

Date of incident:

24th September 2021 - 14th October 2021 (to identify and resolve incorrect current calculations)

- 9th November (to repair all historic data for all customers)

Summary of incident:

Customer impact:

Elements of Snyk's reporting functionality were showing incorrect or inconsistent data. Specifically, any API call starting snyk.io/api/v1/reporting in the reporting API (<https://snyk.docs.apiary.io/#reference/reporting-api>), and the Summary tab of Reporting UI. Issue counts were up to 12 hours stale compared to 3 hours during normal operation, and dips may have been visible on the graphs in the reporting UI that display the issues over time and exposure windows. Finally, the totals on the new issues and fixed issues counts on the Summary tab may have been inflated during this time.

Technical summary of incident:

The extraction service Snyk used in reporting was, on some days, not coping with the volume of load. Instead of loading all of the relevant data within the usual refresh window of 3 hours, on occasions it was taking up to 12 hours to load. The result of this was that at those times when the data was being analysed, it was incomplete. Part of the reporting analysis involves comparing an issue with its state on the previous day. If the issue is no longer present, it is assumed fixed. This means the issue is considered to be fixed and the total historical fixed count is incremented.

Another part of the analysis looks to see if an issue that was not previously listed as existing/unfixed is now present. If this is the case, then it is considered to be a new, open issue, and the new issues and overall identified issues counts are incremented on the Summary tab.

When the data was not fully loaded, the total count of current issues came back lower than it should have been. The next time the data did fully load, the issues count returned back to the correct amount, but the issues fixed and new issues counts were incorrectly incremented, and the issues count over time showed a 'data dip'.

It is a conscious decision to do the calculations in this way because there are instances where a previously identified issue may have been fixed but then reoccurs, such as via code being rolled back or changed. In these cases, it should be counted as a newly found issue, and when fixed, added to the total fixed count.

Summary of incident root cause(s):

Significant growth in the number of projects and issues in Snyk and resulting increase in load caused the temporary failure of a component that handled data extraction. This resulted in the reporting dataset being incomplete a number of times throughout the duration of the incident, in turn causing Snyk's reporting system to return inaccurate data.

Details of incident progress:

Timeline:

- 24th September 2021 - Data extraction component fails; cause identified one-off action to add an additional column to a large data table, causing significantly higher than normal load on this day. This results in data dips over this and the following day.
- 5th October 2021 - Data extraction performance becomes marginal; investigations suggest data volume growth rates have increased dramatically in the previous two months, and that a critical point may have been reached where the extraction tool will no longer be able to keep up with the volume within a matter of weeks. Diverted personnel to investigate and resolve this as a matter of urgency, starting parallel tracks of possible resolutions.
- 8th October 2021 - Unable to reduce data volume, other options still being pursued. Meanwhile, reporting's data freshness within the target of 3 hours has been delayed on more than one occasion. Incident created.
- 12th October 2021 - The new third-party solution is making good progress loading historic data, and is expected to catch up with the live dataset the following day. Other options continue to be pursued in parallel.
- 13th October 2021 - The new third-party data extraction tool has copied the full dataset, and a comparison is made of its data alongside trialling the integration with the reporting system.
- 14th October 2021 - The new third-party tool continues to perform well, and the data proves to be consistent. Development and then production reporting systems are configured to use it, resolving the reporting timeliness part of the incident. The team's focus shifts to repairing the issues in the historic data caused by the incident.

Investigations Undertaken:

The following investigations were worked on in parallel streams:

- Reducing the data volume required by reporting without impacting output.
- Working with the existing vendor to enable the data extraction component to work correctly at the new levels of scale.

- Replacing the existing third-party component with different third-party components, with increased proven scale capabilities.
- Writing a bespoke data extraction component.

Details of resolution/solution implemented:

The original third-party data extraction was replaced with a different 3rd party component for the areas that were experiencing high data load. This led to a significant increase in the component's performance and ability to scale.

Actions undertaken to prevent a recurrence of this or similar incidents:

The failing data extraction component has been replaced with one that has much greater proven scalability.

Additional proactive monitoring of the extraction and loading of reporting data has been added, using the new third-party extraction tool. This will alert on any early indications of delay in data arriving, or of any significant change in throughput.

Work is in progress on replacing Snyk's entire existing reporting infrastructure with one that is more performant and flexible. As part of the new design, there is no longer a requirement for such a third-party data extraction component. This work was ongoing before the incident, and is anticipated to be in place before the end of March 2022.