

Table of Contents

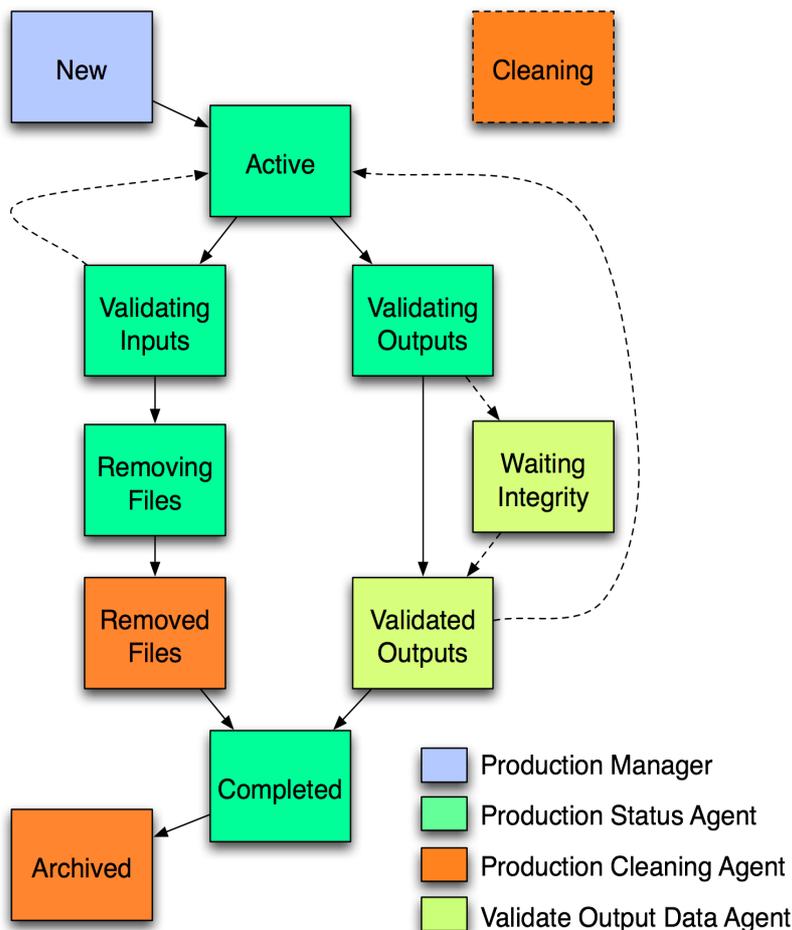
Production Status Machine.....	1
Overview of the production status machine.....	1
New.....	2
Active.....	2
Validating Outputs.....	2
Validating Inputs.....	2
Waiting Integrity.....	2
Validated Outputs.....	2
Removing Files.....	2
Removed Files.....	3
Completed.....	3
Archived.....	3
Special States.....	3
Completing.....	3
Flush.....	3
Cleaning.....	4

Production Status Machine

The purpose of this page is to describe the various states a production can be in and the components controlling each status transition.

Proposal for a new production status machine

Overview of the production status machine



The colour of the boxes in the above figure reflects the component managing productions in the given status. The transition from one box to the other is managed by the component that has the appropriate incoming arrow e.g. both Active to Validating Outputs and Validated Outputs to Active transitions are managed by the Production Status Agent.

The above status machine only applies to simulation requests that have one 'Used' production and one unused production associated with them. These are assumed to be the MC input and Merging productions respectively. Any request or subrequest with more than two associated productions are ignored by the Production Status Agent.

Omitted from the above diagram are productions in the Stopped status as this normally reflects a known problem with the production, preventing any progress.

A more detailed description of each status will now be given below.

New

Productions in New status have been created by the Production Manager and are ready to be started. The transition from New to Active can occur 'by hand' or automatically if using the request page interface.

Active

The Active status means that the production can be extended and jobs submitted. Shifters should be concerned only with the productions in Active status as they will remain there until enough events for the associated request appear in the Bookkeeping.

