

☐ Minimize Risks during Your Millennium® Upgrade

White Paper



Minimize Risks during Your Millennium® Upgrade

White Paper

By David Hallberg

We have seen many clients experience problems after upgrading to version 2007.18, 2007.19 or 2010.01 of the Cerner Millennium® EMR system. If you have not yet upgraded, this white paper offers specific recommendations to help minimize the known risks to your production environment. It's best to execute them before the upgrade. If you have already upgraded, you may want to review these items and see how many of them you can do to improve your current stability.

In general, I've ranked my 14 recommendations in order of importance. The majority of them are MQ and message log issues. They apply whether you are on AIX, HP-UX or VMS (although VMS applies only through the 2007.19 upgrade since Cerner doesn't support this platform on 2010.01 and above). Here's a list of the issues that we detail in this white paper:

1. The hidden MQ queue
2. Java server stress
3. Citrix crashes
4. Exception queue transactions
5. Message log disk I/O increases
6. Confusing MQ log files
7. Extra message log writing
8. Too much extraneous logging
9. MQ information loss
10. Missing or incorrect security settings for a file
11. MQ running in mixed mode
12. Request class routing needs
13. Wrapping issues with message logs
14. Transaction database (TDB)

By performing the following critical tasks before an upgrade, I believe you'll have a much shorter troubleshooting time during the upgrade process.

- 1. The hidden MQ queue.** Set the MaxDepth of the CERN.ADMIN.REPLY queue to 15,000 if you are a two-node configuration and to at least 20,000 if you have more than two nodes on the application side of Millennium. Without this change, the queue will fill up and cause nodal MQ stability problems, and you will have an extremely high probability of experiencing one or more significant production issues after the upgrade. The problem that happens when this queue is full is that no new Millennium servers can be started. This would stop all batch cycling of servers and interactive cycling or stopping and starting of Millennium servers (executables) on the application nodes.

To clear the CERN.ADMIN.REPLY queue in uptime, you have to execute some steps that might not make a lot of sense. Typically, to delete all the messages/transactions from a queue, you simply go into Panther's Queues control, right-click on the queue name, and select *Clear Queue*. Unfortunately, this method does not work with the CERN.ADMIN.REPLY queue because the QCP server (SCP Entry ID 27, Queue Control Panel) maintains connection with the queue. You, therefore, need to go into Panther's Queues control, select the CERN.ADMIN.REPLY queue, right-click it, select *View Messages*, select all of the messages, click the *Delete* button, and, finally, verify the deletion.

If you do not have Panther and are using QCPView either from the back end or front end, you cannot perform a "purge admin.reply-all" command to clear the reply queue. Instead, you need to type the following commands:

- a. ***show queue admin.reply*** (to get the current depth)
- b. ***dir admin.reply***
- c. ***purge admin.reply 1 – (the current depth number, such as 14594)***
Please be patient. This can take several seconds or more.
- d. Run step A again to verify that there are no more messages/transactions in the queue. If there are, redo steps B and C until the queue is empty.
- e. ***exit***

2. **Java server stress.** SCP Entry ID 352 (Clinical Information) has five Java services running in a single Java Virtual Machine (JVM). From what I have seen, this is too many Java services in a single JVM. I recommend you split SCP 352 into two SCP Entry IDs. The parameters section of Clinical Information typically looks like this: *messaging clinical_event messagingAsync FamilyHistory documentation*. I suggest copying 352 to 376 (or the next available SCP Entry ID), leaving 352 with FamilyHistory and documentation in the parameters field. Give 376 a description such as Clinical Info Messaging with the following two entries in the parameters field: messaging and messagingAsync. Select another SCP entry, for instance 377, and copy 352 into that SCP entry and leave only clinical_event in the parameters. Delete all other parameters. We have seen with the 2010.x and 2011.x code a dramatic increase in the usage of this Java service. You would want to change the description field to something like Clinical Event for Java or Java Clinical Event.

The other Java tuning recommended for your system by Cerner Millennium System Settings (CMSS) or whichever upgrade guide you are using should be done to both SCP entries. If you avoid this tuning, you are likely to experience backlogs in MQ's messagingAsync queue, which will eventually back up the Java server (SCP 352) and slow down the clinical_event, FamilyHistory, documentation and messaging services so much that clinicians will find Millennium unusable. The messagingAsync MQ backlog will make the physician's inbox so slow that it too becomes unusable.

3. **Citrix crashes.** You may experience an increase in crashes after the upgrade. To determine how many crashes are normal for you, be sure to use the Millennium crash collector before you upgrade. I would suggest having at least a couple weeks of information from the Lights On Network to determine what normal is for you.

Several sites have had crashes fill up their Citrix server drives, which can have a significant downstream impact on an organization. Watch carefully for how many dump files are created on the Citrix servers. You would want to set up a scheduled task or batch job to delete the dump files on at least a daily basis. You might even have to delete the files more frequently to prevent the Citrix drives from filling up or running out of space.

