the globus alliance
www.globus.org

GRAM4 Architecture

Charles Bacon
University of Chicago

the globus alliance
www.globus.org

GT4 WS GRAM Architecture

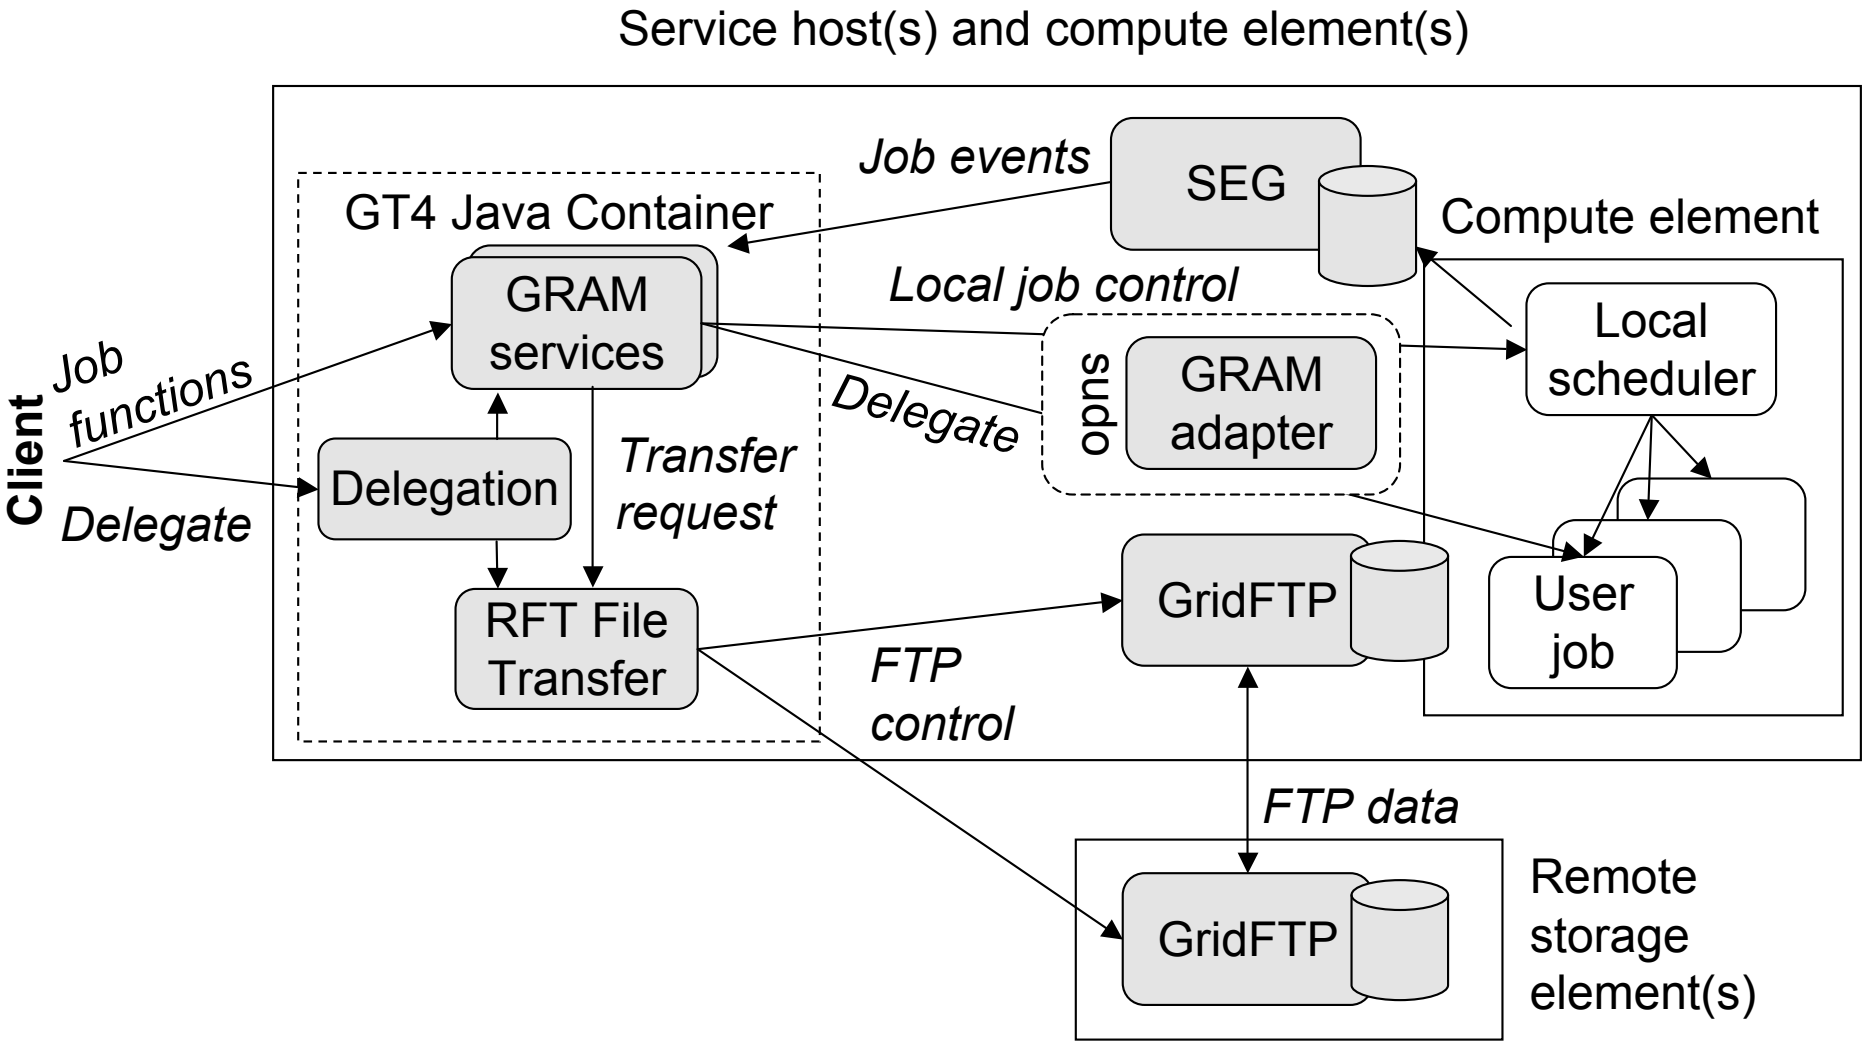

Data Staging

- GRAM uses an RFT service to perform data staging
- RFT uses 3rd party GridFTP transfers to stage the data in or out
- Additionally, RFT performs reliable cleanup of specified files or directories

File System Mapping

- The filesystem layout for your storage nodes may differ from the layout exposed to compute nodes
- The default map is that / maps to /, which works if your GridFTP server and the nodes have the same view of fileystems
- Otherwise, you might want to map /home to /exports/cluster1/home or the like

Compute Node Local Storage

- Your cluster may have some local space setup for temporary use during computation
- Set this value in etc/gram-service-Scheduler/jndi-config.xml under the “scratchDirectory” parameter
- Will be exposed to users as $GLOBUS_SCRATCH_DIR

Job Resource Limit

- GRAM limits its resource usage to 1000 simultaneous job resources by default
- You may increase/decrease this limit as appropriate
  - setup-gram-service-common –max-job-limit="150"

Job Monitoring

- GRAM uses a Scheduler Event Generator (SEG) to keep track of job status
- Parses logfiles kept by the scheduler
  - Allows for restartability if the container fails
  - Higher performance than polling the scheduler
- The location of the logfiles will be picked up at installation time by looking for variables like PBS_HOME

Scheduler Logfile Locations

- Each scheduler also has a file under etc/globus-<scheduler>.conf
- This contains the log_file=/path/to/log map for that scheduler
- You may update this if your scheduler keeps logs in a different location
- The logs should be readable by the globus user