Google Cloud Issue Summary

Multiple Products - 2020-08-19

All dates/times relative to US/Pacific

Starting on August 19, 2020, from 20:55 to 03:30, multiple G Suite and Google Cloud Platform products experienced errors, unavailability, and delivery delays. Most of these issues involved creating, uploading, copying, or delivering content. The total incident duration was 6 hours and 35 minutes, though the impact period differed between products, and impact was mitigated earlier for most users and services. We understand that this issue has impacted our valued customers and users, and we apologize to those who were affected.

DETAILED DESCRIPTION OF IMPACT

Starting on August 19, 2020, from 20:55 to 03:30, Google Cloud services exhibited the following issues:

- **Gmail:** The Gmail service was unavailable for some users, and email delivery was delayed. About 0.73% of Gmail users (both consumer and G Suite) active within the preceding seven days experienced 3 or more availability errors during the outage period. G Suite customers accounted for 27% of affected Gmail users. Additionally, some users experienced errors when adding attachments to messages. Impact on Gmail was mitigated by 03:30, and all messages delayed by this incident have been delivered.
- **Drive:** Some Google Drive users experienced errors and elevated latency. Approximately 1.5% of Drive users (both consumer and G Suite) active within the preceding 24 hours experienced 3 or more errors during the outage period.
- **Docs and Editors:** Some Google Docs users experienced issues with image creation actions (for example, uploading an image, copying a document with an image, or using a template with images).
- **New Google Sites:** Some users were unable to create new Sites, add new pages to Sites, or upload images to Sites. Additionally, there was almost a 100% error rate in creating Sites from a template during the incident period. Impact on Sites was mitigated by 03:00.
- **Chat:** 2% of Google Chat users who tried sending messages experienced errors, and 16% of Chat users who attempted to forward a message to Gmail experienced errors.
- **Meet:** Livestreams were fully down for the duration of the incident, and recordings were delayed due to impact on YouTube. Meet impact lasted from 21:00 to 01:15, and from 01:40 to 02:10.
- **Keep:** Some Google Keep users were served 500 Internal Server Error responses or experienced delays with operations involving media.
- **Voice:** The delivery of some outbound SMS messages with attachments failed. The delivery of some inbound voicemails, call recordings, and SMS was delayed. Impact on Voice was mitigated by 03:20. All voicemails and recordings have been delivered, with a maximum delay of 5.5 hours.
- **Jamboard:** Some users experienced errors when attempting to upload images or copy documents containing images.
- **Admin Console:** Some users experienced errors when uploading CSV files in the G Suite Admin Console. The error rate for these operations ranged between 15 and 40% during the outage period.
- **App Engine:** App Engine Standard apps making calls to the Blobstore API saw elevated error rates. Peak error rates were below 5% in most regions, but peaked as high as 47% in us-west1 and 13% in us-central1. App Engine Standard apps making calls to the Images API saw error rates up to 66%. Inbound HTTP requests served by static files or Blobstore objects saw elevated errors, peaking at 1%. Deployment of apps that include static files failed with a message "The following errors occurred while copying files to App Engine: File https://storage.googleapis.com/.... failed with: Failed to save static file.". Impact on App Engine was mitigated by 03:25.

- **Cloud Logging:** Log messages written to Google Cloud Logging, including logs generated by Google, such as App Engine request logs, activity logs, and audit logs, were delayed up to 4 hours and 43 minutes. The backlog of logs was completely processed by 16:00. During the period of outage, API calls to write and read logs returned successfully, but reads returned incomplete results.
- **Cloud Storage:** API calls to Google Cloud Storage buckets located in the "US" multiregion saw error rates up to 1%. Errors had entirely subsited by 00:31.

ROOT CAUSE

Many Google services use a common, internal, distributed system for immutable, unstructured data, also known as binary large objects, or blobs. This blob storage system contains a frontend which interfaces with Google-internal client services, a mid-layer which handles metadata operations, and backend storage for the blobs themselves. When clients make requests to the frontend, metadata operations are forwarded to the metadata service, which communicates with the storage service.

An increase in traffic from another Google service started overloading the metadata service, causing tasks to become unhealthy and requests to increase in latency. This latency prompted excessive retries of these operations, leading to resource exhaustion. Automated systems attempted to launch new metadata tasks; however, many of these tasks were immediately overwhelmed by the amount of traffic they received, and tasks that did start were allocated insufficient resources due to exhaustion. The issue was exacerbated by the strategies used to cancel and retry failed requests, which caused a multiplicative effect in traffic.

Google Cloud Storage experienced less impact than other products. While Google Cloud Storage is built on the same blob storage system, the metadata layer for GCS is mostly isolated from the affected Google-internal metadata service. The migration for GCS metadata isolation is ongoing for the "US" multiregion, while all other migrations have been completed. As a result, impact to GCS customers was lessened, and this impact was limited to the "US" multiregion.

REMEDIATION AND PREVENTION

Google Engineers were automatically alerted to the issue at 20:58 and began investigating the issue. Engineers for most affected Google Cloud services began their own parallel investigations and mitigations within an hour of the incident start time.

At 23:30, Google Engineers for the blob storage service made a rate-limiting configuration change which mitigated most of the internal errors, though impact continued due to resource exhaustion. Between 00:00 and 04:00, incident responders mitigated the incident by allocating emergency capacity to the metadata service and disabling health checks which were restarting the metadata server. The error rates, latency, and health indicators gradually improved, and the last resource exhaustion errors occurred at 04:03.

To guard against the issue recurring and to reduce the impact of similar events, we are taking the following actions:

- Increase allocation of computational resources to the blob metadata service until the root cause is fully repaired.
- Investigate and improve the health checks performed during startup of metadata service tasks, to avoid stopping tasks prematurely before resources are provisioned for them.
- Evaluate and improve the backoff and retry strategy used when metadata operations fail.
- Repair an issue in which isolated errors can lead to a flood of cancellation requests across replicated resources.

- Improve alerting in the autoscaling service used by the blob storage system, to receive earlier detection of issues when launching tasks and allocating resources.
- Implement more comprehensive rate limiting controls for requests to the blob storage service.
- Add instrumentation to permit more effective debugging of blob operations.
- Improve the speed, efficiency, and automation of transferring production resources between critical tasks.
- Improve internal playbook documentation related to rate-limiting controls in the blob storage service.

Google is committed to quickly and continually improving our technology and operations to prevent service disruptions. We appreciate your patience and apologize again for any impact to your organization. We thank you for your business.

Sincerely,

The Google Cloud Team