashes in jni in libmysqlclient.so during registration	9
20. Default behavior in mysql 8 treats truncation as error	10
21. Cache blob creation task failing on a study with no image objects	10
22. UserManager fails to run separately	10
23. Clear cache from Tech View doesn't update processing status	10
24. Tech view's reprocess doesn't work on a server farm	10
25. Viewer gets the server name variable from server memory (quickfix)	10

New Features and Enhancements

1. Minimize cache footprint

All system components, including viewer streaming, web viewer, technologist view, etc., support the single compressed cache data format (cw3). The creation of data in other formats has been terminated.

2. Track data/processed/cache storage state and manage dead mounts for study data storage

Calls to the repository handler have been replaced with a middle layer that tracks the state of meta data and manages the data accordingly, reporting data location, creating folders, moving data, indicating when folders are inaccessible, etc. The repository handler's dirty file handling and resolving mechanism remains unchanged. See the updated Repository Handler manual for specific details.

3. Refactor database access in application code to be database-agnostic

Performance-critical calls to the database have been encapsulated in an abstraction layer so the database is not directly exposed to medsrv. In addition to providing a common interface, it allows the application to maintain persistent connections to the database.

4. Facilitate runtime server role selection

Servers can be assigned specific roles to play, including stream server, registration server, database server, application server and web server. The setting is defined in `~/etc/.role`. If no specific role is defined, all services are performed.

5. Stream server jit not creating raw files even if the format is explicitly requested

The common stream server code failed to generate raw files when explicitly requested. While this is irrelevant for v9 (because its stream server doesn't use raw files), the change was made to the common code base, which v9 does use.

6. Upgrade java to current stable version

Java has been upgraded to java-17-openjdk-17.0.3.0.7. The system uses the platform's version of Java.

7. Upgrade apache/tomcat to the latest stable version

Apache has been upgraded to httpd-2.4.37. Tomcat has been upgraded to version 9.0.63. The system uses a custom build of Tomcat but uses the platform's Apache.

8. Upgrade mysql to the latest stable version

REVERSIBILITY NOTICE: Once upgraded, the database is modified and no longer compatible with the previous version.

MySQL has been upgraded to version 8.0.26. The system uses the platform's version of MySQL.

9. Upgrade DCMTK

The DCMTK library has been updated to version 3.6.7.

10. Upgrade gwt to the latest stable version

GWT has been upgraded to version 2.9.0.

11. Upgrade openssl to the latest stable version

Openssl has been upgraded to version 1.1.1k. The system uses the platform's version of Openssl.

12. Deleting study when study resides on multiple mounts in the data repository

Studies that exist on multiple repositories (which are possible when a repository was not mounted at some point when the data was updated) cannot be deleted via the user interface or the system. Users are notified of this on the delete review page, and entries are inserted into the log files.

13. UDI for v9 server

The UDI value for version 9.0 has been updated to 0086699400025590. This value is displayed on the appropriate software identification pages.

14. Provide a warning sign on the WL for studies that reside on multiple mounts

The Partially Inaccessible column available to indicate when the study resides on multiple repository mounts. This column is hidden by default. Add it to your layout using the Edit Fields tool.

15. Forwarding study when study resides on multiple mounts in the data repository

Forwarding a study that resides on multiple mount points will result in an error. If initiated from the GUI, the user is notified. If initiated from a forward action, the request will be retried when the action runs again (in five minutes).

16. Editing study when study resides on multiple mounts in the data repository

Editing a study that resides on multiple mount points will result in an error. If initiated from the GUI, the user is notified. If initiated from an edit action, the request will be retried when the action runs again (in five minutes).

17. Editing/adding report and notes when study resides on multiple mounts in the data repository

Editing a report or report notes for a study residing on multiple mount points is not supported. If the condition exists, the report add/edit button and the note add/edit button are disabled in the patient folder.

18. Remove legacy unimplemented jsp-s

Java servlet functions retired or no longer in use in version 9 have been removed from the code base.

19. Web Services notification triggered at child did not get sent

Based on timing, an auto-correction message originating at a child server can jump ahead of the first object registration message, allowing third party devices to believe a study exists before it actually does. Auto-correction messages are suspended until the hub server registers at least one object.

20. Add Study Update to Web Service Device Message Triggers

Web services devices can be configured to receive an order update notification when the study data has been edited. The trigger is enabled when the Study Update setting in the Order Message Triggers section of the web services device edit page is checked. Update sends a notification on new object acquisition, any edit or object re-acquisition. Reindex sends a notification when a study gets reindexed by an admin or the system.

21. Time based warning message incorrect

The wording of notification message indicating the repository handler had to delete data even though the threshold wasn't crossed has been changed to more accurately reflect the cause of the problem.

22. Study with invalid time zone offset value displays empty study date

Objects containing a non-compliant time zone offset value ignore the bad data and present time values as recorded in the object.

23. Serialize (manage) cw3 thumbnail downloads on tech view page

Downloading CW3 images to the technologist view page and the web viewer need to be managed by the client. A maximum of four images are downloaded in parallel to avoid over loading the browser.

24. Include list name in logs generated by actions

Log entries, on the Logs page and in the oper_info log, containing details for events resulting from an action, except the Prefetch action, identify the worklist filter that matched the study.

25. Warn admin when a study is acquired that might nullify the server license

The server's license is checked against multiple events and data. When one of these is detected but not enough to invalidate the license, the system sends a message notification to administrators. Admins can contact eRAD support for details and ways to avoid a license exception.