- 4. Exception queue transactions.** Consider disabling MQ Auto Replay during the upgrade. If transactions are being sent to the MQ Exception queue because of build issues or other environmental issues, they are likely to fail again when Panther replays them, adding more strain to MQ. Once all build issues and environmental issues are resolved, turn the MQ Auto Replay back on.
- 5. Message log disk I/O increases.** Error log writing often increases after an upgrade, so I suggest you determine your current wrap times. Which servers have mlg files wrapping in less than a day? Which servers have mlg files wrapping in less than 10 days? If you know what normal is before the upgrade, you can tell what is not normal after the upgrade and can log points to resolve your issues.
- 6. Confusing MQ log files.** Monitor MQ .FDC and log files in /var/mqm/errors on each application node. For some sites, we have found .FDC files that were months and years old. These add confusion when there are MQ and throughput issues. If an .FDC file is more than two weeks old and you are not actively engaging your support team to resolve the issue, it is just consuming disk space. Please delete these old .FDC files.
- 7. Extra message log writing.** We have seen several servers start logging unnecessary messages for certain events after upgrades. I suggest you consider putting the following events in messages.suppress in \$cer_mgr:
 - a. SRV_UnpackFailure. Because of the potential risk for troubleshooting SCP Entry ID 330 (PM Reg), you may want to verify that there are no PM Reg server issues before suppressing this event.
 - b. CRM_ClientTimeZone.
- 8. Too much extraneous logging.** Set the following servers to LogLevel 0 to reduce extraneous logging.

SCP Description	SCP Entry ID
Search Server	221
MDI Order Server	235

9. **MQ information loss.** Set the following servers to LogLevel 1. Otherwise, you will miss MQ issues these servers can experience.

SCP Description	SCP Entry ID
CPM Application Authorization	50
CPM Script	51
CPM Code Cache Manager	52
CPM Script Private	53
CPM Async Script	54
CPM Process	55
CPM Script Batch	56
CPM Script Batch Priv	57
CPM Script (Read Only)	58
Entity Lock Server	59
CPM Service Resource Security	60
CPM Script001	61
CPM Script002	62
CPM Script003	63
CPM Script004	64
CPM Script005	65
CPM HNAM Agent	66
CPM Async Script Class 001	67
CPM Script008	68
CPM Generic Cache Manager	69
CPM Audit Server	70
CPM Audit Outbound Server	71
CPM Notify Server	72
Time Repository Server	73
CpmScript Report (Read Only)	74
User Services	75
CPM ESI Server	250
FSI ESO Server Load #1	280
FSI ESO Server Load #2	281
FSI ESO Server Load #3	282
FSI ESO Server Load #4	283
FSI ESO Server Load #5	284
FSI ESO Server Batch #1	285

10. **Missing or incorrect security settings for a file.** Please verify that the /cerner/mgr/workflow-orders-definitions.xml exists (or for VMS clients, set def cer_mgr). Verify the ownership and permissions on the file as well.

- 11. MQ running in mixed mode.** I would not run an MQ 6 client with an MQ 5.x back end. Upgrade the application nodes to MQ 6 when Millennium installs the MQ 6 client. It may involve some downtime, but it's worth the scheduled interruption in order to produce a stable production environment.
- 12. Request class routing needs.** The script PM_SCH_GET_PERSONS should be request class routed to a CPM Script server clone whose kill-time property is 5 minutes. Additional instances of the server you move this request to will probably be needed for the go-live volumes.
- 13. Wrapping issues with message logs.** Increase the mlg max records for cmb_0000, cmb_0027, cmb_0029, cmb_0033 and mqualert to at least 16,384. Please verify that you have enough file system space in \$cer_log before increasing these. You will want the information held in these files to assist in troubleshooting post-upgrade issues, especially if you upgraded MQ as well as Millennium.
- 14. Transaction database (TDB).** The TDB is a flat file found on each Millennium application node that tells the executables (application servers) what transactions they can process and if the transactions are Request/Reply (RR), Reliable Delivery Message (RDM) or Remote Procedure Call (RPC). Let me start with some background on these three types:
 - With RR transactions, the clinical information remains intact even if the transactions are lost or discarded. For instance, using CPM Script to retrieve or read data from the database is an RR transaction. If the transaction is lost, the person asking for the information just has to ask for the information again.
 - An RDM transaction is like an order. If the client crashes or the application executable or application node fails, there has to be a mechanism in place to recover or replay that transaction when the system is fully functional again.
 - The RPC transaction was added with Millennium 2007.x. It allows the Java servers and Java client applications to process transactions originating from both the local hardware and remote calls. It was the first extension to the Common Request Model in quite some time.

Until recently, the TDB was a connection-based service. PowerChart or the front-end application connected directly to the TDB's port. The TDB is now a Shared Service. The TDB server communicates with Citrix servers and fat clients over a network port that it allocates when it starts up. Once the communication port is selected or opened, the TDB server tells the Millennium Service Manager on the application node which port it will be listening to for incoming transactions (requests). The Citrix servers and fat clients must be on the same code level as the application node so that when PowerChart or any front-end Millennium application starts up, it will connect to the Service Manager and be told what port to use to connect with the TDB.

By performing this handful of critical tasks before an upgrade, you'll have a much shorter troubleshooting time during the upgrade process.

About the author: *David Hallberg is a systems analyst at Softek Solutions, Inc., the developer of OnTrack Panther and OnTrack Services, which provide third-party monitoring, management and auditing/assurance of the performance of the Cerner Millennium® EMR system. Hallberg has worked with numerous organizations in the United States and around the world to help them implement Millennium and keep it running fast and reliably.*