

OneSTOP Deliverable D4.5 - First prototype of alert system running online

Project: [OneSTOP](#), [Grant Agreement 101180559](#), Horizon Europe. Work package 4. Lead beneficiary: NNIT (The Binary Forest). Type: DEC. Dissemination: PU. Due: M21 (30 September 2026). Submitted: 30 September 2026.

Authors: Nicolas Noé (The Binary Forest). Task 4.3 leader and project coordinator: Quentin Groom (Meise Botanic Garden, MBG). Reviewers appointed under Task 7.3, "Quality control" (January 2025 to June 2028, M1-M42): Reviewer 1, Natalia Joelson (Science for Policy Researcher, European Commission Knowledge Centre for Biodiversity, KCBD), scientific review; Reviewer 2, Tom August (UK Centre for Ecology and Hydrology, UKCEH), technical review. Their comments and our responses are in Appendix A.

Live prototype: <https://onestop.thebinaryforest.net/> Software website: <https://www.gbif-alert.org> Source code: <https://github.com/riparias/gbif-alert>, release v2.8.0 (29 September 2026). Archived release: <https://doi.org/10.5281/zenodo.23041534>

Version history: v0.9 draft for review by the coordinator (MBG), 2026-09-26; v1.0 submitted 2026-09-30.

Licence: this report is released under the [Creative Commons Attribution 4.0 International licence \(CC BY 4.0\)](#). The EU emblem is not covered by this licence. Map backgrounds in the figures: © OpenStreetMap contributors.



Funded by
the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101180559. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

Summary for the reader

What this deliverable is. D4.5 is the first online prototype of the OneSTOP alert system: a running web service, at <https://onestop.thebinaryforest.net/>, that [watches GBIF](#) for new occurrence records of alien species and notifies subscribed users by email and on the web. It is built on [GBIF Alert](#) maintained by The Binary Forest (the commercial name of NNIT, beneficiary 19 of the OneSTOP consortium) and used by several projects. This report describes what the prototype does at M21, what it does not do yet, and what the final deliverable, D4.3 "Results of the alert system" due in February 2028 (M38), will add.

Two tangible products of the period. The running prototype with its public source code, and the Technical Requirements Document written jointly by OneSTOP and GuardIAS, which is the specification both projects build against and can be read on its own (Appendix B).

Development activity since January 2025 ([public repository riparias/gbif-alert](#), counted on 2026-09-29):

Indicator	Value
Releases tagged (release candidates excluded)	26, from v1.7.8 to v2.8.0
Pull requests merged	125
Commits on the development branch	888
Issues opened / closed (closed includes issues opened before 2025)	77 / 59

What changed between v1.7.7 (November 2024) and v2.8.0, in five points:

- *Alerts that fit how people work* (early 2026): records older than a user-set delay no longer raise notifications, a new alert does not flood its owner with past records, and observations can be filtered by basis of record and by verification status.
- *A rebuilt application* (v2.0, June 2026): new single-page interface, sortable table, scientific or common names everywhere at the flip of a switch, alerts on records approaching an area, areas drawn or uploaded by users, a documented public API with personal tokens and rate limits, and container-based deployment so that a new instance can be set up from published images.
- *Tools for instance operators* (v2.1 to 2.3): species managed through the API or a spreadsheet, reusable alert templates, harvest limited to a bounding box, species pictures, and support for [Catalogue of Life](#) identifiers now that the GBIF taxonomic backbone is frozen.
- *Scale and reach* (v2.4 to 2.5): imports about 2.75 times faster, area filters much faster on large databases, a Species tab, a phone-friendly layout, dark mode.
- *Adaptable to other countries, and the review's harvest* (v2.6 to 2.8): base maps configurable per instance, a harvest that can be scoped to a continent, a first Norwegian translation (machine-generated, not yet reviewed) added alongside English, French and Dutch, security hardening, map legend and scale bar, full screen maps, a Datasets tab, accessibility fixes, error emails for every failed step of the nightly import, and the other fixes that came out of the Task 7.3 review.

Feedback gathered. The alert system was presented and discussed with stakeholders at thirteen events between January 2025 and September 2026, within OneSTOP, jointly with [the sister project GuardIAS](#), and at conferences; Appendix C lists them. Two outcomes stand out: the interest of [LifeWatch ERIC](#) in hosting the tools, and a Norwegian team (Artsdatabanken and NTNU) setting up its own instance, which prompted a first, machine-generated Norwegian translation of the interface in v2.6.0; the team has agreed to contribute to the shared code base rather than maintain a separate fork. The peer review under Task 7.3 (section 12) produced 38 comments, all answered in writing. The correctness bugs and the most serious security findings were fixed within two weeks of being reported, in v2.6.0 (14 September) and v2.7.0 (22 September); the account-security points that remain are listed in section 9.

Quality control. Two reviewers from outside the task reviewed the live prototype between 26 August and 14 September 2026, following Task 7.3. Their comments and our responses, sent to them on 25 September, are in Appendix A.

1. Introduction

Task 4.3, "Alert system to disseminate detection events of new IAS" (June 2025 to February 2028), develops an alert system that turns the flow of occurrence records published on GBIF into timely, targeted notifications for the people who can act on them. D4.5 delivers the first online prototype at M21, halfway through the task. Annex 1 describes it as including the partial results of Task 4.3; the complete results are reported in D4.3. The Annex 1 text therefore describes the target system, and this report is explicit about which parts of that target are implemented, partial or deferred (section 3).

Software and instances. The prototype is not a one-off development. It is a deployment of GBIF Alert (www.gbif-alert.org), an open-source web application (MIT licence) first built for the LIFE RIPARIAS project in Belgium, where it has run in production since 2022, and now maintained by The Binary Forest for several projects. In 2023 GBIF Alert won the first prize of the [GBIF Ebbe Nielsen Challenge](#), GBIF's annual open-innovation award, a sign of recognition by the wider biodiversity-informatics community. The software is deliberately species-agnostic: it monitors whatever species list its operator configures, over whatever territory. Each deployment is called an instance and has its own name, species list, territory, languages and areas of interest. What this report reviews is the OneSTOP test instance. The project expects several production instances rather than a single EU-wide one (section 2.1).

Terminology. This report says "alien species" throughout, and "invasive alien species (IAS)" only where the meaning of [Regulation \(EU\) 1143/2014](#) is intended. The Technical Requirements Document (Appendix B) has been aligned to the same convention.

How to read this report. Section 2 states what the system is for and how it relates to EU infrastructure. Section 3 is the traceability table. Sections 4 to 8 describe the prototype. Sections 9 to 11 cover limitations, roadmap and sustainability. Section 12 summarises the Task 7.3 review.

2. Scope and positioning

2.1 What the alert system is

The OneSTOP alert system is a monitoring and notification layer on top of open occurrence data. It continuously harvests occurrence records of a configured species list from GBIF, lets users define alerts by species, geography, data source and verification status, and notifies them by email and through the web interface when new matching records appear. All data it holds is public, all code is open source (MIT licence), and every record it shows links back to its source occurrence and dataset on GBIF.

It is built as generic software (GBIF Alert), deployed as instances. Instances share code and differ in configuration: species list, geographic extent, languages, areas of interest. [The LIFE RIPARIAS instance in Belgium](#) has been in production since 2022, and GuardIAS will deploy the same software for aquatic species. The prototype delivered as D4.5 is a OneSTOP test instance. In production the project expects several instances, each covering a territory and a species list that are meaningful to its users (for example one per country; the breakdown is not decided), rather than a single EU-wide instance. A single instance at that scale would show every user a large majority of species and records that are irrelevant to them, and would be needlessly heavy to operate. This is how the "European alert system" of Annex 1 is delivered; see the note under section 3.1.

2.2 What it is not

- It is not a registry of validated alien species records. It shows what publishers have made available through GBIF, with the verification status they attached, and does not itself validate identifications.
- It is not an official notification channel under Regulation (EU) 1143/2014. Early detection notifications to the Commission and to other Member States [go through NOTSYS](#); the alert system may help a national authority notice a record earlier, but it does not replace or duplicate the official procedure.
- It is not a species information system. Species profiles, pathways, risk assessments and distribution baselines are [the domain of EASIN](#) and of national systems. The alert system carries only the minimum species attributes needed to filter and display records, and links out for the rest.

2.3 Intended roles by level

Local and regional. The primary use case. A site manager, a river contract, a protected-area authority or a municipality defines its own area (drawn, uploaded, or chosen from loaded administrative or Natura 2000 polygons) and receives a notification when a species of concern is recorded inside it, or within a chosen distance of it. This is the level at which early detection translates into rapid response, and the level at which no EU infrastructure currently delivers push notifications.

National. A national competent authority uses an instance covering its territory to watch the whole of it, filtered by the national IAS list or by horizon-scanning species, and as a surveillance complement to its own reporting workflow: GBIF-published records surface here days after publication rather than at the next reporting cycle. National authorities also curate the species list for their instance.

EU. The alert system does not act at EU level, and no EU-wide instance is planned. Its contribution at that level is indirect: it gives data holders (citizen science platforms, regional agencies, research institutes) a concrete reason to publish alien species records to GBIF promptly, since publication is what makes a record visible to alerts. EASIN harvests GBIF, so every record that reaches GBIF for the alert system also reaches EASIN, with no additional integration work on either side.

2.4 Relationship to EASIN and NOTSYS

The three systems occupy different positions in the same flow:

```
data holders -> GBIF -> [alert system: notify subscribers]
                    -> [EASIN: aggregate, baseline, species information]
competent authority -> NOTSYS -> Commission and Member States (official)
```

The alert system and EASIN are both consumers of GBIF. They serve different users (individual subscribers with self-defined areas versus the EU-level baseline) and different time horizons (days versus reporting periods). NOTSYS is upstream of neither; it is the official channel a national authority uses once it has decided that a record warrants notification. The alert system can be one of the inputs to that decision.

Integration direction. The project's primary integration route is GBIF-first: data enters the shared ecosystem by being published to GBIF under an open licence, with [Darwin Core](#) metadata including identification verification status. Any occurrence data source that publishes to GBIF is integrated with the alert system automatically, and reaches EASIN by the same route. This applies equally to national reporting datasets, to EASIN-held data and to Member State baselines. It is the route the project invests in first, because it benefits every downstream consumer at once and requires no bilateral development.

Beyond that, the project will explore with the EASIN team (the JRC, which runs EASIN, is a beneficiary of the sister project GuardIAS) what direct exchanges between the alert system and EASIN or NOTSYS would be useful and feasible: for instance consuming EASIN species or distribution data as reference data, or surfacing alert-system records to EASIN where they are not yet in GBIF. The feasibility of each depends on the availability of a documented, stable and openly licensed interface on the EASIN side, on the licensing of the data exchanged, and on the resources available to both parties within the project. These exchanges are therefore documented as options to be assessed during the remainder of Task 4.3 rather than as committed features of the prototype, and the outcome of that assessment will be reported in D4.3.

Independently of any direct exchange, EU reference data that improves the alert system without creating a runtime dependency can be consumed as configuration today: the Union list and horizon-

scanning lists as species lists, the [EEA reference grid](#) as an aggregation unit for exports, and MSFD marine regions as filterable areas. These are one-off imports by the instance operator.

2.5 Known limitations

- **Latency.** A record is visible to the alert system only after it is published to GBIF and after the instance's next harvest (daily on the OneSTOP test instance). Records that publishers hold back until a reporting deadline are invisible until then.
- **Coverage of verified official records.** In several Member States, verified records from official monitoring reach GBIF late, partially, or not at all. The alert system cannot show them. It labels as "unverified" the records whose publisher did not fill the identification verification field, which is the majority of citizen science records; the label will be reworded to "not marked as verified" (section 9). This is a property of the data, not of the alert system, and the project's response is to promote publication rather than to ingest from alternative channels.
- **First-record detection.** The prototype alerts on records newly present in the latest harvest that match an alert. It does not yet determine whether a species is new to an area. This is a planned event type for the final system (D4.3).
- **No official status.** A notification from the alert system carries no legal weight and creates no obligation. It is a signal to a subscriber.

2.6 Sustainability

The software is maintained as a single open-source codebase serving several projects and instances, which spreads maintenance cost and removes the risk of a project-specific tool that stops at the end of the grant. Operation of OneSTOP instances after the end of the project (M42, June 2028) is addressed in section 11.

3. Scope of the prototype and traceability

This section is the scope statement of D4.5. It lists every sub-action of Task 4.3, every feature named for D4.5 in Annex 1, and every MUST of the Technical Requirements Document, each with a status: implemented, partial, or deferred to D4.3 with a target. For a first prototype at M21, "deferred" is a legitimate status; "silently missing" is not, and nothing is left out. The Task 7.3 reviewers assessed the prototype against Annex 1 and the Technical Requirements Document (Appendix A); the tables below record the state of the prototype at submission.

3.1 Sub-actions of Task 4.3 and features named for D4.5 in Annex 1

Annex 1 describes Task 4.3 as five sub-actions, and D4.5 as a website with four features. Both are traced below: the sub-actions in the words of the Annex, the D4.5 features condensed from its description.

Annex 1 text	Status	How the prototype does it	Target
(1) aggregate occurrence data from IAS in Europe, including those data published in T4.1 and using the GRIIS Europe checklist published in T4.2 as a species filter	Implemented, with a national checklist	The nightly GBIF harvest aggregates every record of the configured species over the configured territory, whatever the publisher; data that Task 4.1, "Publication pipelines of detection data from T2.1-T2.4 to GBIF and EASIN" (June 2025 to December 2026, M6-M24), publishes to GBIF is picked up without any further work. The species filter of the test instance is the 2020 edition of the GRIIS checklist for Romania . Romania is one of the four countries of Task 4.2, whose updated Checklist of introduced taxa in Romania (August 2026, 824 taxa) will replace it on the test instance soon. GRIIS Europe, the aggregation of the national lists, is produced by Task 4.2, "Generate new and updated national GRIIS checklists and GRIIS Europe", and planned for October 2026 with D4.2 "Automated pipelines for the Global Register of Invasive Species (GRIIS)" (M22); a production instance loads the checklist of its territory, and GRIIS Europe can be loaded as soon as it is published	GRIIS Europe: M22
(2) set up an infrastructure that can broadcast those observations in real time (e.g. RSS/API)	Implemented, without RSS	Two layers broadcast. GBIF itself is the primary infrastructure: a record published there is available to every consumer at once. The prototype adds a documented REST API with personal tokens and an OGC Web Feature Service on the aggregated, filtered records. No RSS feed; the need has not come up with users, and the API serves the same purpose	RSS on request
(3) to set up a system that listens to those feeds and aggregates them	Implemented as a scheduled pull	The nightly import is the listener: it asks GBIF for everything new that matches the instance, loads it, detects what is new to each user's alerts, and keeps a stable identity for each record across harvests. It polls once a day rather than subscribing to a push feed, because GBIF offers no push channel	Incremental import: D4.3
(4) to make a user-friendly interface to display that information	Implemented	Single-page interface: map, timeline, species and dataset breakdowns, sortable table, record details with links to GBIF, custom areas, three languages plus a trial Norwegian translation, phone-friendly	-
(5) to set up a taxonomically and geographically configurable email notification mechanism, enabling users to receive immediate updates on new observations of IAS	Implemented	Per alert: species, areas (inside or approaching), datasets, verification status; frequency daily, weekly, monthly or web only. "Immediate" is bounded by GBIF publication plus the nightly harvest (section 9)	-
D4.5: aggregate IAS observations from GBIF, GRIIS Europe as taxonomic filter	See (1)		
D4.5: notify subscribed users of new IAS detections	Implemented	See (5) and section 4	-
D4.5: intuitive and user-friendly interface to access and interpret IAS data and alerts	Implemented	See (4)	-

Annex 1 text	Status	How the prototype does it	Target
D4.5: interactive maps, charts and reporting tools to visualise the occurrence and spread of IAS across Europe	Partial	Maps and charts: yes. Reporting tools: bulk access through WFS only, no export in the interface. "Across Europe": see the note below	Downloads: D4.3

On "across Europe". Annex 1 speaks of a European alert system and of visualising IAS across Europe. The project delivers this as one software and a set of instances, each covering a territory and a species list that are meaningful to its users, rather than as one EU-wide website (section 2.1). The European aggregation of occurrence data is done by GBIF, and the European species list by GRIIS Europe (Task 4.2); every instance draws on both. What the project adds is the notification layer, and that layer only works well when a user is not shown records that are irrelevant to them.

3.2 MUST requirements of the Technical Requirements Document

Condensed from a requirement-by-requirement analysis of the code (July 2026, re-checked September 2026). Section numbers refer to the Technical Requirements Document.

Req.	MUST	Status	Target
1	Detect and report new, expanding or significant observations of priority taxa	Partial: new records matching an alert are detected and reported; expansion and significance are not assessed	D4.3
2	Monitor Darwin Core occurrence records from GBIF	Implemented	-
2	Detect key events	Partial: "new record matching my alert" only	D4.3
2	Configurable filters (taxa, location, ecosystem, triggers)	Implemented, ecosystems through areas: any ecosystem or habitat type is filtered by loading its polygons as areas (for example CORINE land-cover classes); no habitat attribute is stored per record	-
2	Deliver via web, email, API	Implemented	-
3.1	Trigger on new records	Implemented	-
3.1	Trigger on updated records	Deferred: the prototype alerts on new records only; a record whose identification, date or coordinates are corrected on GBIF keeps its viewed status. Whether and how to re-alert on such corrections will be assessed together with the incremental import	D4.3 (assessment)
3.1	Filter by scientific name or taxon key	Implemented, on the curated species list	-
3.1	Filter by geographic area	Implemented, including user-drawn areas and the approaching mode	-
3.2	First record in a defined region	Deferred	D4.3
3.2	Presence in a priority conservation area	Available as an alert filter: a user creates an alert on any loaded protected area (Natura 2000 sites on other instances) and is notified of every record inside it. Not yet a labelled event: the notification does not state "record in a priority area" as such	Event label: D4.3
3.2	Custom rules via the interface	Implemented	-

Req.	MUST	Status	Target
4.1	Filter by terrestrial habitat type	Deferred	D4.3
4.1	Alerts link to protected areas and action zones	Partial: alerts are defined on protected areas loaded as areas; the notification names the alert's areas, not the specific site of each record	D4.3
4.2	Alerts link to marine regions (GuardIAS)	Implemented through areas: marine regions (for example MSFD regions) are loaded as areas and alerts filter on them. OneSTOP focuses on terrestrial species; no marine regions are loaded on the test instance	GuardIAS instance
5.1	Accept Darwin Core archives; ingest from GBIF; automated syncing	Implemented, daily	-
5.1	Nightly runs over 1 to 100 million records	Partial: full-refresh import, comfortable at current scale	D4.3
5.2	Rule evaluation; spatial joins	Partial (filter sets rather than rules); spatial joins implemented	D4.3
5.3	Email and web delivery; frequency options	Implemented	-
5.3	Payload: alert identifier, species names, location	Implemented: the notification names the alert and links to its page, and gives species, coordinates and date for a sample of the new records	-
5.3	Payload: event type	Deferred, depends on the event model	D4.3
5.3	Payload: dataset DOI and publisher	Partial: dataset name and GBIF key only	D4.3
5.3	Alert metadata: GBIF download DOI, temporal scope	Partial: every harvest links, from the "About the data" page, to its GBIF download, which carries GBIF's own DOI, so the provenance of any record is one click away; the alert itself does not carry that DOI in its metadata, and alerts have no temporal scope (only the browse filters have a date range)	D4.3
5.4	Web dashboard; REST API	Implemented	-
5.4	Downloads via API and interface, alert-scoped	Deferred	D4.3
6	Occurrence linked to dataset	Implemented	-
6	Dataset linked to publisher, hoster, licence, EML metadata	Partial: each record links to its dataset page on GBIF, where publisher, hosting institution, licence and metadata are given; the prototype stores only the dataset name and key, so it cannot display or filter by publisher or licence itself	D4.3
6	Darwin Core records with the mandatory fields (occurrenceID, scientific name, coordinates, date, datasetKey, occurrenceStatus)	Implemented: records without an identifier, coordinates or a year are skipped and counted, and only records with occurrenceStatus "present" are imported; a missing month or day is read as 1	-
6	Country code	Deferred: the ISO country code of each record is not imported; a record's country is implied by the instance's territory. Needed once an instance spans several countries or for per-country reporting	D4.3
7	Modular components (ingestion, rules, notification, interface and API)	Partial: logical modules within one application, see 5	-

