Gratia Buffering Points.

Overview

The Gratia accounting data chain has several points where data are buffered in the event that the upstream receiver is offline.
All the specific information below is current and correct as of the 1.04 release of the collector and probes.

Probes

The common probe code saves a backup of each record under the directory labeled, “WorkingFolder” in ProbeConfig before an upload is attempted. In the event of a successful upload the backup file is removed. In the event of a failed backup save (for example the partition containing this area is full), upload is still attempted. If upload fails, the responsibility rests with the individual probe to ensure the data are not lost.

During the next probe run, backups retained due to failed upload attempts will again be uploaded, with the file being removed on successful upload and retained on failure.

Caveats

- Currently, no individual probes are known to preserve the data in the case that backup **and** upload failed. A full disk therefore would lead to data loss.

Configuration items

Location	Item	Details	Default
ProbeConfig	WorkFolder	Location of temporary files	$VDT_LOCATION/gratia/var/tmp
ProbeConfig	MaxPendingFiles	New records will not be sent if there are more than this many unsent records.	100000

Collector

The collector has several servlets that listen for incoming probe data; their sole function is to write those data into a queue area on disk where it will be picked up by one or more threads which then process those data and enter it in the database. There is a configurable maximum threshold: if a thread's queue of waiting files exceeds this limit an instruction will be sent to the servlets to stop receiving new records. The probes will then act as if the collector is down.

Before attempting to process the file in the thread queue, it is written to a backup directory `old-<datespec>`.

Caveats

- In the event of being unable to write the incoming record to the thread queue area, the servlet will reject the record. The probe will retain the record just as if the collector was down.
- In the event of a failure writing the record to the DB: if the connection has failed, the record

will be retained and the listener will wait for connection restoration. If there is some other transaction error (SQL error, etc, etc), the XML record will be written to the DupRecord table for manual handling by an administrator. If the record cannot be written to the DupRecord table then the only remaining record is in `old-<datespec>`.

- Nothing is done if the incoming file cannot be written to `old-<datespec>`. If this happens *and* there is a subsequent failure to process the records to the DB, records from that file would be lost. This would obviously be a very rare occurrence, but can be prevented by a minor change to the collector error handling code.

Configuration items

Location	Item	Details	Default
service-configuration.properties	monitor.q.size	If set, monitor the thread queues and deactivate servlets if max.q.size is exceeded.	1
service-configuration.properties	max.q.size	If monitor.q.size is set, servlets will be deactivated if thread queue exceeds this number. Servlets will be activated only when thread queue drops below 80% of max.q.size.	100000