



RIPE
NCC

Database Update

Denis Walker

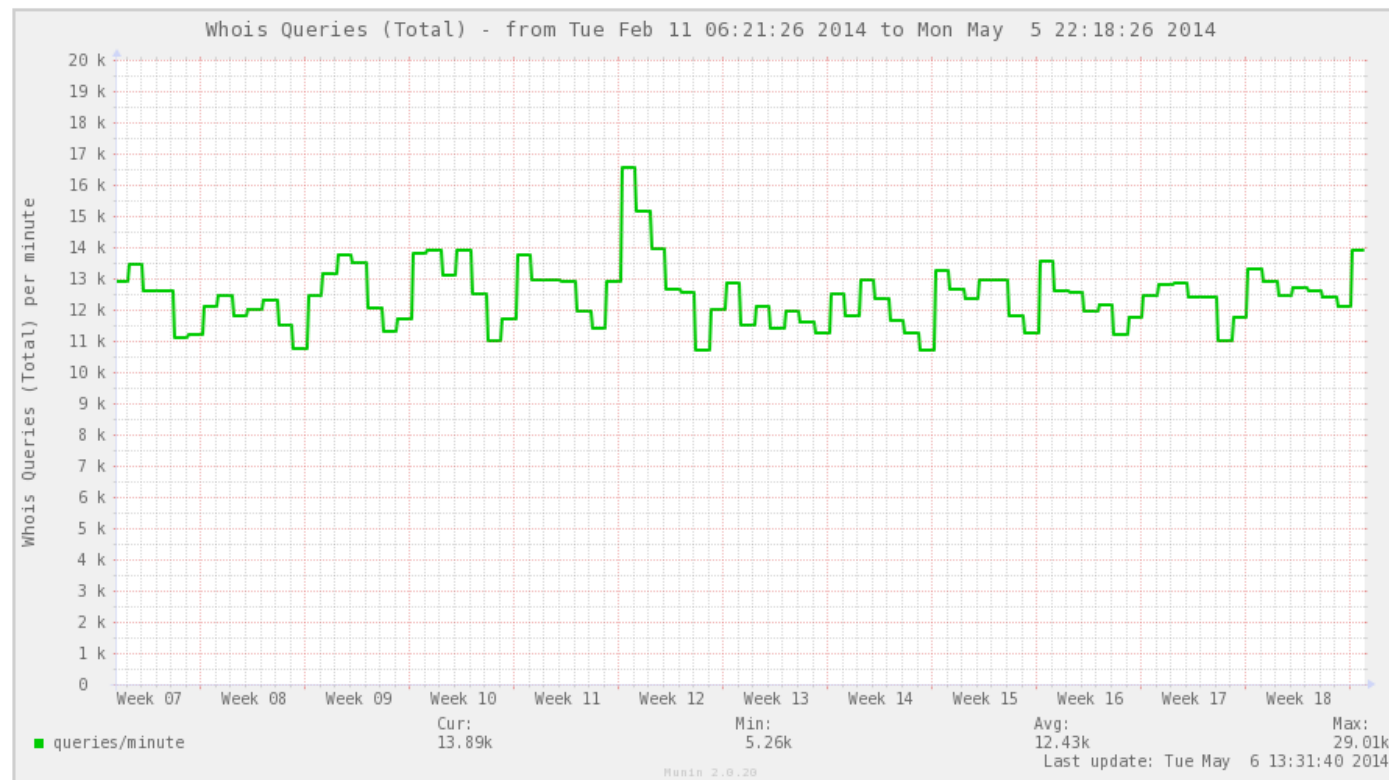
Business Analyst

and

Edward Shryane

Senior Software Engineer

- 2,700 unit tests
- 2,900 end-to-end and integration tests
- Operational stats:
 - <https://www.ripe.net/data-tools/stats/ripe-database>
 - about 206 queries per second



- Issues now tracked on GitHub
 - <https://github.com/RIPE-NCC/whois/issues>
- New issues can be raised by users
- Code base also on GitHub
 - <https://github.com/RIPE-NCC/whois/>
- Users are continuing to contribute to the code
- If you have an idea, we can help you technically

