EclecticIQ Platform release notes 2.1.2

Product information, release 2.1.2

Last generated: March 30, 2018

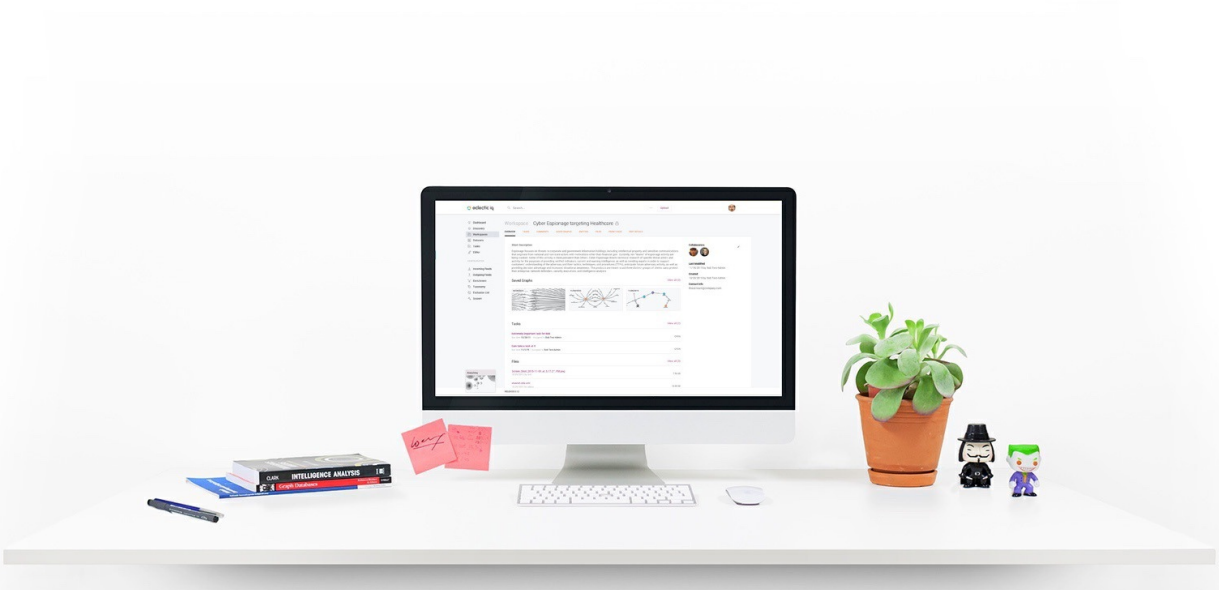

Table of contents

EclecticIQ Platform release notes 2.1.2

Release 2.1.2 — Spotlight: maintenance and bug fixing with new features and enhancements for improving user access management and interoperability with third-party components.

These release notes apply to the following product:

EclecticIQ Platform	
Release version	2.1.2
Release date	2018-03-30

Highlights

As part of our ongoing drive to improve the usability of the platform, EclecticIQ Platform release 2.1.2 contains a mix of maintenance and bug fixes, as well as two new features. The new features provide an extra level of security relating to account policy management and authentication. They enable the platform to be used and configured in line with the company security policies.

What's new

Features

Account Policy

There are two main aspects to the *account policy* feature: firstly, the platform administrator can now configure password policies based on the organizational policies. Secondly, a lockout capability has been introduced. This lets the platform administrator define the lockout settings. More detailed information is contained in the bullets below.

- *Account policy*: platform administrators can configure password policies on the **Account policy** view, based on the organization's security policy. (*16371*, *16747*, *17215*)

  They can configure password settings such as minimum character length, whether users need to include any special characters, numbers, or uppercase characters, as well as an expiry period to prompt users to renew their passwords.

  They can also configure account lockout settings to prevent user access after consecutively failing sign-in for a predefined number of times.

- If users forget their password, platform administrators can change it to a new one, after confirming their own password. (*17329*)

- To prevent account tampering, and to mitigate brute-force attacks, user accounts are automatically locked after a predefined number of consecutive unsuccessful sign-in attempts.
  Account locking is logged, and locked account users cannot access the platform. (*16374*)

  When an account gets locked, the corresponding user receives an email notification. (*17124*)
  Platform administrators receive email notifications as well, so that they can review the **Audit** trail for sign-in log information, and then they can decide whether or not they should unlock the locked accounts. (*16386*)

  Users whose accounts are locked need to contact their platform administrator to request unlocking and to be granted access to the platform. (*16375*)

  Platform administrators can view user account status to determine whether they need to carry out any actions to unlock accounts or to help with account activation. (*17133*, *17165*)

  When a platform administrator unlocks a locked account, the corresponding user receives an email notification. (*17125*)
  Platform administrators receive email notifications as well to inform them about the action.

Authentication

Given our focus on security, in this release we have enabled extra controls around authentication. A new user will now be required to set their own password. They will receive an automated email from the platform which prompts them to create their own password, instead of the password being generated by the platform itself.

- When a platform administrator saves a newly created user account for the first time, the user receives an automatic email notification with an account activation link. New accounts remain in the *pending* status until users activate them by clicking the account activation link in the email message. (*16382*)

  Platform administrators can resend an account activation email to users whose account is in the *pending* status, if users request it. (*17443*)
- When users set or change their password, their input is checked to ensure the new password is strong enough, based on the account policy settings. (*16372*)

  When users change their password, they need to enter the current password before they are allowed to set a new one. (*16378*)

  When users sign out of the platform, their valid web session token is revoked to prevent unauthorized access to the platform through token reuse. (*16750*)

What’s changed

Enhancements

User access management

- When performing a platform installation, the only predefined user account is an administrative one, because the platform requires at least an administrator account.
  During the installation, users are prompted to assign a user name and a password to the administrator account being created. (*16542*)
- When user sign-in fails, the sign-in screen does not display any specific reason for the failure to avoid providing malicious actors with details that could help them refine their subsequent sign-in attempts. (*17372*)

- When a feed, enricher, upload, or rule task run fails, administrators can view traceback content for troubleshooting purposes. Non-admin users are not allowed to view traceback content any longer. (*17193*, *17210*)

Entities

- An `eiq-platform` command line script allows identifying and fixing out-of-sync entities in Elasticsearch vs. the PostgreSQL database: `eiq-platform search sync-data`. (*16663*)
- A workaround is available to manually update the PostgreSQL database and the Elasticsearch indices to address an issue — now solved — where it would not be possible to set half-life values for existing entities . (*17305*)
- Improved checks to prevent possible data inconsistencies between the PostgreSQL and the Elasticsearch databases when reindexing large databases (millions of entities). (*16691*)

System

- During a fresh installation, the platform automatically generates a unique `SECRET_KEY` value. Previously, users were instructed to generate such a key. This step is no longer necessary. Platforms shipped as virtual machine images automatically generate a unique `SECRET_KEY` when they start and boot for the first time. (*16643*)
- Improved checks for compliance with third-party services, components, and dependencies. (*15617*, *15999*, *17571*, *17578*)
- Script-driven platform installation on Ubuntu correctly uses UFW instead of firewalld as the default firewall for the OS. (*17196*)

Documentation

- The procedure to install the platform from a tarball was reworked and updated. (*11287*, *11288*, *16323*)
  Now it includes a maintenance, and a revised upgrade section.
- The guide to install the platform on RHEL OS was reworked and updated. (*16322*)
  Now it includes a maintenance, and a revised upgrade section.
- The user guide was updated. In particular, it includes:
  - New enrichers;
  - A section on configuring account policies for platform users.
- From this release, only authenticated users who have successfully signed in to the platform can access and view platform documentation such as install guides, end-user manuals, and API documentation. (*16540*, *16541*)

Important bug fixes

The following section gives an overview of the *most important bug fixes* to provide context and scope.

Ingestion

- Improved handling of ingested entities with thousands of relationships, resulting in very large ingestion packages (> 20 MB). (*16076*)
- When ingesting data from a source lacking details about the provider's identity, the corresponding field would be automatically populated with the designated identity information for the platform. (*16217*)
- Occasionally, ingestion would partially fail due to relative namespace URI conflicts with XML canonicalization as specified in xml-c14n. (*16986*)

- Mount point incoming feeds where no regex data pattern is specified to define the files to include in the incoming feed would return an error upon execution. (*16262*)
- TAXII poll incoming feeds would pull data from the start of feed instead of retrieving only feed updates, causing unnecessary duplication. (*17417*)

