

Gaia Archive operations for DR1

J. González-Núñez, J. Salgado, R. Gutiérrez-Sánchez, J.C. Segovia, J. Durán, C. Arviset, R. Alvarez, R. Gil, E. Anglada, J. Bakkers, U. Lammers

ESA Science Data Centre

23/10/2016

Issue/Revision: 1.0

Reference: Gaia Archive – VO inside

Status: Issued

ESA UNCLASSIFIED - Releasable to the Public



gaia

European Space Agency

14th September 2016

First release: 14 September 2016

Subject to successful validation, the catalogue will be consisting of:

- Positions (α , δ) and G magnitudes for all stars with acceptable formal standard errors on positions. Positions and individual uncertainties are computed using a generic prior and Bayes' rule (detailed description in "[Gaia astrometry for stars with too few observations. A Bayesian approach](#)"). For this release, it is assumed that at least 90% of the sky can be covered.
- At the beginning of the routine phase, a special scanning mode repeatedly covering the ecliptic poles on every spin was executed for calibration purposes. Photometric data of RR Lyrae and Cepheid variable stars including these high-cadence measurements will be released.
- The five-parameter astrometric solution - positions, parallaxes, and proper motions - for stars in common between the Tycho-2 Catalogue and Gaia will be released. The catalogue is based on the *Tycho-Gaia Astrometric Solution* ([Image of the Week with short TGAS description](#); [paper with a more detailed description](#); [paper describing theory and background](#); [paper describing quasar extension](#)).