Req.	MUST	Status	Target
8	Inputs and outputs follow the FAIR principles	Partial: inputs are GBIF records with persistent identifiers and open licences, every record links back to its source, the API is described in OpenAPI and the software is archived with a DOI; the outputs (alerts, downloads) have no persistent identifier or licence metadata of their own yet (section 10, tranches 1 and 2)	D4.3
8	Integrate GBIF	Implemented	-
8	Integrate national IAS checklists	Implemented by configuration: the test instance runs on the 2020 GRIIS checklist for Romania; other lists, including the updated Romanian checklist of Task 4.2 and GRIIS Europe once published, are loaded the same way	-
8	Integrate EU infrastructure (EASIN)	See section 2.4	Assessment reported in D4.3
8	Authentication and access control on data transfers	Implemented	-
8	Download objects with URL, metadata and DOI	Deferred	D4.3

4. What an alert is

The test instance watches the species of the GRIIS checklist for Romania (2020 edition, see section 3.1) over the Romanian territory. The alert features are available to anyone who creates a free account; sign-up is open on the test instance. An anonymous visitor sees the observations map and filters, which is why the prototype at first sight looks like a data viewer. Its purpose only appears once an alert exists.

Defining an alert. A user chooses one or more species, optionally one or more areas, datasets and a verification status, and a notification frequency (daily, weekly, monthly, or never for a web-only alert). Areas can be the ones the operator loaded (administrative units, protected areas, river basins; on the test instance, Romania and its counties), or areas the user draws or uploads. Three area modes exist: records inside the area, records approaching it within a chosen distance of up to 50 km, or both. Operators can publish alert templates, so that a newcomer starts from a ready-made alert for a species list and a region.

OneSTOP GBIF Alert test | Explore | What's new | My alerts | About | Scientific names | Vernacular names | EN | nnoe

Create alert

[Go back to alerts list](#)

From a template

Approaching crayfishes
All crayfishes observations from GBIF within 10km of the country border.

Use this template | Details

Naturalist records in Cluj

Use this template | Details

Configure the alert manually

Alert name *

My alert #2

Species to include *
At least one species must be selected.

No species selected

Areas to include
If you make no selection here, all observations will be included regardless of their location.

Everywhere

Datasets to include
If you make no selection here, all observations will be included regardless of their origin dataset.

All datasets

Basis of record to include
If you make no selection here, all observations will be included regardless of their basis of record.

All basis of record

Verification filter

All

Alert notifications frequency

Weekly

Create alert | Go back to alerts list

API
Latest data import: #64 - Sept. 24, 2026, 4 a.m.
This site is powered by GBIF Alert version v2.7.0

Funded by the European Union

Figure 1. the alert form, with two operator templates above the manual configuration.

What triggers a notification. The instance harvests GBIF every night. Every record that appears for the first time in that harvest, matches an alert, and was observed within the user's notification delay (below) is flagged "not viewed" for the alert's owner. The next notification email gives the number of such records, a sample of them, and a link to the alert page for the rest. Records are identified across harvests by a stable identifier, so a record that was already there is never reported twice, and a re-published record keeps its viewed status and comments. Two things the prototype does not do: it does not know whether a species was ever recorded before in the area (a "first record" is a planned event type, section 10), and it does not re-alert when a record is corrected.

The notification delay. Each user sets a delay in their profile (default: one year). A record observed longer ago than the delay is never flagged, so that a historical record digitised and published last week does not raise an alert. An alert looks forward: when it is created, the records already in the instance count as viewed, whatever their age, so a new alert starts empty rather than flooded with the past. Choosing "All" on the alert page still shows every record it matches.

Following up. On the alert page, the user sees the not-viewed records on a map, a timeline, a table and per-species and per-dataset breakdowns, opens each record with a link to its GBIF occurrence and dataset, can comment on it, and marks it as viewed one by one or all at once.




Figure 2. an alert page filtered on not-viewed records.




Figure 3. a real weekly notification of 19 September 2026 for the alert "iNaturalist records in Cluj": count, sample table with links to each record, link to the alert page.

5. Architecture and data flows

One flow. Every night the instance asks GBIF for a download of all occurrence records matching its configured query (species list, territory or bounding box), fetches the resulting Darwin Core archive, and loads it into a PostGIS database in a single transaction while the site is in maintenance mode. During the load, each record is matched to the one it replaces, comments and viewed status are carried over, and the "not viewed" flags of every user's alerts are updated. The previous harvest is then deleted. Every afternoon, a second job sends the notification emails that are due. The web application reads the database through a documented REST API (used by the interface itself and by external users, with personal access tokens and rate limits), map tile endpoints that aggregate records into hexagons at each zoom level, and an OGC Web Feature Service for GIS software.




Figure 4. data flow of one instance: nightly harvest from GBIF into the PostGIS database, served to the interface and to external users through the API, map tiles and Web Feature Service; a daily job sends the notification emails.

Stack. Python and Django with GeoDjango on the server, PostgreSQL with PostGIS, a Redis-compatible queue for long jobs, and a single-page interface in TypeScript and Vue.js with OpenLayers for maps. Everything ships as container images and a Compose file, so an instance is a matter of configuration (section 11).

Verification status. Each record's Darwin Core identificationVerificationStatus is stored and mapped, through a classification list, to a verified flag used by the filters and alerts. Publishers often leave the field empty, and such records count as not verified. Since v2.7.0 operators can declare datasets whose records are always considered verified, and iNaturalist research-grade records are verified by default.

Correctness and testing. The code base carries about 900 automated tests, from unit tests of the import to end-to-end browser tests, run on every change together with static type checking of both the server and the interface code. The Task 7.3 review included an independent code review; it confirmed the overall quality of the code base and led to corrections in how "not viewed" records and counts are computed, now covered by automated tests. An alert can be spot-checked at any time: every record it lists links to its GBIF occurrence page, where date, place, dataset and identification can be compared.

On modularity. The Technical Requirements Document asks for modular, distributed components. The prototype has the four components (ingestion, filtering, notification, interface and API) as separate units of one application sharing one database. This is a deliberate choice for a system that several small

teams must be able to deploy and operate; splitting it into services would raise the operating cost without a demonstrated need at the volumes handled so far.

6. User guide

The user guide is an online document at <https://docs.gbif-alert.org>, built from the same repository as the software and published automatically, so that it follows each release. For D4.5 its user part is complete and is the reference for how the prototype is used; it is not repeated here.

- **Exploring the data:** the map and its hexagons, the timeline, the Species, Datasets and Table tabs, the filters (species, datasets, areas, basis of record, verification status, dates, viewed or not viewed), and the scientific or common name switch.
- **Alerts:** creating an alert from a template or by hand, choosing species and areas (inside, approaching, or both), what triggers a notification and how the notification delay works, and following up on the alert page.
- **Advanced use:** custom areas drawn or uploaded, personal API tokens and the REST API, and bulk access through the Web Feature Service for GIS tools.

The operator part of the guide (loading a species list and tagging it, for example by realm; loading areas such as administrative units, protected areas or marine regions; publishing alert templates; configuring base maps and languages) is still being written and will be completed before D4.3; the installation and upgrade instructions already live in the repository. The figures below show what a user meets first.




Figure 5. home page: filters, hexagon map with legend and scale bar.




Figure 6. the Datasets tab with per-dataset counts and shares.




Figure 7. a user's alerts overview.

Usability and accessibility. The interface works on a phone (v2.5.1) and in dark mode (v2.5.3), in three languages plus a trial Norwegian translation, and uses plain wording reviewed with non-specialist

users. On accessibility the prototype has the basics rather than a certification: it is built on a component library that implements the WAI-ARIA patterns and keyboard operation for menus, dialogues, tabs and tables; every form field has a label; icon-only controls carry a text label for assistive technology; the contrast of the filter and alert panels was corrected in v2.5.3. The map and the timeline are visual by nature; the Table tab gives the same records as text. No assessment against WCAG 2.1 has been made yet, and no automated accessibility check runs in the test suite; both are planned before D4.3 (section 10), since production instances run by public authorities fall under the [Web Accessibility Directive](#).

7. Operations, robustness and data protection

7.1 Operations

The test instance is hosted by The Binary Forest on a virtual private server, deployed from the published container images with the project's official Docker set-up through the Dokploy platform, which provides the reverse proxy and TLS certificates. The database is backed up daily to off-site object storage (Backblaze B2), and a restore from those backups has been tested. No uptime monitoring is in place; it was not considered necessary for a first prototype and is listed as a limitation in section 9. This is a test set-up sized for the review and the project's own use; production instances will be hosted under the arrangements described in section 11. The application exposes a health endpoint for the platform's probes, a dashboard for the job queue, and an administration interface. The installation guide covers reverse proxy and TLS set-up, running several instances on one host, periodic tasks, upgrades and backups; it recommends a managed database with built-in backups for production.

7.2 Robustness

- The nightly import runs in maintenance mode, and the database is updated in a single transaction. If any step fails (GBIF download, preflight check, archive reading or database import), the administrators receive an email naming the failed step; a failed database step is rolled back and the site leaves maintenance mode on its own. If the import is interrupted, for example by a redeploy, maintenance mode is cleared when the web application restarts. Unusable rows are skipped and counted rather than aborting the import. Downloaded archives are cleaned up whether the import succeeds or fails. The administrators also receive an email after every successful import.
- Dependence on GBIF: the import waits for GBIF to prepare the download for a configurable time (10 hours by default). If the GBIF download API is slow or unavailable, the import fails for that night, the administrators are emailed, the previous harvest stays online, and the next night retries.
- The public API is rate limited (per instance configuration; on the test instance 60 requests per minute for anonymous clients, 600 for authenticated ones, and 5 sign-in or sign-up attempts per minute per address). Limits cannot be bypassed by forging proxy headers.
- Following the technical review, the stability and security of the API and of user accounts were improved in v2.6.0 and v2.7.0: stricter access control, protection against repeated sign-in attempts, safer error reporting, and up-to-date dependencies. What remains: verification of email addresses, the announcement of API usage limits to clients, and the separation of interface and API capacity under heavy load (section 9).

7.3 Data protection

The instance holds personal data of its registered users: a username, an email address, optional first and last names, a language preference and notification delay, the dates of sign-up and last sign-in, and what they create: alerts, comments, the areas they draw or upload, and personal API tokens (stored

hashed). Passwords are stored hashed only. The data is used only to deliver the service the user subscribed to: the notification emails are the service itself, sent because the user created an alert, not marketing. A user can stop them at any time by setting an alert's frequency to never, deleting the alert, or deleting the account, which removes their alerts and their own areas and empties their comments.

The test instance does not yet display a privacy notice (section 9, [issue #312](#)). It will name the instance operator as data controller (The Binary Forest for the test instance), the purpose (delivering the notifications the user asked for), the lawful basis (performance of the service requested by the user, Article 6(1)(b) GDPR) and retention (account data is kept until the user deletes the account). Three other limitations are listed in section 9: email addresses are not yet verified at sign-up, an address change is not yet confirmed, and stopping notifications requires signing in. Data protection on production instances will be cross-checked with Task 7.4, "GDPR, Gender Balance, and Ethical Issues Management" (January 2025 to June 2028, M1-M42), before they open.

8. Open science

- Licence: MIT. Source code, issue tracker and release history are public at github.com/riparias/gbif-alert. The repository carries a citation file (CITATION.cff) naming the funding.
- Release: v2.8.0, 29 September 2026, [archived on Zenodo with DOI 10.5281/zenodo.23041534](#) (concept DOI [10.5281/zenodo.22935511](#) always resolves to the latest release). Since v2.7.1 every stable release is published on GitHub and archived on Zenodo automatically by the release workflow.
- Standards: Darwin Core and Darwin Core archives, the GBIF API, an OGC Web Feature Service, an OpenAPI description of the REST API.
- Data: everything the instance shows is public GBIF data, and every record links back to its source.
- Data management plan: GBIF Alert is registered in the project's Data Management Plan (D1.5, June 2025, updated in D1.6 in August 2026; the final version D1.7 is due in April 2028) as a software output under the MIT licence; its Zenodo DOI will be added at the next update. The alert system generates no research dataset of its own: it reuses occurrence records published on GBIF under their own licences, links every record back to its source, and stores only the account data described in section 7.3. Its species lists are checklists published on GBIF: the test instance uses the 2020 GRIIS checklist for Romania (CC BY 4.0), and will move to the checklists of Task 4.2, which are DMP datasets (Romania, published in August 2026 under CC BY 4.0; GRIIS Europe, planned for October 2026).
- EU visual identity: the software carries an EU funding acknowledgement (emblem, funding statement and disclaimer, as required by Articles 17.2 and 17.3 of the Grant Agreement) that instance operators switch on for EU-funded instances. It is enabled on the test instance.




Figure 8. About page with emblem, statement and disclaimer, and the emblem in the site footer.

9. Known limitations

- Latency: a record is visible only after publication to GBIF and the next nightly harvest.
- Coverage: verified official records often reach GBIF late, partially, or not at all; the alert system cannot show them (section 2.5).
- "Unverified" covers every record whose publisher left the verification field empty, not only records found wrong; the label will be reworded to "not marked as verified".
- No first-record detection yet: the prototype alerts on records new to the harvest, not on species new to an area.
- A notification has no official status.
- One instance serves one territory and one species list; EU-wide coverage means several instances, by design (section 2.1).
- The nightly import is a full refresh with a maintenance window; adequate at current volumes, to be redesigned for the 100-million-record target.
- No data export from the interface; bulk access through WFS only.
- Account security: email addresses are not yet verified at sign-up, and an email change is not yet confirmed. Stopping notifications requires signing in; there is no unsubscribe link in the emails. The planned authentication revamp (section 10) addresses all three together.
- The interface and the public API share capacity on the test instance, and API usage limits are not yet announced to clients.
- No uptime monitoring on the test instance; an outage is noticed by its users or by the emails of the nightly import, not by a probe. Expected for production instances.
- No privacy notice on the test instance yet (section 7.3).
- Accessibility: no WCAG assessment yet.

- Pickers: long flat lists in which an existing selection is easy to miss; a few small interface improvements suggested by the reviewers are not yet released (Appendix A).

10. Deferred features and roadmap to D4.3

The work toward D4.3 is organised in three tranches, ordered by what unblocks what.

Tranche 1 - provenance. Publisher, licence and metadata of each dataset, the DOI of each GBIF download, and the country code of each record. A prerequisite for the FAIR requirements of the Technical Requirements Document and for meaningful downloads.

Tranche 2 - downloads. Download objects with their own URL and metadata, in the interface and the API, including downloads scoped to an alert, and exports aggregated to the EEA reference grid (1 km and 10 km) for national and local reporting.

Tranche 3 - the event model. Typed events rather than saved filters: first record of a species in a defined area (the change the scientific review asked for most), presence in a priority area, and re-alerting on corrected records, designed together with an incremental import that knows what changed between harvests. Event type becomes part of the notification.

For D4.3 itself. Annex 1 asks D4.3 to report the number of alerts distributed and insights into the user base. From the next release the prototype records every notification sent (alert, date, number of records reported), so that D4.3 can report totals per instance, per period and per user over the whole second half of the task. The prototype at M21 holds only the date of each alert's last email; earlier sends are not recoverable, which is why recording starts now.

Accessibility. A self-assessment against WCAG 2.1 level AA and an automated accessibility check in the end-to-end tests, before D4.3.

Alongside. Stronger sign-up and account security with the option to sign in through an existing account such as Google; photos of the records themselves, read from the multimedia part of the GBIF archive; easier species and area selection, with hierarchy, more tolerant search and lists that adapt to each other; higher taxonomy per species from the Catalogue of Life; the habitat and marine dimension through typed links between alerts and regions; risk scores as species attributes (a MAY in the requirements, candidate sources Harmonia+ and the project's own risk outputs, no commitment at this stage); and the assessment of direct exchanges with EASIN and NOTSYS described in section 2.4.

11. Sustainability

The software is one open-source code base serving several projects and instances (LIFE RIPARIAS since 2022, OneSTOP, GuardIAS, and [a public demo instance](#)). Maintenance cost is shared, each project's improvements benefit the others, and the tool does not stop with any single grant. The Norwegian team setting up its own instance has agreed to contribute to this code base rather than maintain a separate fork. A new instance is deployed from published container images with a configuration file; the installation guide is public and covers upgrades.

The test instance is hosted by The Binary Forest for the duration of the project. For the long-term and official hosting of GBIF Alert instances beyond the project, the first option considered was a transfer of the tools to EASIN, discussed with DG Environment and the JRC in January 2026 (Appendix C). That option was not pursued after March 2026. LifeWatch ERIC then agreed in principle, in April 2026, to host the alert tools on its infrastructure; a technical meeting is planned for October 2026, and the outcome will be reported in D4.3.

12. Quality review (Task 7.3)

Process. Following Task 7.3, two reviewers from outside Task 4.3 assessed the live prototype between 26 August and 14 September 2026, against Annex 1, the Technical Requirements Document and the eight criteria of the review plan. The review material, sent on 26 August, was the live prototype, the Technical Requirements Document and the review plan; this report was written after the review, and the written responses were sent to both reviewers on 25 September. The reviewers were Natalia Joelson (Reviewer 1, European Commission Knowledge Centre for Biodiversity, KCBd), for fitness for purpose from a downstream user's point of view, and Tom August (Reviewer 2, UK Centre for Ecology and Hydrology), for the technical side. EV-INBO acted as arbiter, none was needed.

Reviewer 1 (Natalia Joelson) returned 20 comments. One, the reviewer's first impression of a data viewer rather than an alert system, came from the review material, which did not include a demo of the alert features; section 4 now fills that gap for future readers. Three, on the difference between a record new to GBIF and a species new to an area, named what became the headline item of the roadmap (section 10). Eight concerned the interface and were answered by changes (map legend and scale bar, a clearer "Share" column, a Datasets tab, realm tags on every species), by the user guide, or by roadmap items (a revamp of the species and area pickers, exports on the EEA reference grid). Five concerned scope, data sources and integration with EU infrastructure, and led to section 2 of this report (the question on risk scores to section 10). Three concerned the Technical Requirements Document: its terminology was harmonised, the software and its instances are now distinguished (section 1), and every MUST is traced in section 3.

Reviewer 2 (Tom August) raised 18 points through the public issue tracker, including an independent code review. In his covering email he noted that the code base was "too large ... to do a thorough review in 0.5-1 day", so his points are a sample of a large code base rather than an audit. The correctness bugs and the most serious security findings were fixed within two weeks, in v2.6.0 and v2.7.0; the others, including email verification at sign-up, are on the roadmap or listed as limitations (section 9). He concluded: "Really fantastic work here, congratulations to you and the team."

Outcome. 38 comments, all answered in writing; the full record is in Appendix A. Two changes stand out: the scope and positioning statement of section 2, and the fixes at the data layer that an outside reader found within a day, which is the strongest argument for repeating the exercise at D4.3. One lesson for that round: give the reviewers an agreed scope statement before the review opens, so that a first-prototype gap is not scored as a missing feature. For D4.5 that statement is section 3, written after the review.

Acknowledgements

We thank Damiano Oldoni (Research Institute for Nature and Forest, INBO), who develops the alert system with us on the GuardIAS side: co-author of the Technical Requirements Document, presenter of the system at Living Data 2025, the 1st GuardIAS virtual conference and Neobiota 2026, and organiser of the collaboration with LifeWatch ERIC.

We thank Lissa Breugelmans (Meise Botanic Garden, MBG), who tested the prototype and prepared its demonstration for the Belgian living lab.

