

Sonian View Archive Incident Analysis

Incident: SONIA-81625e-5

Incident Start Date: January 11, 2024

Incident Resolution Date: January 29, 2024

Impacted Region: NEST1-VIEW-US

Time Zone Indicated: EST

Incident Summary

On Jan. 11th, 2024 an issue occurred in the NEST1-VIEW-US stack of Sonian View Archive, resulting in the service continuing to receive emails, but not being able to index them. The general cause of this issue was related to an improperly implemented database type which caused an issue where we could no longer add new emails into a table that is used to pack emails together, before being indexed.

Customer Impact

- All emails received on and after 1/11/2024 were not being indexed. Customers might have noticed that affected emails were not present when they conducted a search.
- Customer Dashboard would have indicated: "Potential Disruption, No data has been indexed since Jan 11, 2024, Please check your collection configuration".

Causation Analysis

Many years ago, at the product's inception, a choice was made to use "int" as a data type for the Primary Key column in several database tables. On 1/11/2024 this incident arose when one of these tables reached the max ID value allowed by the "int" data type which capped the maximum number of entries to 2,147,483,647. By not allowing more entries to be added to the table, emails ingested by the service could not be indexed.

There was no loss of data resulting from this incident.

Remediation

By updating the data type of the Primary Key Column from "int" to "bigint", this would allow for an effectively infinite number of entries into the data tables. Once the database tables were re-indexed with the new data type, new and previously held emails would be given a proper ID and would subsequently be entered into the table for indexing. Due to the number of emails held in the backlog, it took several days to index all emails affected by this incident.

Incident Timeline

Thursday, 1/11/2024

- Issue manifests and new mail is not being indexed.

Monday, 1/15/2024 (Company Holiday)

- Several customers reported being unable to locate newer emails and these reports were received by our support teams."
- Support took a preliminary look at the situation and escalated the issue to DevOps teams.
- An operations team member took an initial look to see what could be happening and reported back that the engineering team would need to continue the investigation.

Tuesday, 1/16/2024

- 10:00am: Engineering team begins investigation.
- 10:45am: Engineering confirms the cause of the issue and shifts to investigating remediation options.
- 11:00am: An incident response was created by Product Manager and all communication organized into a related Slack Channel. A notification was sent out alerting users of the on-going issue.
- 12:00pm: Two remediation options are identified, and engineering begins work on POCs for both.

Wednesday, 1/17/2024

- 12:00pm: POCs results are checked and verified, leading to a plan of action.
- 12:10pm: A fix is pushed to a staging environment for final PDV validation.
- 2:20pm: PDV testing is validated, and deployment of a fix begins in production environment.
- 3:00pm: The first step in the fix is to re-index several Database tables.
- Everything seemed on track for the re-indexing to be completed by morning.

Thursday, 1/18/2024

- 8:30am: Upon review of the progress, engineering discovers that one of the tables did not index properly. Further inspection revealed an additional table which did not index properly.
- 9:00am: Engineering decides the best course of action is to halt any further re-indexing and fix this new issue before restarting the table re-indexing again.
- 9:30am: A new fix goes into QA for testing and PDV.

- 11:00am: New fix passes QA and PDV and is deployed into production.
- 11:30am: Tables begin re-indexed successfully.
- Engineering continues to monitor progress closely and detects no further issues or regressions.

Friday 1/19/2024

- 5:00pm: Re-indexing of the tables is completed successfully.
- 5:15pm: Ops introduces extra nodes to begin processing (indexing) the backlog of emails.
- 5:30pm: Emails which had been unprocessed since 1/11/2024 begin to process.

Wednesday 1/24/2024

- By 4:00pm, all but one partner's emails were fully processed and indexed.

Monday 1/29/2024

- 8:00am: Confirmed that all backlogged emails fully processed, and this issue would be marked as fully resolved.

Corrective Actions

We will be addressing the following areas to improve our response of this incident.

- We have conducted an audit of all relevant tables across the remaining View Archive stacks and will be replacing any data type "int" with the "bigint" data type. This work is in progress with no immediate concerns of a near-term resurgence of this exact issue.
- We will be reviewing and updating our monitoring and alerting to help capture events of this nature more quickly in the future.