- Releases made to Release Candidate (RC) Test Environment for at least 2 weeks
- Users can test against this with live copy of data
- For major changes will also release draft documentation with RC deployment
- Then deploy to production (close down RC)
- Is this working for you? How can we improve?
- RC is a very lonely database...it's missing you

- Continued working on RDAP, the Registration Data Access Protocol – more later
- Added “via” attributes in **aut-num**
- Member resources are covered by “abuse-c:”
- Provided wizard for adding “abuse-c:” to PI objects
- Currently 33% of IPv4 and 44% of IPv6 PI objects covered by “abuse-c:”
 - deadline for End Users/LIRs setting it is end Q3 2014
 - after deadline RIPE NCC will set to sponsoring LIR

- Introduced SSO – Single Sign On for updates
 - When signed in, no passwords required for database operations on your data if your **mntner** object(s) include “auth: SSO xyz”
 - Currently 979 “auth: SSO” in use in 210 **mntner** objects
 - More features/services to come built on SSO...
 - authenticated queries
 - certification of PI resources (RPKI)
- Introduced pending **route** authorisation – 2 step hierarchical auth process

- Continued to improve documentation
 - Manuals linked to software version
 - Draft updates to manuals released with RC
 - Working with Comms on new documentation style
 - Focusing on 2 options – will choose 1 after meeting
- Hot node in Stockholm up and running
 - load balanced with Amsterdam servers
 - accepting queries now
- Cleanup of references to returned ASNs ongoing

- Added mandatory “status:” to all **aut-num** objects to clearly identify authoritative source (‘ASSIGNED’, ‘LEGACY’, ‘OTHER’)
- Set all legacy **inetnum** objects to “status: LEGACY”
 - all top-level (4k) and all more specific objects (35k)
- [FAQ](#) with more details
- Added “sponsoring-org:” to End User resources, **inetnum**, **inet6num**, **aut-num**
- Software and converted data available for testing – [Release Candidate Environment](#) – check it out!!

- Deprecate “referral-by:” attribute
- New flag to *request* personal data
 - Reverse logic to not provide personal data by default
- Only block **person** objects when limits reached
 - Blocking query behaviour currently blocks all or nothing
- Internationalisation of data in RIPE Database
 - UTF-8 technically possible
 - still no policy on how to handle this data

- Replace “changed:” with new auto-generated “created:” and “last-updated:” attributes
 - 7.4M objects, 8.6M “changed:”, 105K unique emails
 - 27.4% of “changed:” dates do not correctly document created or last updated dates
 - Real values available for most object types with history query ‘–list-versions’
- Improved dummification algorithm removing slightly less data and standardised wherever personal data is removed

- Improved user experience (requested from member survey 2013)
- Review of all error messages (do you parse them?)
- History query to show deleted data (no personal)
- **Role** object no longer considered personal data
- Convert **irt** object to **role** object
- Change syntax of “role:” to match “organisation:”
- Allow change of “person:” and “role:” names

- Authenticated queries
 - If we know who you are and which organisation you are with, we can do more for you
 - Creating personal views of your data
 - allow you to look at your own data in ways that are meaningful to you, but views restricted to you and your organisation
 - Unlimited access to your personal data
 - no rate limiting, no limits, no blocking

- Integrate RIPE Database better with other RIPE NCC services
 - Log in once and use all services
 - We know who you are, which organisation you are with, what is registered to you, what your permissions and preferences are
 - Seamlessly carry out your business with any of these functions - look, request, create, change, delete/cancel, download, sync with your data
 - No longer think in terms of separated services with own passwords – LIR Portal, RIPE Database, RIPEstat, DNS checks, IP analysers, various wizards

- Full re-implementation of auto cleanup including clusters – **person, role, organisation, mntner**
- Deprecating object editors in LIR Portal by use of business rules to lock down fixed object parts
- Server side **as-set** expansion
 - single API call, replaces large number of queries