We thank the two reviewers appointed under Task 7.3, Natalia Joelson (European Commission Knowledge Centre for Biodiversity, KCBd) and Tom August (UK Centre for Ecology and Hydrology, UKCEH), for their careful reviews. The distinction between a record new to GBIF and a species new to

an area, raised in the scientific review, now leads the roadmap (section 10), and the independent code review in the technical review led to the fixes at the data layer described in section 12.

We also thank Lien Reyserhove and Peter Desmet (Research Institute for Nature and Forest, INBO), who reviewed the tool and provided valuable feedback.

Appendix A - Task 7.3 review record

A.1 Review plan (Task 7.3), as sent to the reviewers on 26 August 2026 (4 pages). A.2 Reviewer 1 (Natalia Joelson): comments and response table (3 pages). A.3 Reviewer 2 (Tom August): issues and response table (3 pages).

Appendix B - Technical Requirements Document

"European/National Biodiversity Alert System - Technical Requirements Document" (Noé, Groom, Oldoni), attached. Written jointly by OneSTOP and GuardIAS, it is the specification against which section 3.2 is traced.

Appendix C - Stakeholder engagement and dissemination events

Events at which the alert system was presented, discussed or evaluated, January 2025 to September 2026.

Date	Event	Setting	What happened for the alert system
21 Jan 2025	Joint OneSTOP-GuardIAS-EASIN technical workshop, JRC Ispra	In person, both projects and EASIN/JRC	Kick-off of the joint alert-system work; summary reviewed with EASIN in March 2025
25 Sep 2025	Meeting with EASIN on the alert system requirements	Online	Requirements document discussed
30 Sep - 2 Oct 2025	GuardIAS annual meeting, Thessaloniki	In person (The Binary Forest remote)	Alert system presented to the GuardIAS consortium; on 2 October, a workshop with EASIN and project partners on upscaling and integrating the early alert tool, with a demo, account creation and feedback groups
21-24 Oct 2025	Living Data 2025, Bogota	Conference	Joint GuardIAS-OneSTOP talk on the alert system (D. Oldoni); origin of the Norwegian contact
14 Nov 2025	1st GuardIAS virtual conference	Online conference	Talk on the development and integration of GBIF Alert as a European early warning system (D. Oldoni)
Nov 2025	Meeting with EASIN, Ispra	In person (INBO, EASIN)	Data flows EASIN to GBIF, relevant to the alert system's coverage
19 Dec 2025	GuardIAS stakeholder workshop	Online, decision-makers at national, regional and EU level	Early Alert System presented with a demo; stakeholder feedback collected; LifeWatch ERIC expressed interest in hosting
15 Jan 2026	GuardIAS-OneSTOP session with DG Environment and JRC	Online	Early warning system and stakeholder feedback presented; DG ENV feedback; discussion of a transfer of the tools to EASIN, not pursued after DG ENV's written reply of March 2026

Date	Event	Setting	What happened for the alert system
19-22 Jan 2026	OneSTOP annual general meeting, Almeria	In person, whole consortium	Dedicated GBIF Alert session on 19 January (The Binary Forest, WP4)
Jan 2026	Finnish living lab core group meeting	Living lab	Survey on what an early warning system should do, results received February 2026
29 Jan 2026	Meeting with the Norwegian Biodiversity Information Centre (Artsdatabanken) and NTNU	Online	Norwegian needs for an alert system; scope and plans of GBIF Alert; the Norwegian team has since been working on a Norwegian version
21 Apr 2026	GuardIAS-OneSTOP-LifeWatch ERIC meeting	Online	Agreement in principle to deploy the alert tools on LifeWatch ERIC infrastructure; technical meeting planned for October 2026
7-11 Sep 2026	Neobiota 2026, Brussels	Conference	Speed talk on the alert system (Q. Groom, with The Binary Forest and INBO); at the LIFE RIPARIAS end conference on 8 September, a joint talk on CamAlien (S. Delva) and the OneSTOP early alert tool (D. Oldoni); in-person session with the Norwegian team on 10 September

Not listed: internal coordination meetings (biweekly OneSTOP-GuardIAS development calls, WP4 pillar meetings, the 2 June 2026 alignment meeting with Task 5.6, "Automating communication of information to support effective management" (June 2026 to April 2028, M18-M40)) and the OneSTOP European checklist webinar of 11 May 2026, which concerns Task 4.2.

OneSTOP D4.5 review plan

Date: 2026-08-18

Status: draft, for agreement with MBG before end of August 2026

Deliverable: D4.5 "First prototype of alert system running online"

WP4, lead beneficiary 19 - NNIT, type DEC, dissemination PU

Due: M21 = 2026-09-30. Draft freeze: 2026-09-02.

Grant Agreement 101180559 (OneSTOP), Annex 1.

Related: [2026-07-29 alert system requirements gap analysis](#)

Scope of this plan

This plan covers the **internal consortium peer review** required by Task 7.3, which is completed *before* submission. The separate assessment by the Project Officer and EU expert reviewers takes place *after* submission, at the periodic review - D4.5 falls in RP2. The engineering deadline is therefore 2 September, not 30 September.

Basis

Annex 1, Task 7.3 (Quality control, lead MBG):

Quality control will be ensured through peer review of deliverables by at least two reviewers who are not part of the task. Task leaders responsible for the delivery of the deliverables will need to prepare a draft at least four weeks before the delivery deadline and report it to the Coordinator. The management team will appoint a reviewer and ensure that the draft is reviewed, and feedback is provided to the task leader before the final submission of the deliverables.

Four weeks before 2026-09-30 gives a draft freeze of **2026-09-02**.

Point to settle. Task 4.3 is led by MBG with EV-INBO. Under a strict reading of "not part of the task", neither can act as an independent reviewer for D4.5, so reviewers come from outside T4.3. Reviewers appointed: Natalia Joelson (scientific) and Tom August (technical) - confirm both sit outside T4.3 before the review opens.

Maturity

D4.5 is a *first prototype* at M21, mid-way through Task 4.3 (M6-38). Annex 1 states it "will include the *partial* results of Task 4.3", the complete results being reported in D4.3 at M38. The Annex 1 text therefore describes the target system, not the M21 state.

To avoid reviewers scoring against the wrong bar, NNIT submits a **scope baseline** - which features are in the prototype, which are deferred to D4.3, and why - for endorsement by MBG as task leader *before* the review starts, and reviewers assess against it. Depth may be provisional, but the end-to-end path (ingest -> filter -> publish -> notify) must work.

1. Roles

Role	Who	What they do
Author / deliverable lead	NNIT (Nicolas Noe)	Builds the system, writes the accompanying document and scope baseline, resolves comments.
Task 4.3 lead	MBG	Contributor and requirements owner - not an independent reviewer.
Reviewer 1 - scientific / fitness for purpose	Natalia Joelson	Downstream user of the alerts. Checks fitness for purpose against Annex 1 and O4.3.
Reviewer 2 - technical / IT	Tom August	Code, deployment, data flows, security, documentation.
Optional	GuardIAS counterpart; 2-3 end users	Cross-project check evidences the synergy committed in the DoA; a 30-min user test doubles as dissemination evidence.
Arbiter	EV-INBO (WP4 lead)	Decides if author and reviewers disagree.
Sign-off & submission	MBG (Coordinator)	Confirms QC completed; uploads to the Funding & Tenders Portal.

Expected reviewer effort: 0.5-1 person-day each - please book it when accepting.

2. What is reviewed

D4.5 is a DEC deliverable - a live service - so review is anchored to a frozen snapshot:

- ♦ the live URL
- ♦ an accompanying document (8-12 pp: scope baseline, architecture, data sources, user guide, limitations and deferred features)

- ♦ git tag **v0.9-D4.5** , public repository and dated screenshots

Final release tagged **v1.0-D4.5** and archived to Zenodo with a DOI.

3. Review criteria

- 1. Traceability to Annex 1** - a table listing each feature named in the D4.5 / Task 4.3 text: GBIF aggregation; GRIIS Europe as taxonomic filter; real-time broadcast (RSS/API); feed listener and aggregation; interface with interactive maps, charts and reporting tools; email notification configurable by taxon and geography. Each marked implemented / partial / deferred to D4.3, with a one-line justification and target for anything not complete. Deferred is a legitimate outcome for a first prototype; silently missing is not.
- 2. Functional correctness** - scripted acceptance tests; spot-check alerts against the source GBIF records; false positive/negative check on the GRIIS filter; alert de-duplication.
- 3. EU visual identity** - EU emblem, funding statement and disclaimer present and correct (GA Article 17.2).
- 4. Open science** - open licence, public repository, tagged release with DOI, consistency with the project DMP, open standards (Darwin Core, GBIF API).
- 5. Data protection** - email subscriptions are personal data: privacy notice, lawful basis, one-click unsubscribe, retention, hosting arrangement. Cross-checked with Task 7.4 (MBG / CU).
- 6. Robustness & operations** - behaviour when the GBIF API is slow or down, rate limiting, backups, uptime monitoring, error handling.
- 7. Usability & accessibility** - mobile rendering, accessibility basics, plain language for non-specialist users.
- 8. Documentation & sustainability** - user guide, deployment/admin notes, API documentation, and a short statement on operation after M42.

Criteria 2-8 scored OK / minor / major; majors resolved before sign-off, minors logged as known limitations. A feature deferred under the agreed scope baseline is not a major.

4. Timeline

Date	Step
by 2026-08-31	Reviewers confirmed (done: Natalia Joelson, Tom August); this plan agreed; MBG endorses the M21 scope baseline.
2026-09-02	Feature freeze. NNIT sends the live URL, draft document, tagged release, criteria checklist and test script to reviewers (Task 7.3 four-week rule).

Date	Step
2026-09-02 - 09-15	Reviewers work and return written comments in a shared table.
2026-09-16 - 09-23	NNIT resolves comments and completes a response table: comment -> action taken -> where.
2026-09-24 - 09-26	Reviewers confirm resolution; Coordinator signs off.
2026-09-29 - 09-30	Tag v1.0-D4.5 , archive to Zenodo, Coordinator submits to the Portal.