gaia

→ THE BILLION STAR SURVEYOR


GAIA DATA RELEASE 1

ESAC

European Space Astronomy Centre
Madrid, Spain

14 September 2016

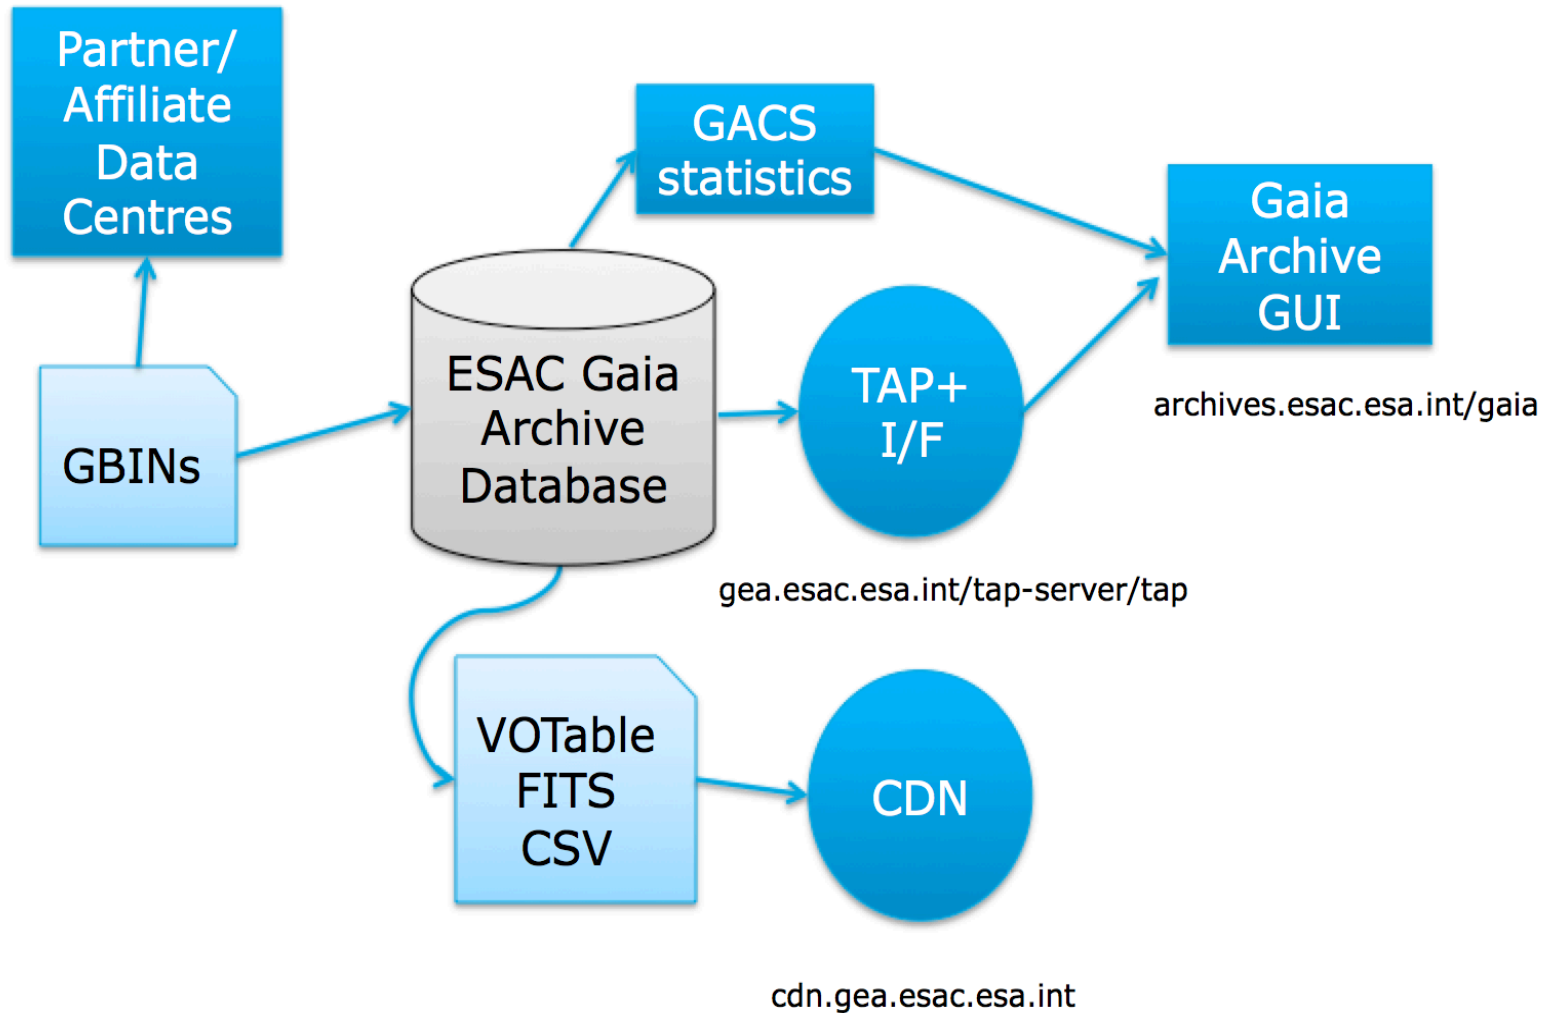
11:30 CEST



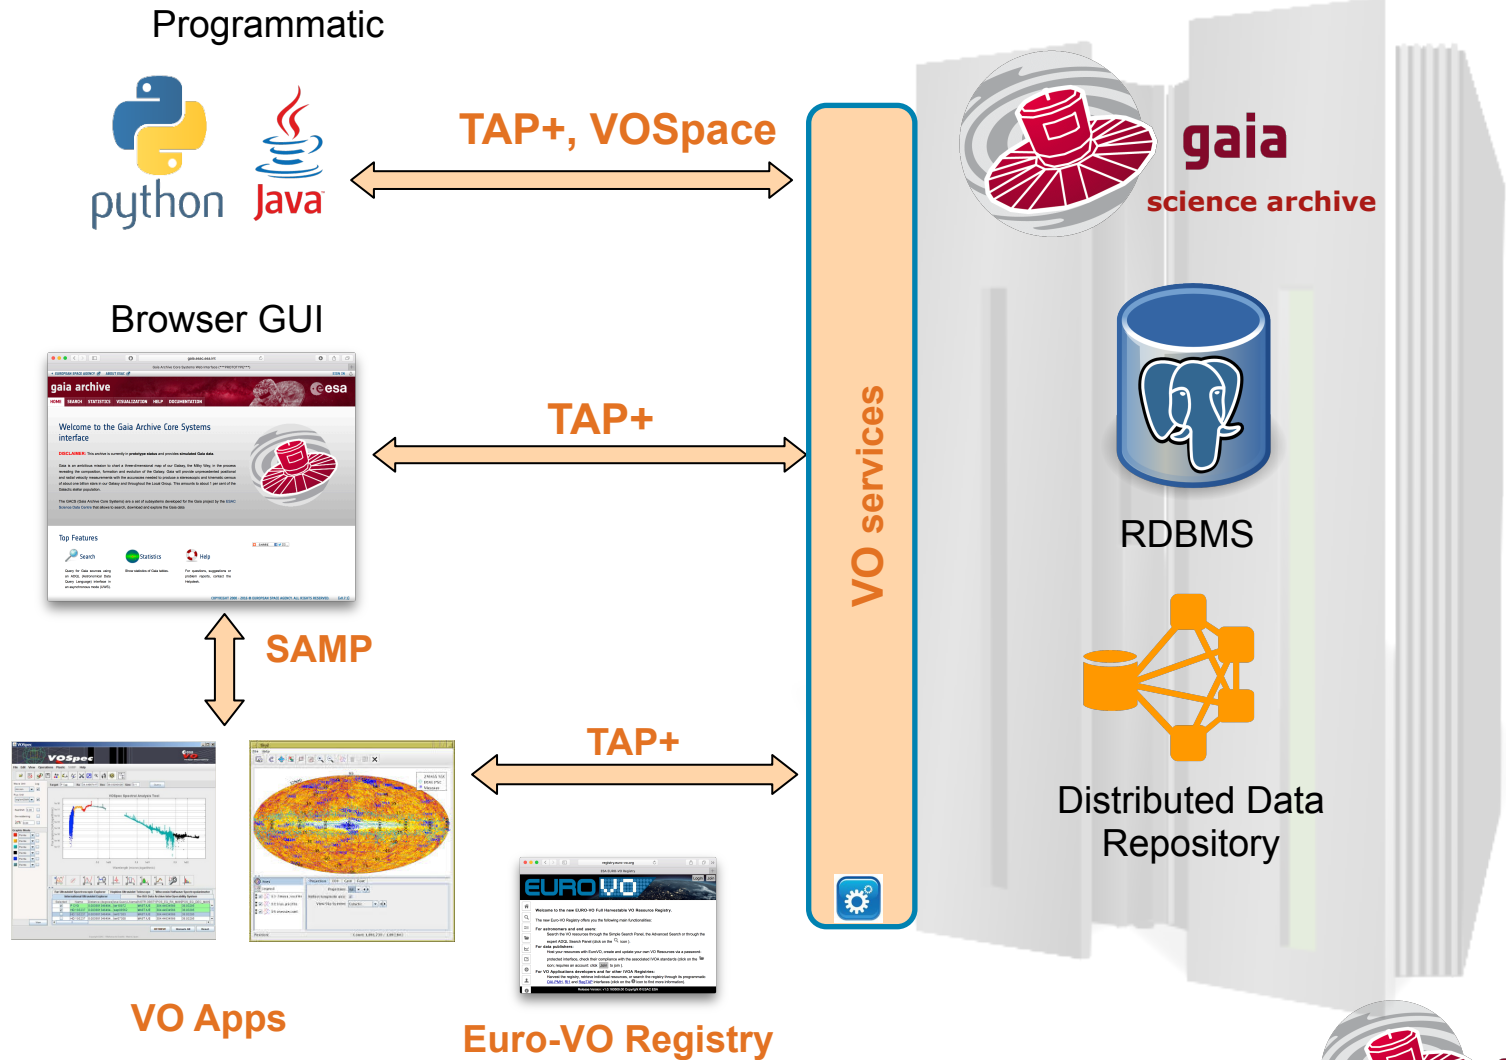
High level scalability measures

- Replication of data with Associate & Affiliate data centres
 - AIP, ARI, ASDC, CDS, IRSA, GAVO
- Bulk download infrastructure based on a Content Delivery Network

Ops for Gaia DR1



Ops for Gaia DR1



Scaling up S*AP vs scaling up TAP

- Potentially large processing time for requests
- Requests can be quite CPU intensive, both front-end and DB sides