Entities and observables

- After selecting the latest version of an entity, users would be redirected to a previous version, unless they perform a hard refresh of the active web browser view. (*15685*)
- It would be possible to edit and update the content of a placeholder entity referencing an external entity. (*16182*)
- Relationships stored in the platform as JSON would include unnecessary elements that may add data noise. (*16877*)
- For entity views with entity lists spanning across multiple pages, removing all items from the last page would not redirect users to the previous populated page. (*16538*)
- Deleting multiple taxonomy entries shared across a large number of entities (tens of thousands) may leave Elasticsearch in an inconsistent state compared to PostgreSQL. (*16219*)
- Leading and trailing spaces would not be automatically stripped from newly created tags and taxonomy entries upon saving them. (*16660*)
- Executing an ignore observable rule would not correctly update the connection counter on the observable view. (*16454*)
- Filtering observables based on their connections would return an Elasticsearch error. (*16520*)

Search

- Search views would sometimes display inconsistent behavior because view state handling was defined at tab level instead of globally at search level. (*15262*)

UI

- Addressed several issues to improve UI consistency and usability. (*15183*, *15262*, *15249*, *15313*, *15404*, *15447*, *15552*, *15558*, *15630*, *15636*, *15643*, *15686*, *15781*, *15795*, *15804*, *15835*, *15950*, *16314*, *16436*, *16452*, *16674*, *16878*, *16946*, *17082*, *17224*, *17230*, *17231*, *17235*, *17242*, *17361*, *17434*, *17485*, *17499*, *17500*, *17711*)

System

- The install script for CentOS and Ubuntu would assume firewalld to be already installed by default on the target system. (*16816*)
- The install script for CentOS and Ubuntu would display a message prompting users to check supervisor tasks even if they are up and running normally. (*16949*)
- Supervisord would start before the PostgreSQL, Redis, Elasticsearch, and Neo4j services could become available. (*17001*)

Misc

- Manually copy-pasting a URL pointing to a platform page would redirect users outside of the platform. (*13373*)

Known issues

- When sharing content through an outgoing feed, memory usage increases after publishing a package (*10207*, *10557*)

- After editing a feed and saving the changes, the **Last updated** time is not updated accordingly. (*16663*)
- When exchanging data including sightings with observables between two EclecticIQ Platform instances, sightings loose their observables during the data transfer. (*16853*)
- When exchanging data between two EclecticIQ Platform instances through feeds, the source platform does not export complete entity version histories. (*13642*)
- When exchanging data between two EclecticIQ Platform instances through feeds with the TAXII poll transport type, and when the receiving platform instance is configured to check incoming feed content for a valid signature, the receiving platform instance accepts incoming feed data from the source platform instance even if there is no signature. (*13642*)
- When enriching items in very large packages (> 20 MB) the Fox-IT InTELL Portal enricher completes execution, but it does not produce results. (*16224*)
- Enrichment rules do not currently have an option enabling users to run them from the rule detail pane. (*16133*)
- When creating a new rule, users can access groups they are not part of through the **Criteria selection > Source** rule configuration option. (*16142*)
- It is not possible to select rules on the **Rules** view by clicking the corresponding checkbox. (*16155*)
- After creating a new version of an entity, users are not automatically redirected to the **Versions** tab on the updated entity detail pane to check the changes. (*15420*)
- A delete option is available also for entities that cannot be deleted from a static dataset. (*15591*)
- Bulk creation of indicators from many observables (> 100) at the same time hangs and fails. (*14879*)
- When creating a dynamic dataset, users are allowed to define and save invalid regex data patterns. (*16329*)
- Occasionally, saving content to reports would behave unexpectedly, and content originally saved to the **Analysis** section of the report would not be available any more when opening the report at a later time. (*17034*)
- The graph button allowing users to access the graph from the left-hand navigation sidebar is not available to users lacking the *read graphs* permission. (*16007*)
- The graph may not refresh the view correctly after adding, and then removing an item. (*15645*)
- Occasionally, loading all observables related to an entity on the graph through the right-click context menu **Load observables > All** option would not produce the expected result. (*15911*)
- The search result view may hang or when users repeatedly cycle through the pagination options. (*15351*)
- Sorting entities by name on the **Production > Draft** view fails. (*15823*)
- MS IE and MS Edge compatibility issues slow down user experience through the web-based GUI. (*14966*)
- Mozilla Firefox compatibility issues negatively affect user experience through the web-based GUI. (*16893*)
- The install script does not log errors when it hangs during execution, and users need to exit it by pressing **CTRL + C**. (*16916*)

Contact

For any questions about the content of this document or to request assistance, you can contact EclecticIQ at the following email address: support@eclecticiq.com

The Support Team