Record: reviewer comments and the response table are annexed to the deliverable - evidence that Task 7.3 was applied, and the template to reuse for D4.3 at M38.

European/National Biodiversity Alert System - Technical Requirements Document

Nicolas Noé, Quentin Groom, Damiano Oldoni

Version of 25 September 2026, terminology harmonised following the D4.5 internal review.

Executive Summary

This document defines the functional and technical requirements for a European or national-level biodiversity alert system. It builds upon an existing alert prototype by incorporating new features that enhance usability, scalability, and alignment with regional conservation goals and policy obligations.

The system will provide real-time or near-real-time detection of biodiversity occurrence events such as new species records, range shifts, and alerts related to conservation priorities or legal mandates (e.g. Regulation (EU) 1143/2014 on invasive alien species). Alerts will be delivered via multiple interfaces including web interface, email notifications, and APIs (Application Programming Interface). The system will be designed using FAIR principles, ensuring Findability, Accessibility, Interoperability, and Reuse.

A modular architecture will allow for ongoing enhancements, such as integration with eDNA data, satellite-based indicators, or national risk lists. The system will serve multiple stakeholder roles: data publishers, field ecologists, policy makers, and the public, via configurable filters, multilingual interfaces, and interoperable standards.

Introduction

An Alert System will be designed to automatically detect and report significant biodiversity observations, particularly the appearance of alien species, in a timely, structured, and scalable way. It will process occurrence records from platforms like GBIF and apply configurable filters and rule-based triggers to generate alerts relevant to conservation practitioners, alien species managers, policy makers, and biodiversity researchers. This system will support early detection, rapid response, and coordinated data-driven decision-making across regional or national landscapes.

The system will build on [GBIF Alert](#), the core functionality of the [LIFE RIPARIAS alert system](#), which provides early alerts for alien aquatic and riparian species in Belgium. Lessons learned from that system's implementation and use will inform the expanded design, enabling broader taxonomic coverage, additional alert triggers, and wider geographic applicability at the national and European scale.

This system will be jointly developed by the **OneSTOP** project, focusing on terrestrial alien species, and the **GuardIAS** project, focusing on aquatic alien species. While the two projects share a common infrastructure, their specific objectives require that some alert types, filters and user needs be addressed separately. OneSTOP will emphasize terrestrial habitats and land-based surveillance networks, while GuardIAS will focus on aquatic ecosystems and riparian surveillance. Each component will incorporate tailored filtering criteria, alert triggers, and stakeholder reporting aligned with the goals of the respective project.

Key Takeaway Messages

- Expands alert into a modular, scalable system suitable for European or national implementation.
 - Enables automated, role-based alert generation for alien or high-priority taxa.
 - Supports filtering by country, marine regions, habitat, ecosystem, protected area, and species list.
 - Embeds FAIR principles, persistent identifiers, and metadata traceability.
 - Built on open APIs, modular architecture, and extensible components.
 - Designed for integration with national and EU biodiversity infrastructures.
-

1. Purpose and Scope

This document defines the technical and structural requirements for a new biodiversity alert system suitable for implementation across Europe or at a national scale. It expands on the foundational GBIF Alert system and integrates additional features to accommodate broader geographical, institutional, and functional scope.

The system **MUST** detect and report new, expanding, or significant observations of priority taxa, particularly Invasive Alien Species (IAS), in support of conservation, early warning, and biosecurity strategies.

The document uses RFC 2119-style requirement language to express levels of obligation in technical documents, particularly in specifications and requirements. It helps ensure that everyone involved clearly understands what is mandatory versus optional.

‘MUST’ 🚩 Indicates an absolute requirement. If this requirement is not met, the system will fail to meet its core specifications.

‘SHOULD’ 🚩 Indicates a strong recommendation. There may be valid reasons to deviate from this recommendation, but any such deviation should be carefully considered and justified.

‘MAY’ 🚩 Indicates a permissible option. The feature is truly optional, and implementers can decide whether to include it based on context or need.

Terminology. This document says “alien species” for any species introduced outside its natural range, and “invasive alien species (IAS)” only in the sense of Regulation (EU) 1143/2014, i.e. alien species whose introduction or spread threatens biodiversity and ecosystem services. “Non-native” and “invasive species” are not used.

2. System Overview

The alert system:

- **MUST** monitor biodiversity occurrence records following the DwC standard (e.g. GBIF)
 - **MUST** detect key events (e.g. first record, expansion, presence in sensitive areas)
 - **MUST** allow configurable filters based on taxa, location, ecosystem, and user-defined triggers
 - **MUST** deliver actionable alerts via multiple channels (e.g. web interface, email, API)
 - **SHOULD** serve stakeholders across countries, institutions, and roles
-

3. Functional Requirements

3.1 Event Detection & Filtering

- The system **MUST** allow alerts to be triggered by new or updated occurrence records.
- Filters **MUST** include at minimum: scientific name, taxonKey, geographic areas.
- The system **SHOULD** support additional filters, such as:
 - Habitat or ecosystem type
 - Administrative region

- Protected area membership
- Identification verification status
- Basis of record (human observation, machine observation, preserved specimens, ...)

3.2 Alert Types

- The system **MUST** support at least the following event types:
 - First record in a defined region
 - Presence in a priority conservation area
- The system **MAY** support at least the following event type:
 - Temporal increase in frequency (expansion)
- Users **MUST** be able to define their own custom alert rules via the user interface.
- User **MAY** be able to define their own custom alert rules via an API

4. Project-Specific Requirements

4.1 OneSTOP: Terrestrial Alien Species

- The OneSTOP component of the alert system will focus on detecting and managing alerts for terrestrial alien species. To meet its specific goals, the following requirements apply:
 - The system **MUST** support filtering of occurrence data by terrestrial habitat types (e.g. forests, grasslands, agricultural zones) where such data are available.
 - Alerts **MUST** include links to relevant terrestrial protected areas and biodiversity action zones.
 - Data ingestion pipelines **SHOULD** integrate with land use/land cover datasets (e.g. [CORINE Land Cover](#)).
 - The rule engine **MAY** incorporate landscape-level risk models and buffer zones for dryland spread.

- Notifications **MAY** be integrated with national IAS platforms and land management systems.

4.2 GuardIAS: Aquatic Alien Species

- The GuardIAS component of the alert system will address aquatic and riparian alien species, drawing from lessons of the LIFE RIPARIAS system. Its specific requirements include:
 - Alerts **MUST** link to marine regions.
 - Alerts **SHOULD** link to hydrological units and water management regions
-

5. Technical Requirements

5.1 Data Ingestion & Standards

All occurrences of interest must be delivered as a single, coherent batch from the same system and in a consistent format. The system **MUST** support nightly downloads of all occurrences of interest—ranging from 1 to 100 million records—with good performance.

- The system **MUST** accept records delivered as Darwin Core Archive (DwC-A).
- Data **MUST** be ingested via API endpoints or harvested from GBIF downloads.
- It **MUST** support regular, automated syncing (e.g. at least weekly).

5.2 Rule Engine

- The rule engine **MUST** evaluate occurrence data against predefined or custom criteria.
- It **MUST** support spatial joins (e.g. overlays with Natura2000 protected areas, marine regions).
- It **MAY** integrate Harmonia+ risk scores for filtering.
- It **MAY** integrate GuardIAS (WP3) risk scores(?) for filtering.

5.3 Notification Engine

- The system **MUST** be able to deliver alerts via email and web interface.

- Users **MUST** be allowed to choose the frequency of the email alerts. The following timespans **MUST** be implemented: *never*, *weekly*, *monthly*.
- Alert payloads **MUST** include:
 - Unique alert ID
 - Species names
 - Location
 - Event type
 - Source dataset info (DOI, publisher, etc.)
- Alert metadata **SHOULD** be exportable as JSON metadata.
- Alert metadata **SHOULD** be shareable.
- Alert metadata **MAY** have a DOI
- Alert metadata **MUST** contain:
 - DOI of the GBIF download
 - Title of the alert
 - Taxonomic scope
 - Geographic scope
 - Temporal scope

5.4 System Interfaces

- A web-based dashboard **MUST** be available for human users.
- The system **MUST** provide a REST API for programmatic access.
- The system **MUST** support data downloads through both the API and the user interface (UI).
- The system **MUST** allow data downloads related to a specific alert.
- UI **SHOULD** be multilingual where relevant.
- The interface will support multilingual content, but the translations effectively available will depend on the availability of volunteer translators.

6. Data Requirements

6.1 Hierarchical Structure

- Each occurrence **MUST** belong to a dataset.
- Each dataset **MUST** belong to a publisher, e.g. INBO.
- Each dataset **MUST** be hosted by a data hoster. Publisher and data hoster MAY be not the same organisation.

6.2 Minimum Required Fields

Occurrence records **MUST** follow the Darwin Core (DwC) standard.

The following fields are mandatory and **MUST** be present:

- `occurrenceID` (stable and unique identifier),
- `scientificName`,
- `decimalLatitude`, `decimalLongitude`
- `eventDate`, `year`, `month`, `day`,
- `datasetKey` (dataset identifier),
- `occurrenceStatus`,
- `countryCode` (if relevant)

The following fields are not mandatory, but improve the quality of the occurrence records and so they **SHOULD** be present:

- `coordinateUncertaintyInMeters`,
- `municipality` (if relevant),
- `identificationVerificationStatus`,

6.3 Metadata

- The source metadata **MUST** conform to EML (Ecological Metadata Language).
- Each dataset **MUST** include publisher identifiers and license.

- Alert bundles **SHOULD** include reference to all source datasets.
-

7. System Architecture

- The system **MUST** be modular and support distributed services.
 - Key components **MUST** include:
 - Data Ingestion Module
 - Rule Evaluation Engine
 - Notification & Distribution Layer
 - User Interface and API Gateway
 - Each module **SHOULD** expose interfaces for monitoring and debugging.
-

8. Interoperability & FAIR Compliance

- Data inputs and outputs **MUST** follow FAIR principles (Findable, Accessible, Interoperable, Reusable) from both technical and legal perspectives.
- The system **MUST** support integration with:
 - GBIF
 - National alien species checklists (e.g. GRIIS)
 - EU infrastructure (e.g., EASIN)
- Each alert and dataset **SHOULD** be associated with a persistent identifier (e.g., DOI). Public web pages must allow users to trace back each entity to its source, including:
 - Occurrence
 - Dataset
 - Publisher
- All data transfers **MUST** support standard authentication and access controls.