Scaling up TAP+

- Persistent uploads
- Server-side crossmatches

Cloud vs dedicated HW

- Dedicated HW can be selected and configured in ways Cloud of virtualized infrastructures can't
- But it sets limits to scalability

Infrastructure dimensioning

- Conservative engineering considering worst case scenarios
- High performance hardware
- Measures to ensure graceful service degradation if capacity is exceeded

Stress test campaign

- Stress test plan defined covering all server-side elements of the Gaia Archive developed at ESAC (GACS), ie. TAP+ interface
- Built with support of IT support team (SITU) & vendors engineering teams support.
- Executed in 4 iterations (IT1, IT2, IT3, IT4) with an increasing level of systems monitoring
 - Full system performance analysis at each Iti
 - Incremental system infrastructure implemented to IT(i+1)
- Tested high data volume scenarios, with good performance up to network limits
- Tested high CPU consumption scenarios, with efficient use of archive resources (up to 80% occupation factor)

Resource limits

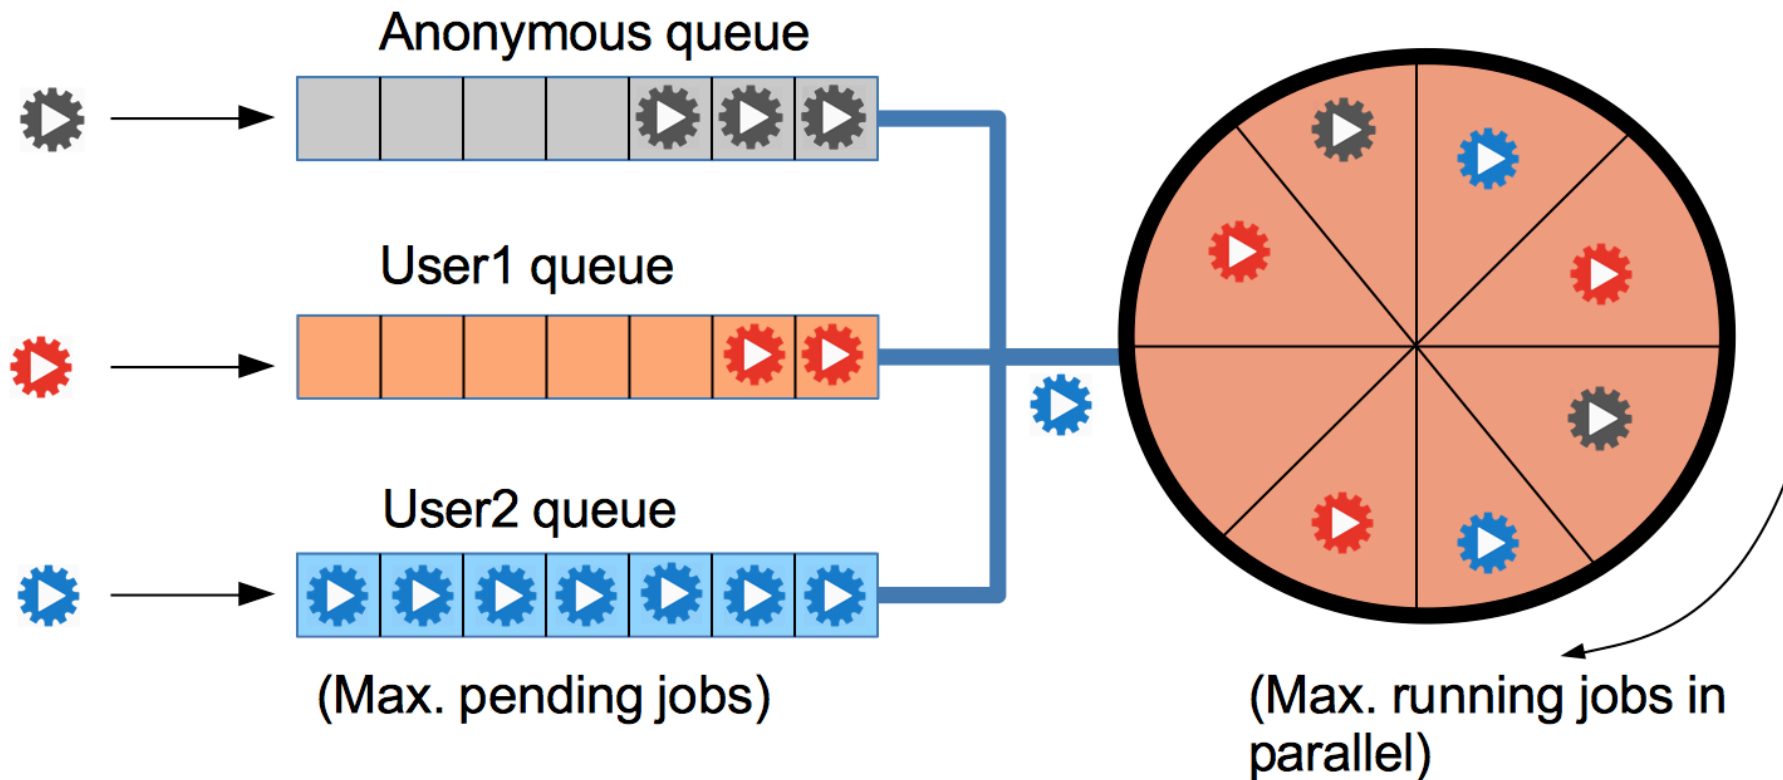
	Sync rows	Async rows	Sync time	Async time	DB space	Job results
Anonymous	100K	100K	1min	30min	none	N/A (np)
Registered	unlimited	unlimited	1min*	30min*	1GB*	1GB*

