

Ops Session

Thursday June 5 2025 11:00-11:30, Room: room 1309

Euro-VO Registry weather report

Speaker: Henrik Norman

This time at the interop, Henrik presented a new twist on their usual report by taking a more proactive approach to actively improving the VO ecosystem.

The EURO-VO registry is a fully searchable registry that aggregates data from the entire VO ecosystem. It also includes validators that check services twice a month to ensure compliance with VO resource standards.

For several months now, they have been digging deeper into the specific issues users encounter with their services, reaching out directly to providers to ask them to fix particular problems.

They contacted around 40 data providers managing approximately 150 different services. The response rate was very impressive—over 80%—and most providers are actively fixing issues with their services.

They haven't yet deeply investigated Cone Search, but they have reviewed TAP, SIAP (versions 1 and 2), and SSA services, reaching out to the relevant providers. Cone Search is planned for review later this summer.

During this process, they also discovered problems with some of their validators, which they have since addressed and improved.

They faced an issue with the ESA firewall, where 1,015 different services were flagged as suspicious, blocking all requests. This has been largely resolved, and they can now access nearly all of those services.

SSL certificates remain a recurring issue, whether due to expiration or trust problems during renewal, which the teams have been able to resolve by coordinating with the service providers.

Another common problem is improper redirection from HTTP to HTTPS. In many cases, services redirect all HTTP traffic to HTTPS, but the original redirect then downgrades the connection back to HTTP, causing failures. This has affected several data providers.

Several services have been offline for various reasons. For example, the Chile data center suffered a major outage and is expected to take several months to recover. In cases where services are permanently unavailable—such as when funding has been lost—the registry asks the responsible publishers to mark these resources as deleted or inactive to keep the system clean. However, a challenge arises when these services are published through their own registries, which may also be offline, leaving no way to flag the resources as inactive. This has been a problem for the Euro-VO

registry. To address this, if a service remains continuously unavailable, it should be removed from the registry of registries upon request, which then propagates to other registries.

Regarding Cone Search, the number of services continues to grow by around 100 new ones each month. Most issues are concentrated with just two data providers. Common problems include missing the required ID column or having two different columns labeled as the main ID. These providers will be contacted soon.

For SIA version 1, they contacted owners of 90 services and observed a slight improvement. A frequent issue is missing coordinate data, which limits validation. Providers are encouraged to add this data in their VOResource descriptions or migrate to SIA version 2.

SIA version 2 has seen impressive improvement over the last few months, with compliance now being very strong.

For SSA, they contacted owners of 15 services. Most issues stemmed from one main provider who has already fixed about half of them.

TAP required the most effort, with 28 providers contacted to fix specific issues. Many services had multiple, varied problems, so the team focused on a single category of issues and reached out accordingly. Compliance improvements are less obvious here because although some issues were fixed, other problems may remain.

Overall, people are responding well and actively improving their services.

This effort will continue into 2025.

Q&A:

Tamara Civera: To determine if a resource is truly dead, what time frame do you use before deciding it's no longer active?

Henrik Norman: In some cases, maintainers have reported that their services have been offline for over 10 years. They perform these checks roughly twice a month to monitor service availability.

Christophe Arviset: Last year, 1,400 new Cone Search services were registered. During the major version transition session led by Marcus, there was discussion about what to do with Cone Search, including whether to drop support for it altogether. The high number of new registrations indicates that Cone Search remains actively used.

Room: Some Cone Search services are intentionally failing because they added a DataLink service descriptor to all their results. The basic Cone Search version 1 doesn't allow a second resource, causing these failures on purpose. We hope to update the standard to fix this.