- Each download must have a unique URL, descriptive metadata, a data file link, and a DOI.
-

9. Future Enhancements

- The system **MAY** support extensions for:
 - eDNA signal integration
 - Predictive modelling of spread, risk or impact
 - Citizen science feedback loops (e.g. iNaturalist, ObsMapp)
- Alert analytics **MAY** be integrated into dashboards to track trends.

10. Additional Notes

Format Flexibility (Limited)

While GBIF Alert can accommodate minor variations in machine-readable formats (e.g. alternate DwC dialects or DCAT metadata instead of EML), the underlying content structure—based on the occurrence/dataset/publisher model with unique identifiers—must remain unchanged without a complete conceptual redesign of the tool.

Publication Pipeline Requirements

To enable end-to-end data flows, the ecosystem must also support robust publication mechanisms, including:

- Data publishing tools
- Training materials
- Technical support This ensures that partner organizations can easily and automatically publish their biodiversity data to the source system (as a prerequisite to integration with GBIF Alert).

Baseline vs. Future Enhancements

The requirements listed above represent the minimum technical baseline for GBIF Alert's current functionality.

However, we aim to enhance capabilities during the OneSTOP and GuardIAS projects. Supporting additional data types (e.g. eDNA) and extended use cases aligned with program goals would further improve outcomes and value for all stakeholders.

11. References

- [RIPARIAS Alert List Report 2022](#)
- [EU Regulation 1143/2014](#)
- <https://www.riparias.be/338/>
- <https://www.gbif.org/developer/occurrence>
- <https://ec.europa.eu/environment/nature/invasivealien/>

OneSTOP D4.5 - Response to the scientific review

Reviewer: Natalia Joelson. Review received 2026-09-14. Response sent 2026-09-25; this version, updated on 2026-09-30 for Appendix A of the deliverable report, gives the section numbers of that report. Prepared by Nicolas Noé, The Binary Forest (NNIT), Task 4.3.

How to read the table

- **FIX**: small change to the prototype, already released or to be released soon. Any not yet released is listed as a known limitation in the deliverable report (section 9).
- **CLARIFY**: already exists; made visible or documented.
- **DOC**: addressed in the D4.5 deliverable report; the "Where" column gives the section.
- **ROADMAP**: deferred to D4.3, "Results of the alert system", due in February 2028 (M38); stated as such in the deliverable report.
- **SCOPE**: outside the deliverable; answered in section 2 of the report, first sent to the reviewer as a standalone scope note.

Response table

#	Reviewer comment (condensed)	Assessment	Action	Where
0.1	Functions as a GBIF viewer with spatial/temporal filters rather than a targeted alert system.	Fair for an anonymous visitor: alerts, seen/unseen tracking, delay and notifications only exist behind login. The end-to-end path (ingest -> filter -> notify) is implemented and was the M21 objective.	CLARIFY: the report has a section "What an alert is", with screenshots of the alert form, an alert page and a notification email, and states that sign-up is open to anyone.	report, section 4
2.1a	Unclear whether the tool alerts on records recently added to GBIF or on species genuinely new to a region.	Correct. The prototype alerts on records newly present in the latest GBIF import that match the alert's filters. It does not know whether the species was already recorded in the area.	DOC: stated in the report. FIX: one sentence on the alert form.	Alert form help text; report, sections 2.5 and 4
2.1b	Three situations (new introduction, new record of an established species, recently digitised historical record) should be distinguished.	Case 3 is already handled: each user sets a notification delay and records whose event date is older than that delay are never flagged. Cases 1 and 2 are not distinguished today.	CLARIFY case 3 (delay setting, report). ROADMAP cases 1 and 2.	report, section 4; D4.3
2.1c	"Observations = 1" filter is a poor proxy for a new introduction; implement a "First occurrence" flag as the primary trigger.	Agreed. No such filter exists; the reviewer presumably used the count column of the Species tab. "First record in a defined region" was already a MUST in the requirements document and a known gap (section 3.2 of that document). It needs an event model plus knowledge of what changed between imports, which is why it is scheduled with the incremental-import work.	ROADMAP: "first record in area" event type, D4.3 tranche 3. Named as such in the roadmap.	report, sections 3.2 and 10; D4.3
2.2a	Detects only intra-national introductions; should account for arrivals from a selected neighbouring area.	Exists: the area filter has "inside", "approaching" and "both" modes with a 0-50 km buffer, both for browsing and for alerts. Not discoverable from the public map.	CLARIFY: screenshot and paragraph in the user guide; mode selector gets a one-line hint on the alert form.	report, sections 4 and 6; alert form

#	Reviewer comment (condensed)	Assessment	Action	Where
2.2b	Relying on GBIF alone misses recent verified occurrences, e.g. EU baseline reporting; A. melas shows fewer records than EASIN and all unverified.	Both observations are correct and are consequences of the GBIF-first design, not defects. Records that are not published to GBIF cannot be shown. "Unverified" is derived from the Darwin Core identificationVerificationStatus field; when a publisher leaves it empty the record is classed unverified, so the label overstates what is known.	FIX: relabel to "not marked as verified" with a tooltip naming the source field. Since v2.7.0, operators can declare datasets whose records count as verified, and iNaturalist research-grade records are verified by default. DOC: latency and coverage listed as known limitations. SCOPE: data-source policy explained in section 2 of the report.	Filter sidebar; report, sections 2 and 9
2.2c	Integrate Horizon Scanning lists, latest official reporting occurrences and neighbouring-country IAS lists.	Species lists are configuration, not code: an instance operator loads any list (GRIIS, horizon-scanning, national) through the admin or the species API. Occurrence data from official reporting is welcome as soon as it is published to GBIF, which is the integration route the project promotes.	CLARIFY (species lists, user guide). SCOPE (occurrence sources).	report, sections 2.4 and 6
2.3a	Clarify how the planned integration with EU infrastructure (EASIN) is envisaged, and acknowledge limitations of centralising timely verified IAS data in GBIF.	Agreed that this was missing. The new section distinguishes the GBIF-first route (in place) from direct exchanges with EASIN or NOTSYS, which are to be explored with the EASIN team during the rest of Task 4.3 and reported in D4.3, subject to interface availability, licensing and resources on both sides. The limitations of GBIF as a channel for timely verified records are acknowledged explicitly.	DOC: new "Scope and positioning" section.	report, section 2
2.3b	Define scope and use cases at national, local and EU level, and the main limitations for integration with existing infrastructures.	Same.	DOC: same section.	report, section 2
2.3c	Where would landscape-level risk model data come from?	This is a MAY item in the requirements document (4.1). No source is committed; the concrete candidates are the OneSTOP WP risk outputs and Harmonia+ scores, as species-level attributes. Not part of D4.5.	DOC: mark as MAY / deferred with candidate sources named.	report, section 10
2.4a	Map needs a legend and a scale bar.	Correct. The single-observation map has a scale bar, the main map has neither. The hexagon colour scale is logarithmic with min and max recomputed per filter.	DONE (v2.7.0, 2026-09-22): scale bar plus a colour legend for the hexagons.	Observations map
2.4b	Why hexagons? Official grids such as the EEA 10x10 km grid would be preferable for national or local reporting.	The hexagon grid is a zoom-dependent visualisation aggregate (a different cell size per zoom level, in the map projection), not a reporting unit. The EEA grid is a fixed 10 km grid in ETRS89-LAEA and only makes sense at one scale. Reporting on the EEA grid is a data-export need, not a map need.	DOC: explain the hexagons in the user guide. ROADMAP: an export of filtered or alert-scoped observations aggregated to the EEA 1 km / 10 km grid, as part of the downloads work already planned for D4.3 (tranche 2).	report, sections 6 and 10; D4.3
2.4c	"Share" bar is not defined; unclear what the reference total is.	Correct. It is each species' count divided by the total number of observations matching the current filters.	DONE (v2.7.0, 2026-09-22): column renamed "Share of filtered observations".	Species tab

OneSTOP D4.5 - Response to the technical review

Reviewer: Tom August. Review conducted 2026-09-09 to 2026-09-11 through 18 issues on the public issue tracker (riparias/gbif-alert #416, #417, #429, #430, #432 to #445), including an independent code review (#417). Answers were given on the issues; this table records them as of 2026-09-30. Prepared by Nicolas Noé, The Binary Forest (NNIT), Task 4.3.

How to read the table

- DONE: fixed, with the release that carries it. v2.6.0 was released on 2026-09-14, v2.7.0 on 2026-09-22 and v2.8.0 on 2026-09-29.
- FIX: small change, to be released soon. Listed as a known limitation in the deliverable report until it is.
- ROADMAP: planned beyond D4.5, mostly for D4.3, "Results of the alert system", due in February 2028 (M38).
- CLARIFY: already exists; explained.
- EXPLORE: agreed in principle on GitHub, feasibility or design to be checked before it is planned.