(np) not persistent

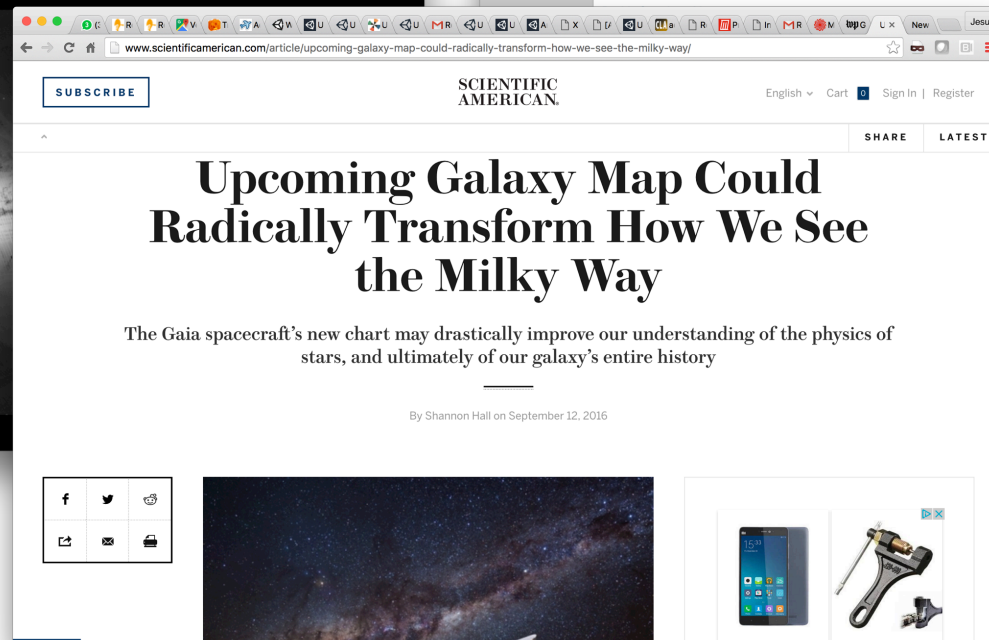
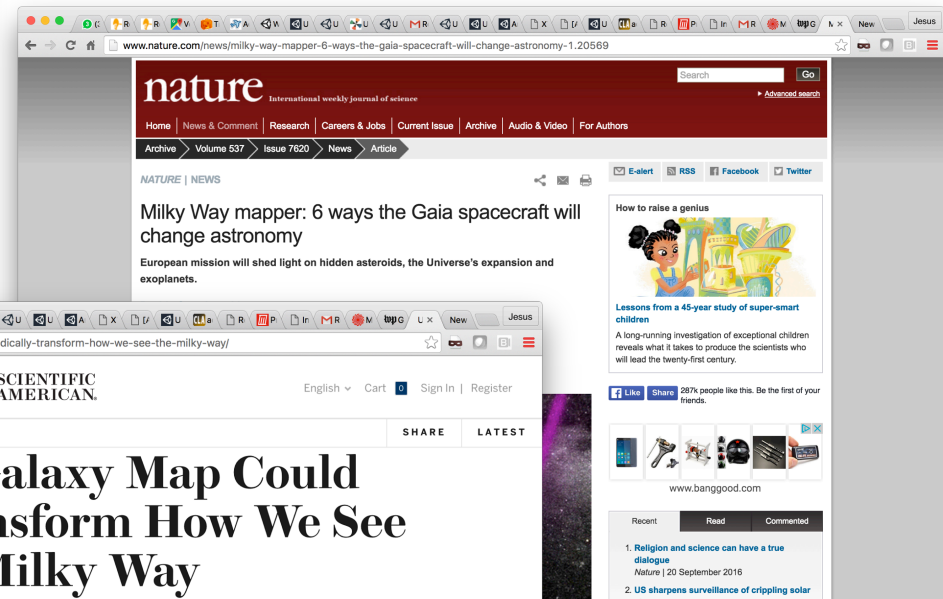
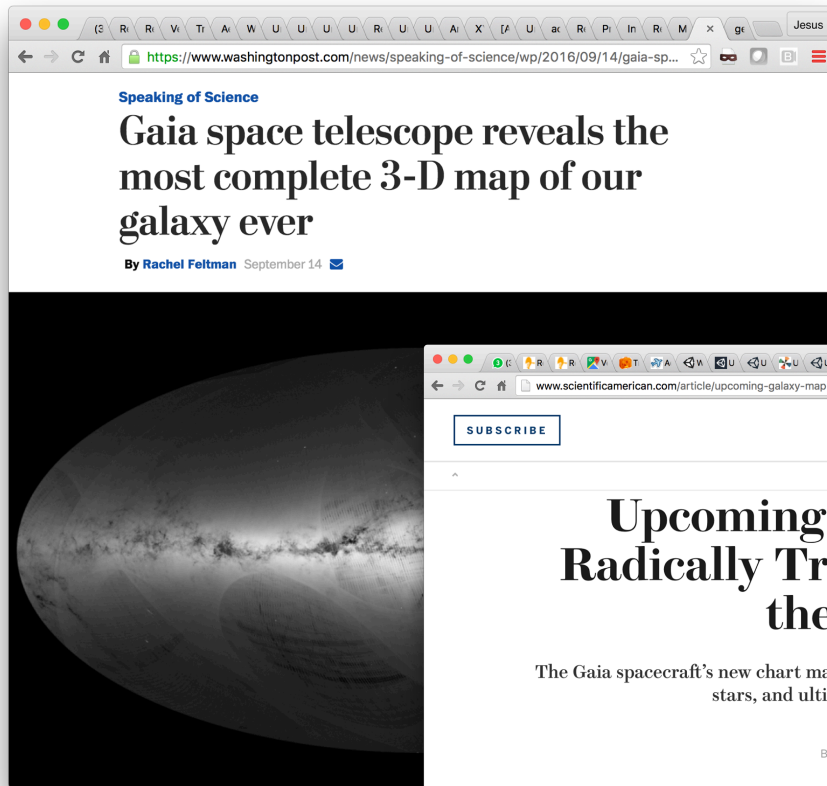
* Can be upgraded on demand




Job Scheduling



Ops for Gaia DR1



Archive UI usage: First 24 hours



UI Usage


- Usage sessions: **12,005**
- # of users: **10,959**

Ops for Gaia DR1



File CDN download: First 15 days

Bandwidth



File download

- Total volume downloaded: **73 TB**
- # download requests: **9 millions**



gaia

European Space Agency

TAP Interface: First 15 days

TAP Interface

- **Number of queries**
 - Synchronous: **174,807**
 - Asynchronous: **95,090**