ANTHROP\C

Enterprise Frontier Safeguards: Technical overview

Published
November 23, 2026

Contents

1. Executive summary 3

2. How EFS works 4
- 2.1 End-to-end flow 4
- 2.2 Data categories 6
- 2.3 What EFS covers 7

3. Storage and retention 8
- 3.1 Writing to your bucket 8
- 3.2 How records are encrypted 8
- 3.3 What Anthropic holds 9
- 3.4 Retention and deletion 9
- 3.5 Residency 10

4. Access to your data 11
- 4.1 Access to your bucket 11
- 4.2 How a read is authorized 12
- 4.3 Legal requirements 12

5. Automated safety monitoring 13
- 5.1 Tasks 13
- 5.2 Scope screening 14
- 5.3 Workers 14
- 5.4 Results 15
- 5.5 Fail-closed behavior 16

6. Safety Case Review 17

7. Shared responsibility 19
- 7.1 Who does what 19
- 7.2 When requests may be refused 19

8. Customer-Managed Encryption Keys 21
- 8.1 What your key covers 21
- 8.2 Key hierarchy 21
- 8.3 Revocation 22
- 8.4 Rotation and setup 22

Appendices 24
- A. Controls reference 24
- B. Infrastructure security 25

1. Executive summary

Enterprise Frontier Safeguards (EFS) provides state-of-the-art detection of sophisticated misuse, while remaining compatible with zero data retention (ZDR): neither Anthropic nor cloud providers retain your prompts and responses.

Real-time classifiers screen individual requests, but sophisticated attacks generally occur across multiple queries, and any single one might look harmless in isolation. Our September 2026 threat intelligence report describes misuse of this kind. To catch it, EFS adds a second layer, automated safety monitoring, which looks for patterns of misuse across many requests without Anthropic storing any prompts and responses. It analyzes data stored in your cloud account, and your security team reviews any findings.

Our EFS commitments

- **Zero data retention.** Neither Anthropic nor your cloud provider retains your prompts and responses. You bring your own storage and store your own data, under your own access policies and audit logging. Anthropic encrypts records kept for monitoring before writing them to your bucket, where they are encrypted at rest. You can also require your own key to decrypt them. Each record's key expires within 30 days, after which Anthropic's systems no longer decrypt the record.
- **No human review by Anthropic.** EFS is designed so that neither Anthropic nor cloud provider personnel can access your prompts and responses. When monitoring finds a potential violation of Anthropic's Usage Policy, Anthropic may open a Safety Case, which is a record of the finding that lists the related Transcripts by identifier. It is delivered to your bucket, and your security team investigates and responds within 7 days.
- **We never train on your data.** Prompts and responses in your storage are used for automated safety monitoring and your review of Safety Cases. Monitoring runs screened tasks on isolated compute and returns closed-schema results, such as policy categories, harm scores, and record identifiers.

2. How EFS works

EFS combines two required components and one optional component, which you configure once your organization is approved for EFS.

Component	Capability	Customer requirement
Customer Storage (required)	Data is stored in your cloud account, subject to your access policies and audit logs.	Provide a bucket with a supported cloud provider, apply Anthropic's access template, configure access logging, and keep the bucket available.
Customer Safety Case Review (required)	Safety Cases are delivered to your bucket for your security team to review. Anthropic's controls are designed to prevent Anthropic personnel from accessing your prompts and responses.	Review and respond to Safety Cases within 7 days, using Customer Storage.
Customer-Managed Encryption Keys (CMEK, optional)	Sealed Safety Data is encrypted under a key you hold, so reading it depends on that key and each use can appear in your logs. If you revoke the key, within approximately one hour Anthropic will no longer be able to decrypt that data.	Create a key in a supported key service, set a key policy that lets Anthropic use it, and keep the key available.

EFS is currently available in public beta on the Claude Platform, Claude Platform on AWS, Amazon Bedrock, and Claude on Google Cloud. Features and functionality may change before EFS becomes generally available. Support for EFS in Microsoft Foundry is coming soon.

2.1 End-to-end flow

EFS stores each prompt and response encrypted in your bucket after the response is returned, outside the request path. Automated safety monitoring then analyzes the stored records separately, through recurring tasks and tasks written for a specific investigation. The systems that write to your bucket, encrypt, and analyze Safety Data make up the Safety Environment.

EFS end-to-end flow

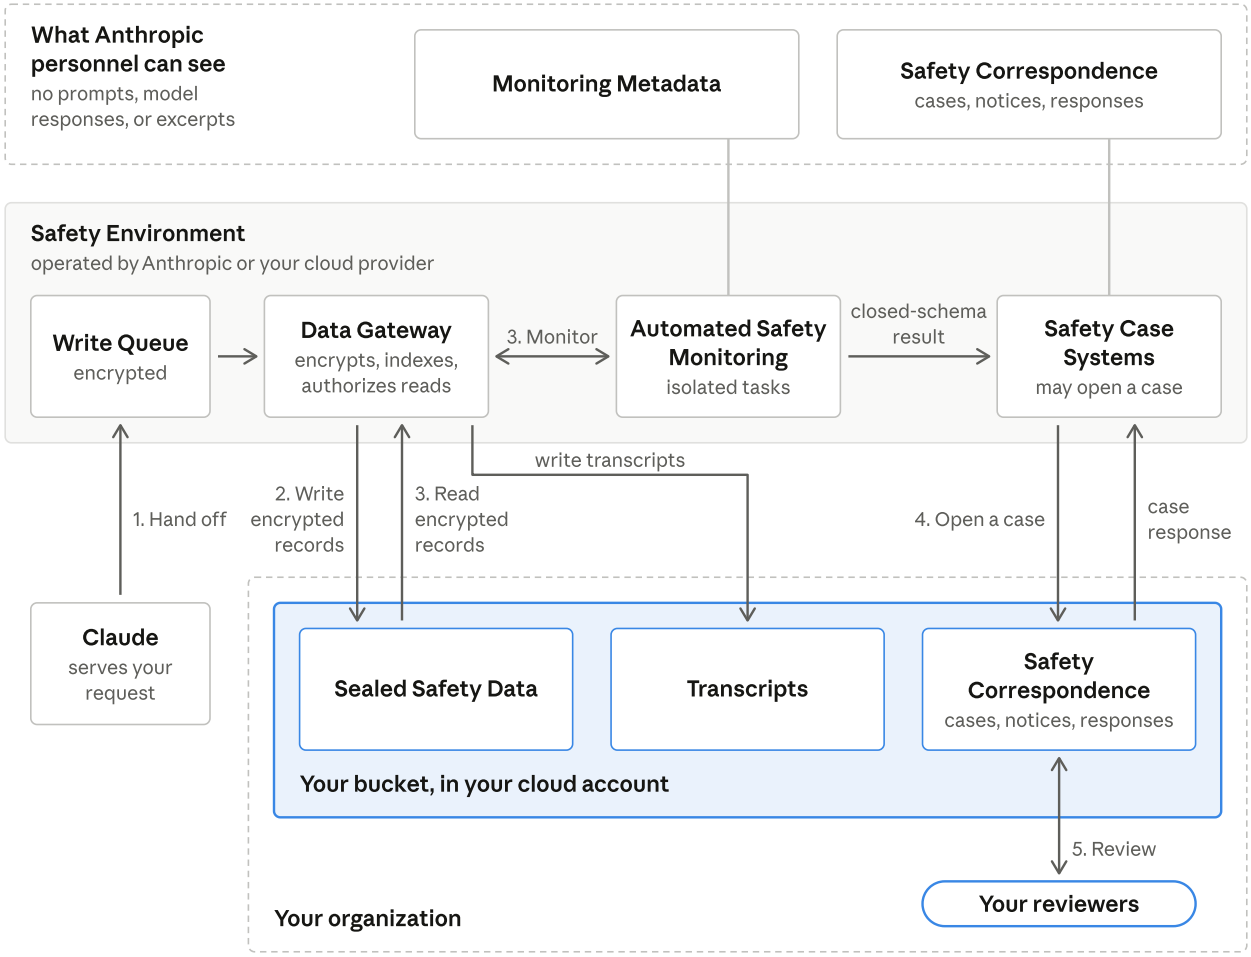

1. **Hand off.** When an exchange completes, Claude places a copy of the prompt and response on the encrypted Write Queue.
2. **Write.** The Data Gateway encrypts each record under its own key and writes it to your bucket. With CMEK, the record is also encrypted under your key. The Data Gateway keeps an index so that it can find and decrypt the record later. It also writes readable Transcripts for your team.
3. **Monitor.** Automated safety monitoring reads records through the Data Gateway and analyzes them.
4. **Open a case.** When automated safety monitoring finds a potential violation of the Usage Policy, Anthropic's Safety Case Systems may open a Safety Case and write it to your bucket. Anthropic may also send Safety Notices, which are messages such as a reminder to respond or additional context.

5. **Review.** Your security team investigates and writes a Safety Case Response to the same bucket. The Safety Case Systems read your response, record it, and close the Safety Case.

When you use EFS through Amazon Bedrock or Claude on Google Cloud (Google Managed), the Safety Environment is operated by your cloud provider. During public beta, some controls and data retention on these services may differ from what this paper describes. Refer to your cloud provider for details.

2.2 Data categories

Category	Contents	Where it is held	Who can read it
Sealed Safety Data	Prompts, responses, and model reasoning, encrypted in Anthropic's format	Your bucket	Anthropic's automated systems during automated safety monitoring
Transcript	A readable copy of a prompt and response, for your team to read	Your bucket	Your personnel
Safety Correspondence	Safety Cases, Safety Notices, and your Safety Case Responses. Never prompts and responses. A Safety Case may include a case ID, schema version, status and any close reason, and timestamps, such as when it was created, sent, updated, due, answered, and closed; organization, workspace or project, user profile, and storage IDs; the platform and cloud region; policy categories, such as cyber; any enforcement action; and the ID and file path of each related Transcript or file. Your Response includes the case ID, your determination, policy categories, and the time submitted.	Your bucket	Your personnel
		Anthropic	Authorized Anthropic personnel can read Monitoring Metadata and Safety Correspondence.
Monitoring Metadata	Includes identifiers, timestamps, token counts, object references, checksums, wrapped (encrypted) keys, classifier scores, and monitoring results	Anthropic	Anthropic

2.3 What EFS covers

EFS covers requests to the Messages API in a workspace configured for EFS regardless of model. It does not cover data that is stored by Claude Platform features for their own use, even within a Messages API request, such as code execution containers, Files API uploads, Batch API results, and Claude Managed Agents sessions. See API and data retention for each feature.

EFS applies from the time the configuration takes effect. Data written earlier remains where it was written until it expires.

3. Storage and retention

3.1 Writing to your bucket

You provide Customer Storage for Sealed Safety Data, Transcripts, and Safety Correspondence. Data is written as follows:

1. **Hand-off.** Claude's inference service places a copy of the prompt, the response, and metadata on the Write Queue in the Safety Environment. The path that serves requests does not persist this data. The queue is transient and encrypted. Only Anthropic-approved services can place records on the Write Queue, and only the Data Gateway reads from it.
2. **Write.** The Data Gateway encrypts each record and writes it to your bucket. Anthropic deletes prompts and responses promptly after the relevant response has been generated and they are written to your bucket.
3. **If a write fails.** Anthropic retries the write automatically. If the retries stop before it succeeds, the record stays in Anthropic's queue for up to 72 hours. A record that is not written in that time is deleted and cannot be recovered, whether the failure is on Anthropic's side or yours.

Anthropic writes only under fixed prefixes, which keep Sealed Safety Data apart from Transcripts and Safety Correspondence. Sealed Safety Data is written as ciphertext under `safeguards-data/` and should be left unchanged. Transcripts and Safety Correspondence are covered in more detail in Section 6.

3.2 How records are encrypted

Sealed Safety Data is encrypted at the application level by the Data Gateway, before it is written to storage, where it is encrypted at rest. Each record, including attached images and files, is encrypted with AES-256-GCM under its own random key. That key is stored only in wrapped form. The wrapped keys and the index stay outside your bucket, so the bucket alone does not yield plaintext. Object names and unencrypted fields carry identifiers and dates, not prompts or responses.

You can optionally use CMEK. With CMEK, each record is also encrypted under your key, and decrypting it depends on that key (Section 8).

Anthropic records the exact version of each Sealed Safety Data file it writes, and it reads records using that version. Authenticated encryption binds each record to its identity, so a file that has been altered, replaced, or restored from an older copy is not used in place of the original.

Transcripts and Safety Correspondence are written in readable form so that your team can review them, and they rely on the encryption configured on your bucket. On Claude Platform, the setup template creates a key for the bucket that is separate from your CMEK key. These files are kept apart from Sealed Safety Data, and each kind of file has its own prefix, so you can grant triage staff and reviewers separate access.

3.3 What Anthropic holds

Anthropic or your cloud provider holds Monitoring Metadata and a record of Safety Correspondence, including your Safety Case Responses. Neither contains prompts or responses. The Data Gateway keeps an index of each record, which contains no prompts or responses. The index is held in Anthropic's own cloud account or your cloud provider's. It is encrypted at rest, and the wrapped keys have an additional layer of application-level encryption. For each record, the index holds:

- The record's location, storage version, and checksum
- Organization and workspace identifiers, request and conversation identifiers, timestamps, and data-handling labels
- Each approved purpose, with its wrapped key and expiry time
- Labels added later, such as a review status

3.4 Retention and deletion

Sealed Safety Data is available for automated safety monitoring for up to 30 days. After that, the key and data are deleted.

Transcripts and Safety Correspondence expire under lifecycle rules on your bucket. The setup template expires Transcripts no later than 30 days after they are written, and keeps Safety Correspondence in your bucket. To keep Transcripts longer, copy them to another location. When a Safety Case is created, your team reviews the related Transcripts, not the Sealed Safety Data directly, so the Safety Case does not extend the retention of the records it references.

Early deletion. Anthropic requires you to retain all Sealed Safety Data for the full 30-day period, so that automated safety monitoring can analyze it. If data is deleted early repeatedly, requests may be refused. To stop Anthropic decrypting it, revoke Anthropic's access to the bucket or, with CMEK, to your key. After sustained unavailability or repeated access revocations, requests may be refused until access is restored (see Section 7.2).

Changing buckets or leaving EFS. When you register a new bucket, new records are written there, and the previous bucket must remain in place until its last record expires. If you stop using EFS, writes stop, existing records expire on schedule, and the bucket remains yours to keep or delete. Anthropic keeps Monitoring Metadata and Safety Correspondence.

3.5 Residency

You select a supported region for your bucket, and content is stored there, apart from records waiting to be written. Automated safety monitoring may run outside your bucket's region, but only for processing: Workers have no persistent storage, and any temporary storage is destroyed with the Worker (Section 5.3).

4. Access to your data

Anthropic's safety systems must maintain read and write access to your data to support automated safety monitoring. Anthropic's controls are designed to prevent any access outside of automated safety monitoring. Appendix A summarizes how these controls address the key threats to retained data.

4.1 Access to your bucket

Anthropic provides an infrastructure template that you review and apply in your own account: CloudFormation or Terraform for Amazon S3, and a setup script for Google Cloud Storage. On Amazon S3, it creates a narrow role for each function, limited to the bucket and tied to your organization by an external ID. On Google Cloud Storage, it gives Anthropic's service accounts, one for each function and shared by all customers, object access limited to your organization's folder in the bucket.

Anthropic's software uses these identities with short-lived credentials. These credentials cannot change the bucket's configuration or policies. Removing these identities or grants ends Anthropic's access. For Sealed Safety Data, requests in the affected workspace may then be refused (Section 7.2).

Before any data flows, Anthropic verifies the bucket, encryption, and each identity against the template. For Sealed Safety Data, these checks are repeated daily, and any change from the template is reported to Anthropic. Anthropic has no permission to correct it.

Cloud storage providers may offer controls for object-level logging. If configured, your logs record each write, read, and deletion by Anthropic's systems. These logs show that an access occurred, not what was done with the content afterward.

4.2 How a read is authorized

Every read of Sealed Safety Data through the Data Gateway follows the same order.

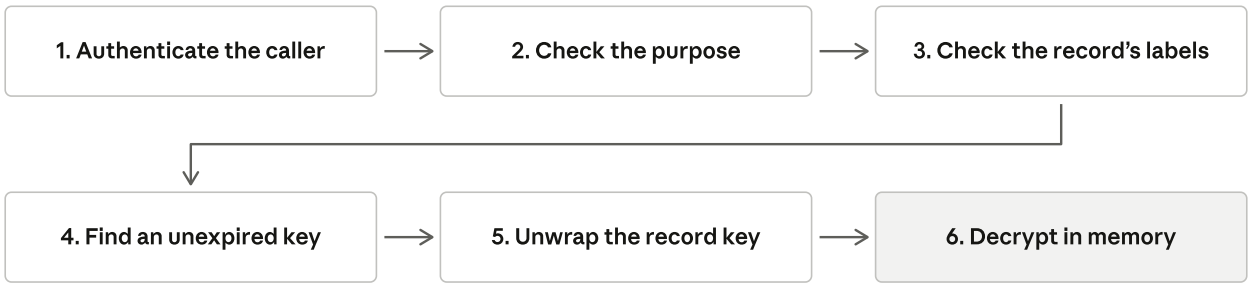

A failed check denies the read. With CMEK, step 5 depends on your key.

The caller is authenticated and must be approved for the stated purpose. A policy check is made against the record's labels, such as Customer Safety Case Review or CMEK, and a denied record is reported as not found, with no request made to your bucket. The record must still have an unexpired key for that purpose. The record's key is then unwrapped and the record is decrypted in memory. With CMEK, unwrapping depends on your key.

The following controls apply to each read:

- **Purpose-bound access.** The purpose is checked when the access happens, so a caller cannot reach data for a purpose it was not granted.
- **Fail-closed enforcement.** Access is denied if the policy check cannot be completed.
- **Mandatory justification.** No content can be read without a recorded reason.
- **Access records.** The Data Gateway, not the caller, records each permitted read in its own access log.

4.3 Legal requirements

The exception to these access limits, and to the storage and retention described in Section 3, is where Anthropic is required by law to access and/or retain data (for example, material reported to NCMEC under 18 U.S.C. § 2258A).

5. Automated safety monitoring

This section describes how the inputs and outputs of automated safety monitoring are constrained. Incoming safety monitoring tasks are checked against the Scope Policy before they run, each task runs on an isolated Worker, and results are validated against a closed schema. These controls require no configuration.

5.1 Tasks

A safety monitoring task is a request to analyze retained records, written as a prompt or fixed-schema description of what to analyze and for which harm. Some task prompts are written for a specific investigation, because emerging threats may not match an existing template. Inputs are therefore only partially constrained by schema, and every task is subject to the Scope Screener.

Task governance. The Anthropic Control Plane accepts tasks only from Anthropic's automated safety systems and from authenticated, authorized Anthropic personnel. In both cases, the analysis itself is automated. Submitting a task does not let the human requester view your prompts or responses. The requester receives only the task's result.

A dedicated access role allows personnel to submit tasks. It is granted individually with a second person's approval and does not itself provide access to your prompts or responses. For recurring tasks, task definitions are maintained in Anthropic's source control and changed through Anthropic's standard code review process. Tasks written for a specific investigation are not gated by this code review process before they run, given the need to respond quickly to emerging threats. Both are subject to the Scope Screener, and Anthropic retains task prompts internally for a limited period for post-hoc review. Task prompts are not shared with customers because they can contain sensitive information.

Tasks that span organizations. Each Safety Case covers only the receiving organization's records, and no other customer sees your records or the results. Most tasks concern a single organization, but a task may analyze records from

more than one, because the same actor can split activity across several accounts so that no single one looks harmful. Abuse coordinated in this way can only be detected through aggregate analysis. The closed-schema results go only to Anthropic.

5.2 Scope screening

Scope Policy. Anthropic maintains a Scope Policy that lists the purposes for which a task may run. Tasks may analyze retained content only to detect and investigate violations of Anthropic's Usage Policy, or when a customer has explicitly requested or consented to a review of its own traffic.

Scope Screener. Before a Worker executes any task, the Scope Screener evaluates it against the Scope Policy and produces a verdict of ACCEPT or REJECT. It is designed to reject a task in full if any part of it is out of scope.

Task audit log. The Scope Screener records each verdict in a task audit log, and an accepted task does not proceed until that record is written. The record contains no task text and none of your prompts or responses.

Overrides. A rejected task can run only through a two-person override, in which a second person at Anthropic reviews and approves that exact task. Only personnel who hold the access role described in Section 5.1 can request or approve an override, and no one can approve their own request.

5.3 Workers

Per-task ephemeral Workers. Each task runs on a Worker, a container or virtual machine provisioned for that task alone and terminated when the task completes or reaches its time limit. Workers mount no persistent storage. Any temporary local storage is task-scoped and destroyed with the Worker.

Reads. Workers read retained content only through the Data Gateway, as described in Section 4.2, and hold no decryption keys or direct access to storage. The Data Gateway logs each read with the identifier of the task that made it.

Network. Egress is deny-by-default, and a Worker has no general path to the internet. It can reach only an allowlist of services, which is enforced at the network layer outside the Worker.

Inference. Workers may use Claude to analyze retained content, such as summarizing a set of records and assessing them for potential harm. These inference calls do not retain the content they process.

Change management. Worker and Scope Screener software is changed through Anthropic's standard code review process, and is built and deployed by automated pipelines.

5.4 Results

The result schema is fixed and is designed to not contain prompts or responses, and every result is validated against the schema before it is transmitted. String fields are constrained to predefined, enumerated values, numbers are bounded, and identifier fields are validated against the format for their identifier type. A result can cite only records that the task actually read. Fields vary by task, within the fixed schema. An illustrative example follows:

```
{
    "task_id": "<identifier>",
    "harm_label": "<none | low | medium | high>",
    "confidence_label": "<low | medium | high>",
    "policy_categories": ["<enumerated value>"],
    "harm_score": 0.82,
    "record_ids": ["<identifier>"]
}
```

Results are Monitoring Metadata, as defined in Section 2.2. Workers also emit operational metrics and logs, which are designed to hold identifiers, counts, and error types, not content.

5.5 Fail-closed behavior

Each enforcement point is designed to fail closed. Automated safety monitoring is an offline system and is not in the path of customer inference, so a failure affects only the tasks involved, not customer requests. Section 7.2 covers loss of access to a customer's key or bucket.

- If the Scope Screener is unavailable or returns no verdict, the Worker does not execute the task.
- If an ACCEPT verdict cannot be recorded in the task audit log, the Worker does not execute the task.
- If an override is invalid or revoked, the Worker does not execute the task.
- If the inference endpoint is unreachable, the task fails.
- If a result fails validation, it is rejected and not transmitted.
- If a result cannot be delivered, delivery or the task is retried.

6. Safety Case Review

Your designated team (such as your security team), not Anthropic, reviews Safety Cases. Anthropic's controls are designed to prevent Anthropic personnel from reading your Transcripts and Sealed Safety Data.

If content is flagged by automated safety monitoring, Anthropic's systems may open a Safety Case. It contains identifiers and details, such as status, timestamps, policy categories, any enforcement action taken, and references to the related Transcripts, but no prompts or responses. Your team investigates and responds, and Anthropic personnel do not access the related Transcripts.

Safety Notices. Anthropic may also send a Safety Notice, a message that usually relates to one or more Safety Cases, such as a reminder to respond or additional context. A Safety Notice contains no prompts or responses.

Safety Case lifecycle

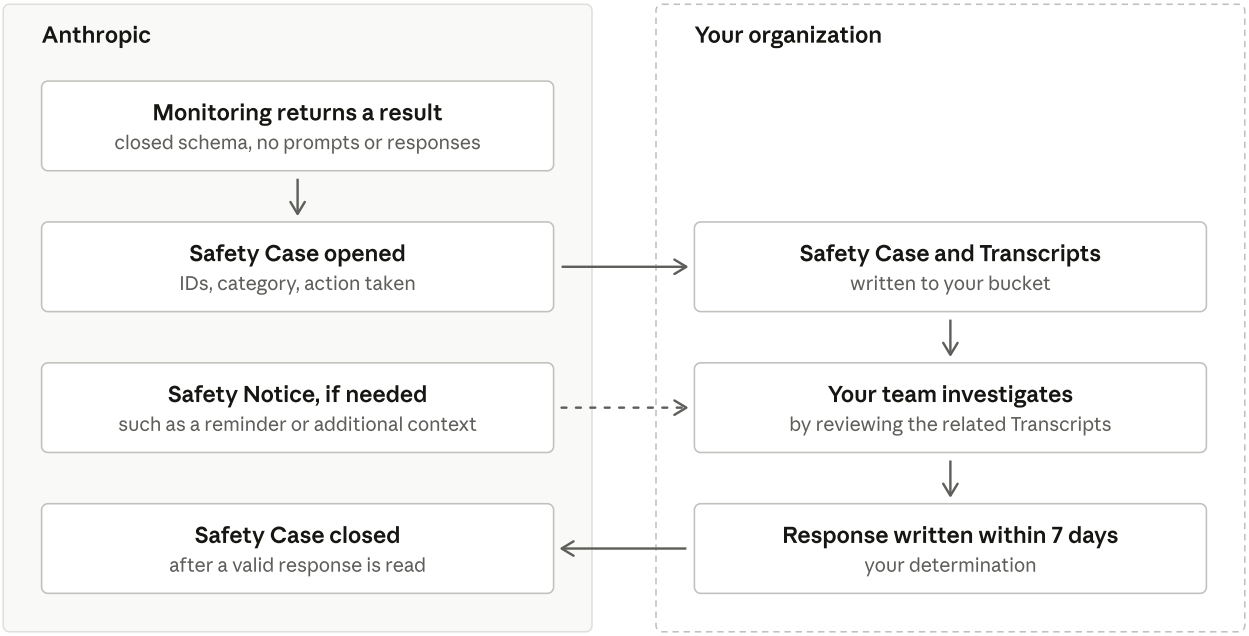

To investigate a Safety Case, you receive the information you need in your storage destination. You should expect the following content, written as JSON files under the following paths.

Path	Contents	Suggested access
`transcripts/`	Readable copies of the Transcripts that Safety Cases reference and, depending on your storage provider, of other requests in the workspace, in the public Messages API format	Reviewers
`review-cases/<case_id>.json`	One file per Safety Case, covering only the receiving organization's records. It contains no prompts or responses.	Triage staff and reviewers
`review-notices/<notice_id>.json`	Safety Notices from Anthropic, such as a reminder to respond to a Safety Case. They contain no prompts or responses.	Triage staff and reviewers
`review-responses/<case_id>.json`	One file per response, written by you. It contains your determination on the Safety Case, with optional determinations for individual Transcripts.	Reviewers

You are responsible for access to this content, for example by limiting bucket access to reviewers, encrypting the bucket with a key you manage, and enabling access logging. Each kind of file has its own prefix, so access to Safety Correspondence and to Transcripts can be granted separately. Anthropic's identities cannot read Transcripts and cannot write or change your Safety Case Responses.

7. Shared responsibility

The properties described in this paper depend on both Anthropic and you.

7.1 Who does what

Area	Anthropic	Customer
Safety Environment	Operates the Data Gateway, Automated Safety Monitoring, and Safety Case Systems. Maintains the Scope Policy and screens each task.	Configures EFS on its workspaces
Review	May open Safety Cases, and delivers them to you	Reviews and responds to Safety Cases within 7 days, and monitors Safety Notices
Storage	Writes to the configured destination and verifies it against the template	Keeps Customer Storage available, and protects Transcripts and Safety Correspondence held there
Keys (optional)	Uses a customer-managed key when provided	Keeps the key available, and treats key policy changes as production changes

7.2 When requests may be refused

Monitoring depends on access to your bucket and, with CMEK, your key. If either fails, Anthropic does not fall back to its own storage or its own keys. Requests continue for a bounded period and may then be refused until access is restored. Requests resume once access returns.

Condition	Window	What to expect
Your key is unavailable, if configured	About 1 hour on Claude Platform	Requests are served on cached key material, and then refused with an error.
Your bucket for Sealed Safety Data is unavailable	24 hours	Requests are served while records wait in the Write Queue. After 24 hours, plan for requests to be refused until the bucket is available.
An unexplained bucket failure specific to your account, or an explicit denial of access	8 hours from notice	Acknowledge the notice and give an expected resolution time, or requests may be refused before the 24 hours end.
A Safety Case is unanswered	7 calendar days	Anthropic may stop serving requests.

On Claude Platform, the unavailable key window is shortest because a record can wait in a queue for a bucket, but cannot be encrypted later under a key that Anthropic cannot reach. If the key and the bucket fail together, the key window applies. These windows apply to reads as well as writes. These windows apply to any cause, including a cloud outage. Anthropic may end a window early on signs of misuse or repeated failures.

8. Customer-Managed Encryption Keys

CMEK is optional. With it, decrypting Sealed Safety Data requires a key in your own key management service (KMS), and each use of that key can be recorded in your logs. EFS uses Anthropic's existing CMEK capability, so this section covers only what is specific to EFS.

8.1 What your key covers

Your key covers Sealed Safety Data, including records in the Write Queue. It does not apply to Transcripts or Safety Correspondence, which rely on the encryption configured on your bucket.

With CMEK, a record opens only with both your key and Anthropic's keys. Neither your key alone nor the stored data alone is sufficient to read a record, so a compromised key does not decrypt records by itself. Your key is tied to your organization, and only Anthropic's Key Service can call it.

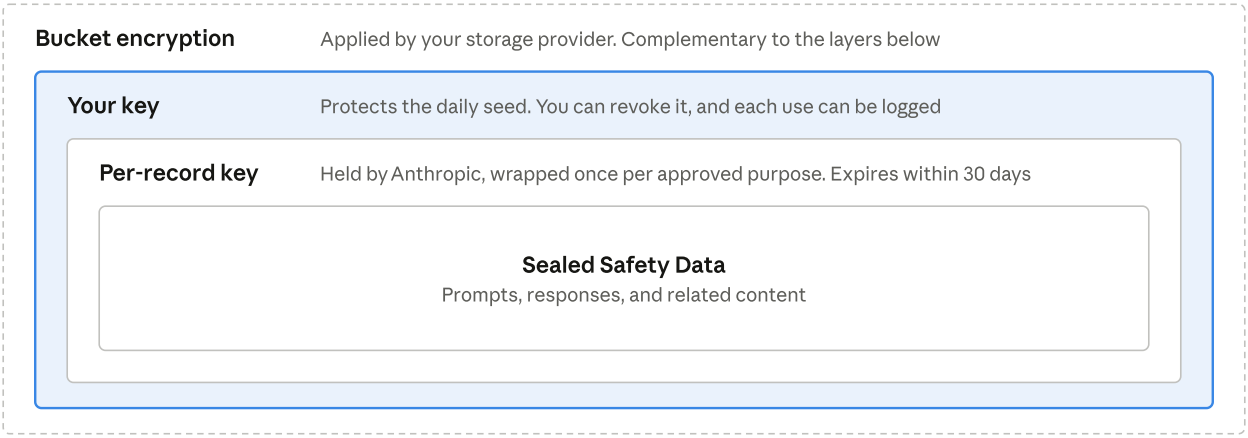

8.2 Key hierarchy

- **Per-record key.** Each record has its own random key. It is wrapped first under a key derived from the daily seed, then once per approved purpose under a

key from Anthropic's own hierarchy. A wrapped key copied to another record or used for another purpose cannot be unwrapped.

- **Daily seed.** For each workspace and day, the Key Service generates a random seed and encrypts it under your key. The keys for the first wrap are derived from the seed using HKDF-SHA-256 and are never stored. A compromised seed alone reveals no prompts or responses. It is useful only to a party that also holds Anthropic's purpose keys, the index, and the ciphertext, and then covers about one day of one workspace.
- **Your key.** Your KMS sees only operations on the daily seed, never content or content keys.

The Customer-Managed Encryption Keys: Cryptographic Design paper describes the underlying design. EFS adds the per-record key and the per-purpose wrap.

8.3 Revocation

On Claude Platform, if you disable the key or remove Anthropic's access, within approximately one hour Anthropic will no longer be able to decrypt Sealed Safety Data, and requests will stop. Revocation does not delete data or reverse reads that have already occurred. See What happens when you rotate, disable, or delete your encryption key.

Anthropic recommends that you alert on denied key calls and unexpected key use, treat key policy changes as production changes, and test by disconnecting a non-production workspace.

8.4 Rotation and setup

You rotate the key in your own KMS, with no change to your EFS configuration. Earlier key versions must remain enabled. A key cannot be removed or replaced once attached, so on Claude Platform you change keys by moving traffic to a new workspace.

Anthropic validates a key when it is added. The setup guides for AWS KMS, Google Cloud KMS, and Azure Key Vault list the requirements, and Troubleshoot CMEK covers validation failures.

Appendix A. Controls reference

This table summarizes how EFS addresses the key threats to retained data.

Threat	Control	Section
Records are copied from your bucket	Application-level encryption, with keys and index held apart	3.2
A stored file is altered or replaced	Versioned reads and authenticated encryption	3.2
Records are kept too long	Reads end at expiry, then key and object deletion	3.4
Confused deputy: one organization causes Anthropic to act on another's storage or key	Identities, folder-limited grants, and keys bound to your organization	4.1, 8.1
A bucket or key is misconfigured	Checks before use, and daily for Sealed Safety Data	4.1, 8.4
A person uses a bucket identity	Identities assumed by services, not issued to people	4.1
Data is read for an unapproved purpose, such as training	Purpose-bound, fail-closed reads	4.2
Anthropic personnel seek to read content	Data Gateway reads refused for records under Customer Safety Case Review	4.2, 6
A task is submitted for another purpose	Scope Screener, with two-person override	5.2
A Worker is compromised or manipulated	Per-task compute and deny-by-default egress	5.3
Retained content contains a prompt injection aimed at automated safety monitoring	A manipulated analysis could miss misuse or flag it incorrectly, but its result must still pass schema validation and the Worker can reach only allowlisted services. These controls limit what it could signal.	5.3, 5.4
Content leaves through results	Closed-schema validation	5.4
Another organization's data appears in a Safety Case	Each Safety Case covers only your records	6
Your key is compromised	The key alone does not decrypt records	8.1
You need to end Anthropic's ability to decrypt	Key revocation	8.3

Appendix B. Infrastructure security

The services Anthropic operates run with the following baseline controls:

- **Encryption in transit.** External connections use TLS 1.2 or higher, and internal services that handle retained data use mutual TLS.
- **Tenant isolation.** Each retained record carries organization and workspace identifiers. Customer-facing access is authorized against the requesting organization, and cross-organization reads are denied.
- **Hardened infrastructure.** Production environments use least-privilege service accounts and role-based access control, pod security standards, node and image freshness monitoring, continuous vulnerability management, host-level execution telemetry and detection, safe-deserialization enforcement, and secure build and supply-chain controls.