26. Change default media creation engine to local

The default media creation engine defaults to the local MCS. This applies to new installs and upgrades.

27. Handle lost SQL connections/reconnects from C/C++ more reliably

When the underlying connection to the database is lost, the software transparently reconnects and retries the pending operation.

28. Remove configuration options for mandatory v9 features

Some features optional prior to version 9 are no longer optional. They are hard configured by default. The settings for these features have been removed from the GUI.

29. Local cache usage support for registration

The initial registration creates the compressed image files on the local cache repository, before adding them to the blob. This requires the creation of a local cache repository (~/.var/localcacache.repository)

Resolved Issues

1. Configured device timeouts did not get picked up properly

A change in the DCMKT toolkit required connection timeouts to be assigned earlier than they were. As a result, all but the first send request and all the receive requests used the built-in timeout value.

2. Users can forward or delete studies that aren't on their access list

Some operations could be copied from the browser's network panel and invoked from another browser by a user with different permissions, allowing users to perform unpermitted operations. The missing permission check has been applied.

3. MySQL escaping problem in Viewer's previous/next study

If the Institution Name value contains an apostrophe and the field is used in the filter criteria applied on the worklist, the open next/previous study command results in an exception due to an improperly formed query.

4. User's forward targets can be ignored

Some operations could be copied from the browser's network panel and invoked from another browser by a user with different permissions, allowing users to perform unpermitted operations. The missing permission check has been applied.

5. Global restriction is ignored by Viewer's previous/next buttons

The system ignored global restrictions when selecting the next/previous study. If the user does not have permission to view the study selected, he/she ends up with an invalid study error.

6. Users can delete their own .access list

Some operations could be copied from the browser's network panel and invoked from another browser by a user with different permissions, allowing users to perform unpermitted operations. The missing permission check has been applied.

7. The forward study process is misspelled in the epws.log

A typo existed in the term "ForwardStudy" in log file entries for a web services study forward event.

8. Page does not refresh after editing/creating DocType

Newly created document types failed to show up on the document type configuration page until after a browser page refresh.

9. Viewer log version information parsing wrong

The viewer version number in the uploaded viewer logs was incorrect because the data was taken from the wrong object.

10. Extra exception thrown by CacheStateToolbox

When a registered viewer device issues a cache state request to the server and the server does not have an explicitly defined prior study cache state setting, the parsing algorithm misinterprets the parsing results and registers an unnecessary exception in the log file.

11. Messaging recipients' limits can be ignored

Some operations could be copied from the browser's network panel and invoked from another browser by a user with different permissions, allowing users to perform unpermitted operations. The missing permission check has been applied.

12. User Account page not displayed for support user

The initial dotcom setup processed failed to include the server's self ID setting in the default configuration file. As a result, the support account was not recognized.

13. AE title keeps reverting back to PBUILDER

A reorganization of the component start up scripts broke the setup of the default pb-scp configuration file, resulting in appending the wrong defaults to the end of the configured settings which were then taken as the configured value.

14. Curl doesn't exist where it should, actions can't run

Actions failed to run because the path used to identify the curl script used the removed custom component. The path has been updated to use the OS-supplied curl tools.

15. Persistent native database connections not reconnecting after timeout

The persistent database connections would not reconnect if the connection was lost or timed out, resulting in retried attempts to register objects, among other incomplete database requests.

16. Herelod daemons stop processing

When attempting to acquire and register large numbers of objects in a short period of time, herelod processes failed to terminate cleanly, waiting unnecessarily on the release of conditional variables, resulting in failed registration tasks and dropped objects.

17. Possible class/jar file missing

After upgrading java, an incompatible JAX-WS file caused all web services commands to fail. Upgraded JAX-WS to version 2.3.5.

18. Exportusers.sh won't function - java doesn't exist where it should

The location of java has moved but the path variable using in multiple scripts, including the user account import and export tools, still pointed to the former location.

19. Taskd crashes in jni in libmysqlclient.so during registration

Persistent and non-persistent database connections would release the SQL library object when terminating, invalidating the persistent connection and causing unstable behavior in other threads.

20. Default behavior in mysql 8 treats truncation as error

The updated version of MySQL, using the carry-over settings, treats truncation as an error as opposed to truncating the value automatically. The settings have been updated to default to the previous truncation behavior.

21. Cache blob creation task failing on a study with no image objects

A study with no cache or processed data directories, e.g., a study with just a PbR object, encountered an exception when preparing the meta data because of a failure to examine the return code value.

22. UserManager fails to run separately

If the user manager runs on an independent server, login attempts fail because the software was not properly prepared to pass data between discrete objects correctly.

23. Clear cache from Tech View doesn't update processing status

Clearing cache from the Technologist view page deleted the data but failed to update the internal processing state value because the feature to track processing status across multiple servers was not yet implemented.

24. Tech view's reprocess doesn't work on a server farm

Reindexing a study fails when the ingestion and application components run on separate servers because the application server has no registration abilities. The registration request is submitted to one of the registration servers.

25. Viewer gets the server name variable from server memory (quickfix)

The stream server the viewer uses to download the data is defined in the herpa data but the early v9 viewer does not use the value yet. Until this is available, the server will leave the setting empty if it determines the stream service and the web service are running on the same server, forcing the viewer to fall back on its assumption they are the same. Note this solution only works when running stream and web services on the same server. When the services are separated, an updated viewer is required.