Validating Outputs

The Production Status Agent determines from the request database which productions (designated 'Used') have produced enough events compared to what has been requested. The list of merging productions with enough BK events are then updated to this status. The Validate Output agent then starts to process the output data of the merging productions and reports any problematics to the integrity DB. If actions are pending then the production is moved to the Waiting Integrity status. After the resolution of any pending integrity DB actions (or if none are required) the production then moves to the Validated Outputs status.

Validating Inputs

For each merging production updated to Validating Outputs the corresponding MC input production is updated to Validating Inputs at the same time. The MC productions remain in this status until the outputs of the merging production have been validated. At this point the total number of events of the merging production is rechecked, if some files were lost during integrity checking the MC production may return to the Active status (at the same time as the associated merging production) for the shifter to extend. Otherwise the transition to Removing Files is made such that any leftover files can be cleaned.

Waiting Integrity

Productions in the Waiting Integrity status have reported problematic files to the integrity DB and are awaiting the resolution of data management operations to recover some files. As soon as the problematic cases are resolved (at worst by removing the files) the production is moved to the Validated Output status.

Validated Outputs

Merging productions in the Validated Outputs status are rechecked for the total number of produced events in the Bookkeeping. If files were lost during the integrity checking phase the production (and associated MC production) are returned to the Active status. If sufficient events remain in the Bookkeeping then no more jobs have to be run. For 'flat' requests the merging production is set to Completed status and the associated MC production to Removing Files. The request for these productions is updated to Done status by the Production Status agent at this point. For requests with subrequests the same transitions occur only when all subrequests are in the Validated Outputs status. The Production Status Agent knows to skip over any problematic subrequests (for Active parent requests) that have zero BK events available.

Removing Files

The Removing Files status is set by the Production Status Agent to trigger the cleaning of any unmerged files from the MC input production. The Production Cleaning agent processes the productions in this status.

Removed Files

The Production Cleaning agent polls for MC productions in the Removing Files status and cleans any leftover outputs. As soon as this operation is finished the status is updated to Removed Outputs.

Completed

Productions in the Completed status have nothing further to produce. MC input productions reach this status after any leftover files are removed. Merging productions are put into Completed status whenever the requested amount of events are present in the Bookkeeping.

Productions in 'Completed' are only considered by the TransformationCleaningAgent, that will trigger archival. The archival happens after a grace period of 7 days. All production types can be cleaned.

Archived

This is a terminal state. The only remaining transition after Completed is when the production jobs are removed from the WMS / Production DB etc. Productions associated to simulation requests that are in this status will not have any job metadata in the production monitoring page.

Special States

Completing

A production is set to 'Completing' when it is derived by another production. In this case, the derived production contains a copy of what is in the TransformationFiles table of the original one, excluding files with status 'Unused'. This means that deriving a production is a safe operation, as long as there are no running jobs, nor pending requests for file updates on the transformation DB from jobs in 'Completed' status.

Each production in 'Completing' status executes:

1. Bookkeeping Query: executed by the BookkeepingWatchAgent. **False**
2. Tasks Creation: executed by the TransformationAgent. **True** (for MC: the MCExtensionAgent only extend 'Active' productions)
3. Tasks Submission and Monitoring: executed by the WorkflowTaskAgent(s) (inherits from TaskManagerAgentBase). **True**
4. Data Recovery: executed by the DataRecoveryAgent. **True**

Flush

A production can be set to 'Flush' manually. It is used to create tasks that otherwise would not be created (e.g. create the "last" merging tasks, when not enough files are available, and probably never will)

1. Bookkeeping Query: executed by the BookkeepingWatchAgent. **False**
2. Tasks Creation: executed by the TransformationAgent. **True**
3. Tasks Submission and Monitoring: executed by the WorkflowTaskAgent(s) (inherits from TaskManagerAgentBase). **False**

4. Data Recovery: executed by the DataRecoveryAgent. **False**

Cleaning

Productions can be set to Cleaning status from any of the above statuses. This is a 'by hand' operation and reflects a terminal problem that has been understood and implies any outputs generated by the production should be removed from all catalogs and storage elements.

Productions in 'Cleaning' are only considered by the TransformationCleaningAgent, that will trigger the real clean. All production types can be cleaned.

This topic: LHCb > ProductionStatusMachine

Topic revision: r5 - 2011-05-02 - FedericoStagni



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback