How do you know Data is Actually Deleted?

Proving Data Deletion in AWS

May 8, 2026
Flock AWS Account Team

Summary

One of the questions customers ask is: "How do I know my data is actually deleted?" This is a fair question — especially for organizations handling sensitive data subject to compliance requirements like CJIS. This white paper explains how AWS ensures data is permanently removed from Amazon S3 buckets, and how Flock's use of these services in their AWS environment — S3 Lifecycle Policies, AWS CloudTrail, Amazon CloudWatch, and Amazon OpenSearch — work together to provide verifiable proof that deletion has occurred.

The short answer: AWS provides multiple layers of evidence — from API-level audit logs to automated deletion policies to real-time monitoring — that together create an airtight chain of custody that your data has been permanently removed.

How S3 Deletion Actually Works

What Happens When You Delete an Object

When an object is deleted from Amazon S3, the behavior depends on whether versioning is enabled:

- **Without versioning**: The object is permanently deleted. It cannot be recovered. AWS confirms the deletion by returning a "deleted" result in the API response.
- **With versioning**: A "delete marker" is placed on the object, making it invisible to normal requests. The previous version still exists until you explicitly delete it or a lifecycle policy removes it.
- **Permanent deletion of a version**: When you delete a specific object version (or when a lifecycle policy expires it), that version is permanently removed from AWS infrastructure. It cannot be recovered by anyone — including AWS.

Key Point: Deletion Is Irreversible

AWS documentation is explicit: "Deleting an object can't be undone." Once an object or object version is permanently deleted, there is no mechanism — not even for AWS engineers — to retrieve it. This is by design and is enforced at the infrastructure level.

S3 Lifecycle Policies: Automated, Guaranteed Deletion

What They Do

S3 Lifecycle Policies are configurable rules that automatically manage objects throughout their life.

How You Know It Worked

1. **The policy is declarative and auditable**: Your lifecycle configuration is stored as part of the bucket's settings. You can retrieve it at any time to confirm what rules are in place.
2. **S3 executes lifecycle rules automatically**: Once configured, AWS guarantees execution. Objects that meet the expiration criteria will be deleted — there is no manual step that could be skipped.
3. **CloudTrail logs the lifecycle actions**: Every object expiration triggered by a lifecycle policy generates a CloudTrail event, giving you a timestamped record of what was deleted and when.
4. **You can verify absence**: After expiration, attempting to retrieve the object returns a "404 Not Found" or "NoSuchKey" error — direct proof the object no longer exists.

EXAMPLE-A typical compliance-oriented lifecycle policy might include:

- *Permanently delete noncurrent versions after 30 days*
- *Remove expired delete markers after 1 day*

This ensures that even versioned objects are permanently purged within a defined, auditable timeframe.

AWS CloudTrail: Your Deletion Audit Trail

What CloudTrail Captures

AWS CloudTrail records every API call made in your AWS account. For S3 deletion, this includes:

- **DeleteObject** — Records who deleted what object, from which bucket, at what time, and from what IP address.
- **DeleteObjects** (batch) — Same as above, but for bulk deletions.
- **Lifecycle-triggered expirations** — In addition to Amazon CloudTrail, Amazon S3 Server access logging provides detailed records for the requests that are made to a bucket. When S3 automatically deletes objects based on lifecycle rules, they are captured in Amazon S3 server access logs.
- **PutBucketLifecycleConfiguration** — Records when lifecycle policies are created or modified, and by whom.

Why This Matters for Proving Deletion

AWS CloudTrail and S3 Server Access Logging gives you:

- **Who** performed or triggered the deletion
- **What** was deleted (bucket name, object key, version ID)
- **When** it happened (precise timestamp)
- **How** it was initiated (console, CLI, SDK, or lifecycle automation)
- **Tamper-proof records**: CloudTrail supports log file integrity validation. You can cryptographically verify that log files have not been modified, deleted, or forged after delivery. This means your deletion records are forensically sound.

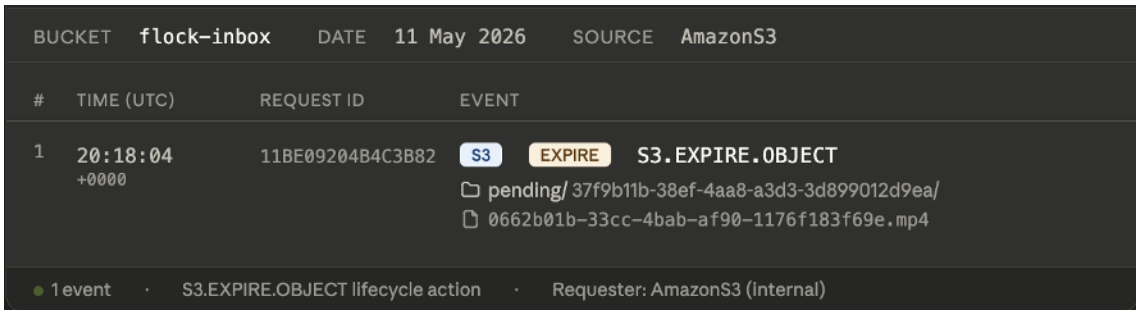

Amazon CloudWatch: Real-Time Deletion Monitoring

Metrics That Prove Deletion Is Happening

CloudWatch provides S3 bucket-level metrics that give you visibility into deletion activity:

- **NumberOfObjects**: Track the total object count in a bucket over time. A declining count after lifecycle policies take effect is direct evidence of deletion.
- **BucketSizeBytes**: Monitor total storage size. Decreasing storage confirms objects are being removed.
- **Custom Metrics**: You can publish custom metrics from CloudTrail events to track deletion rates, failed deletions, or anomalies.

Alarms and Anomaly Detection

You can configure CloudWatch to:

- **Alert you if deletions stop**: If your lifecycle policy should be deleting objects daily but the object count plateaus, an alarm can notify you of a potential issue.
- **Detect anomalous deletion patterns**: Unusual spikes in deletion activity could indicate unauthorized access — CloudWatch anomaly detection can flag these automatically.
- **Create dashboards**: Build visual dashboards showing deletion trends over time for compliance reporting.

Amazon OpenSearch Service: Audit logs

Enabling Audit logs in Amazon OpenSearch cluster enables tracking user activity on OpenSearch clusters, including authentication success and failures, requests to OpenSearch, index changes, and incoming search queries. These audit logs are automatically vended to CloudWatch Logs. When indices are deleted in Amazon OpenSearch Service, the audit log captures this as an INDEX_EVENT category on the transport layer. Additional key fields for index deletion events include the index/indices being deleted, who performed the deletion and when it occurred providing chain of custody documentation.

Auditability for Amazon S3 and Amazon OpenSearch objects

When all four services are combined, you get an end-to-end deletion verification system:

	Layer	What It Proves	Evidence Type
1	**S3 Lifecycle Policy**	Deletion rules are in place and enforced	Configuration (auditable, versioned)

2	**S3 API Response**	Individual objects were deleted successfully	Real-time confirmation
3	**AWS CloudTrail**	Who deleted what, when, and how	Tamper-proof audit logs
4	**Amazon CloudWatch**	Deletion is happening at expected rates	Metrics, trends, and alarms
5	**Audit Logs for Amazon OpenSearch**	Who deleted what, when, and how	Tamper-proof audit logs

AWS Infrastructure-Level Guarantees

AWS provides infrastructure-level assurances:

Physical Media Handling

- AWS storage devices are decommissioned using techniques detailed in NIST 800-88 ("Guidelines for Media Sanitization").
- Storage media is cryptographically erased or physically destroyed when decommissioned.
- AWS data centers undergo regular third-party audits (SOC 1, SOC 2, SOC 3, ISO 27001) that verify these processes.

Encryption at Rest

- Flock's is encrypted with AWS KMS (which is recommended and relevant to your CJIS compliance posture), deleting the encryption key renders all encrypted data permanently unreadable.
- This provides a "crypto-shredding" capability: destroy the key, destroy access to the data.

The Nitro System and Zero-Trust Architecture

As noted in your CJIS compliance documentation, the AWS Nitro System ensures that AWS personnel cannot access customer data on the underlying infrastructure. This means:

- No AWS employee can recover your deleted data
- The hardware itself enforces isolation
- Customers retain complete control and ownership of their data

Compliance and Regulatory Alignment

CJIS Compliance

AWS is the only cloud provider that is CJIS compliant in all 50 states. The zero-trust security model built on AWS KMS and the Nitro System means that when you delete Criminal Justice Information (CJI), you can be confident that:

- The data is permanently removed from S3
- No AWS personnel had access to it before, during, or after deletion
- The deletion is logged and verifiable through CloudTrail

Additional Compliance Frameworks

AWS's deletion mechanisms align with requirements from:

- **GDPR** (Right to Erasure / Right to be Forgotten)
- **SOC 2** (Data retention and disposal controls)
- **ISO 27001** (Information security management)
- **NIST 800-53** (Security and privacy controls)
- **FedRAMP** (Federal security requirements)

Conclusion

The question "How do you know if data is actually deleted?" has a clear, multi-layered answer on AWS:

- **You configured it**: Lifecycle policies define exactly when and how data is removed.
- **AWS executed it**: S3 automatically and irreversibly deletes objects per your policies.
- **You have proof it happened**: CloudTrail provides tamper-proof, timestamped records of every deletion.
- **You can see it in real time**: CloudWatch metrics show deletion activity as it occurs.
- **You can search the evidence**: OpenSearch gives you instant access to your complete deletion history.
- **The infrastructure guarantees it**: The Nitro System, KMS encryption, and NIST 800-88 media handling ensure data cannot be recovered by anyone — including AWS.

For an organization like Flock Safety operating under CJIS requirements, this combination of automated policies, cryptographic verification, real-time monitoring, and

infrastructure-level guarantees provides the highest level of assurance that deleted data is truly, permanently gone.

References/Supporting Documentation:

- **Deleting objects from a directory bucket** — docs.aws.amazon.com/AmazonS3/latest/userguide/directory-bucket-delete-object.html
- **Deleting multiple objects** — docs.aws.amazon.com/AmazonS3/latest/userguide/delete-multiple-objects.html
- **Deleting objects (S3 console)** — docs.aws.amazon.com/AmazonS3/latest/userguide/delete-objects.html
- **Nitro**- https://aws.amazon.com/ec2/nitro/
- AWS Digital Sovereignty Pledge: Control without compromise