Response table

#	Reviewer comment (condensed)	Assessment	Action	Where
#416	With two alerts, are their filters merged, so that an observation matching neither alert is flagged "not viewed"?	Correct, a real bug. Introduced by a recent optimisation and masked in most views because results are further filtered by the selected alert.	DONE v2.6.0: an observation is flagged only when it matches at least one alert in full (PR #420), with a regression test.	Import / unseen logic
#417-1	Timeline and species counts inflated when an area filter is active and areas overlap.	Correct.	DONE v2.6.0. The same double count on the map colour scale was fixed in v2.5.3.	API v2
#417-2	Filters merged across a user's alerts.	Same as #416.	DONE v2.6.0.	-
#417-3	Server error in "approaching" / "both" mode when the area no longer exists.	Correct.	DONE v2.6.0: returns no observations, as "inside" mode did.	Filtering
#417-4	Comment migration costs one statement per re-imported observation.	Correct.	DONE v2.6.0.	Import
#417-5	A row without basis of record aborts the import with a misleading message.	Correct.	DONE v2.6.0: the row is skipped and counted.	Import
#417-6	Downloaded archive left behind when the import fails.	Correct.	DONE v2.6.0.	Import
#417-in	Unvalidated input on the tile and legacy endpoints reaches the database and returns a server error instead of a "bad request" (zoom out of range, non-numeric ids, free-form order, limit=0, ...).	Correct. Robustness rather than security: the values are parameterised, the outcome is an error page.	DONE v2.8.0 (#474): malformed parameters on the tile and legacy endpoints return a "bad request" (400).	Tile and legacy API

#	Reviewer comment (condensed)	Assessment	Action	Where
#417-7	Image upload endpoint of the admin page editor reachable without authentication.	Correct, the most serious finding of the review.	DONE v2.6.0: staff account required.	Admin editor
#417-8	API tokens keep working after an account is deactivated.	Correct.	DONE v2.6.0.	API auth
#417-9	Publishing an alert as a template can expose another user's private area.	Correct.	DONE v2.6.0: alerts can only use areas their owner can see; templates with a private area are refused.	Alerts, templates
#417-10	Latent cross-site scripting in a link-rendering template filter.	Correct, and the filter was unused.	DONE v2.6.0: filter removed; the page configuration block is now escaped as Django's own helper does.	Templates
#417-11	No TypeScript checking anywhere, including CI.	Correct.	DONE: type check added and run in CI on every push.	Frontend, CI
#417-12	No request sequencing: a slow earlier response can overwrite a newer one.	Correct.	DONE v2.6.0, with a retry message instead of stale or empty results when a request fails.	Frontend
#417-13	CSRF cookie lookup matches on a substring of the cookie name.	Correct.	DONE: exact name comparison.	Frontend
#417-14	Interface library shipped as one 1.2 MB block on every page.	Correct.	DONE v2.6.0: per-page loading, about a third less JavaScript outside the home page.	Frontend build
#417-sm	Eleven smaller remarks (dead branch, misleading comment, N+1 in the admin, pickled queryset in the job queue, extra counts in the notification email, ...).	Reviewed.	Applied where actionable, in v2.6.0 (for example the mark-as-viewed job payload, PR #452).	Various
#429	Map popup shows the scientific name and the occurrence id. Suggest common name first, then scientific name, then date.	Partly covered: a switch in the navigation bar chooses scientific or common names everywhere, to keep tables and popups short. Adding the date is agreed. The GBIF id stays useful to tell apart several records of the same species and day.	CLARIFY (name switch). FIX: date in the popup; GBIF id shown only when needed.	Observations map
#430	The image shown on hover is a generic species image, not the observation's. Are observation images shown when they exist?	Correct, and the interface does not say which kind of image it is. Users asked for observation images in the same week.	FIX: label the species picture as such. ROADMAP: observation images, read from the multimedia file of the GBIF archive at import time.	Species name hover; import
#432	Full screen map.	Agreed.	DONE v2.7.0: every map has a full screen button.	Maps
#433	Fuzzy matching in the species and dataset search.	Nice to have. Today the search matches any part of the scientific or common name as one types, which limits the need.	ROADMAP, with the picker revamp (see Reviewer 1, 2.4g).	Pickers
#434	No way to download the map or the tables.	Correct. Bulk access exists only as an OGC Web Feature Service, aimed at GIS users. Downloads are a MUST of the requirements document.	ROADMAP D4.3, downloads tranche (same answer as Reviewer 1, 2.4b). CLARIFY: WFS documented in the user guide.	D4.3

#	Reviewer comment (condensed)	Assessment	Action	Where
#435	Filter lists should adapt to each other: only the datasets that contain the selected species, and the reverse.	Agreed, and generalised on GitHub: facet-style filtering, where every filter shows the number of matching records and updates at each change. Performance to be tested first.	EXPLORE, then ROADMAP with the picker revamp.	Pickers
#436	The GBIF dataset key in the dataset list is not informative; link the name instead.	Agreed. The key was shown because dataset names were sometimes missing; v2.5.1 fixed empty names.	FIX: a quick look at the dialogue for an interface improvement (answered on GitHub 2026-09-22); likely drop the key and keep a small link icon so that clicking the name still selects the row.	Dataset picker
#437	Sign in with Google.	Agreed, already considered.	ROADMAP: authentication revamp (with #438, #445).	Auth
#438	No email verification at sign-up: anyone can register somebody else's address, who then receives notifications; typos go undetected.	Correct.	ROADMAP: authentication revamp, answered under #445. Partial mitigation in v2.7.0: sign-up limited to 5 attempts per minute per IP. Known limitation until then.	Auth; report, sections 7.3 and 9
#439	Token creation dialogue unreadable in dark mode.	Correct, a spot missed when dark mode support was added in v2.5.3.	DONE v2.8.0 (PR #484): the new-token box and the sample command are readable in dark mode.	API tokens page
#440	In the alert form, is "All species" different from selecting every species?	Correct, the label was misleading: an alert needs at least one species, unlike datasets and areas where an empty selection means "no filter".	DONE v2.7.0: reads "No species selected".	Alert form
#441	900 API calls in a minute: the site stayed up but showed no data and empty filter lists.	Acknowledged. Interface and public API share the same application workers on a small test server, and throttled requests still occupy a worker.	EXPLORE (answered on GitHub 2026-09-22): on the boundary between code and deployment, but since the project ships a production Docker setup, it should give a good enough base or at least warn operators in the installation guide. Listed as a limitation in the report.	Deployment; report, sections 7.2 and 9
#442	Rate limits are documented nowhere and no Retry-After header is returned.	Correct on both counts: limits are configurable per instance but not shown in the API documentation, and no Retry-After header is returned.	FIX (confirmed on GitHub 2026-09-22): state the limits on the API documentation page and return Retry-After with the 429 response.	API docs
#443	Sign-in not rate limited; the general limit can be bypassed with a forged X-Forwarded-For header.	Correct.	DONE v2.7.0 (PR #466): 5 attempts per minute per IP on sign-in and sign-up, and only the last proxy hop is trusted.	API auth
#444	Error emails to administrators could contain a user's password in clear text.	Correct, though it required a crash inside an authentication request.	DONE v2.7.0 (PR #467): the HTML debug report is no longer attached. Operators are advised to delete old error emails.	Settings
#445	Email address can be changed without password, confirmation or notice to the old address: a hijacked session leads to account takeover.	Correct.	ROADMAP: authentication revamp, probably on django-allauth. Known limitation until then.	Auth; report, sections 7.3 and 9

#	Reviewer comment (condensed)	Assessment	Action	Where
2.4d	Datasets used to generate the map are not visible or clickable.	Partly a visibility issue: every observation detail links its occurrence and its source dataset to gbif.org, and the dataset filter lists all datasets. What is missing is an overview of which datasets contribute to the current selection.	CLARIFY (user guide). DONE (v2.7.0, 2026-09-22): a "Datasets" tab next to the Species tab lists every dataset present in the current results with its count and share.	report, section 6; Datasets tab
2.4e	Order species tags by taxonomic level; add marine, freshwater and terrestrial realm filters; no marine regions filter.	Tags are free-form labels set by the instance operator, so realm tags need no code. Marine regions are polygons and can be loaded as filterable areas today, as Natura 2000 sites and river basins already are on other instances. A true taxonomic hierarchy (kingdom, class) requires storing higher taxonomy per species, which the prototype does not.	DONE (2026-09-23): every species of the test instance is now tagged marine, freshwater and/or terrestrial, from GBIF species-profile data reviewed by hand; the tags appear as filter chips in the species picker. CLARIFY: marine regions loaded as areas if the reviewer or MBG names the preferred source (e.g. MSFD marine regions). ROADMAP: higher taxonomy from Catalogue of Life for hierarchical filtering.	OneSTOP instance data; D4.3
2.4f	The "Fish" filter does not work in the general species view.	The tag chips live inside the species picker and narrow the list of species to choose from; they are not an observation filter by themselves. Choosing a tag and then "select all" applies it. Confusing, agreed.	FIX: selecting a tag chip selects the matching species directly (or the picker states "narrow, then select all"). A broader revamp of the picker is on the roadmap, see 2.4g.	Species picker
2.4g	With Ameiurus melas selected, the count does not change between "Romania" and "Tulcea".	Checked on the OneSTOP test instance: 31 records in Romania, none in Tulcea. The filter itself is correct. Selected areas are combined as a union ("in any of these areas"), so if Romania was still selected when Tulcea was added, the result stays at the 31 Romanian records. We believe that is what happened, and the interface makes it easy: the area picker is a long list in which an existing selection is easy to miss, and nothing signals that a new choice is added to the previous one rather than replacing it.	CLARIFY (user guide: several areas mean "any of"). ROADMAP: the area and species pickers share the same design and both will be revamped, with hierarchy (administrative nesting for areas, taxonomy for species, see 2.4e) and a clearer display of what is currently selected.	report, sections 6 and 9; D4.3
Doc 1	Harmonise "alien", "invasive", "non-native" across the document.	Agreed. The comment concerns the Technical Requirements Document.	DOC: "alien species" throughout, "invasive alien species (IAS)" only where the Regulation's meaning is intended. Same convention in the deliverable report.	Requirements document; report, section 1
Doc 2	Name is ambiguous: "Biodiversity Alert System" vs "Alien Species Alert System".	The requirements document uses the generic name because the software (GBIF Alert) is deliberately species-agnostic and also serves non-IAS uses. Each deployed instance carries its own name. What was reviewed is a OneSTOP test instance; production instances (several are expected, not a single EU-wide one) should be named for their purpose and territory.	DOC: the introduction of the report distinguishes the software (GBIF Alert) from its instances.	report, section 1
Doc 3	Not all MUSTs in the system overview are covered.	Correct, and expected for a first prototype at M21: the requirements document describes the target system, not the M21 state. The report carries a traceability table marking each MUST as implemented, partial or deferred with a target.	DOC: every MUST of the requirements document appears in the traceability table with a status; the "first record" and "updated record" gaps are listed explicitly.	report, section 3.2