



THREAT ADVISORY

Threat Advisory: Public Cloud Storage Exposure

Object storage left world-readable remains one of the most common causes of data exposure we find. The detection and prevention controls are inexpensive and rarely switched on.

PUBLISHED

January 20, 2026

READING TIME

8 minutes

AUTHOR

SCIAT AFRICA Research

ABSTRACT

Publicly readable object storage continues to account for a meaningful share of the data exposure we encounter, and it is almost always a configuration default rather than an attack. This advisory sets out how the exposure arises, how to detect it across an estate, and the preventive controls that make it structurally difficult to recreate.

Topics: Cloud Security · Data Protection · Configuration

WHAT TO TAKE AWAY

- Account-level public-access blocks are available on the major platforms, cost nothing, and are frequently left off.
- Exposure is usually created by a convenience decision during development and never revisited.
- Pre-signed or time-limited URLs cover almost every legitimate reason a bucket is made public.
- Detection belongs in continuous configuration monitoring, not in an annual review.
- An exposed bucket holding personal data is a notifiable breach, and the clock starts at awareness.

CONTENTS

01 How the exposure happens

02 Finding it across an estate

03 Preventing recurrence structurally

04 If data was exposed

01 How the exposure happens

Very little of this is adversarial. A developer needs an image served to a browser and makes the container public because that works immediately. A data team shares an export with a partner and grants broad read because scoping it properly takes a ticket. A legacy static site is migrated with its permissions intact. A backup lands in a bucket whose policy was written for something else.

In each case the exposure is created once, satisfies an immediate need, and is never revisited, because nothing breaks and nothing alerts. Discovery is trivial for anyone looking — storage endpoints are predictable and continuously enumerated by third parties.

02 Finding it across an estate

Every major platform exposes the information needed to answer "what is publicly readable" through its own API, and the native configuration-compliance services will report it continuously. The work is less technical than organisational: enumerate the accounts, subscriptions and projects that actually exist, including the ones created for a proof of concept and never decommissioned.

Shadow accounts are the common gap. An inventory built from the billing relationship rather than from memory is the only reliable starting point, and it frequently surfaces environments nobody currently owns.

- Enumerate all accounts, subscriptions and projects from the billing root, not from documentation.
- Enable the platform's configuration-compliance service in every region in use, not only the primary one.
- Alert on the transition to public, rather than reporting the state monthly.
- Record an owner for every storage container; unowned containers are the ones that stay exposed.

03 Preventing recurrence structurally

Turn on the account-level public-access block. On the major platforms this is a single setting that prevents individual containers from being made public regardless of their own policy, and it is the highest-value control available here. Where a genuine public-serving use case exists, isolate it in its own account with an explicit exception, so the default stays safe everywhere else.

For the legitimate need that drove most of these decisions — serving a file to a specific person for a limited time — use pre-signed or time-limited URLs, or a content delivery network with an origin that is not itself public. This covers the overwhelming majority of cases and removes the incentive to loosen the policy.

Finally, put the check in the pipeline. Infrastructure-as-code policy scanning catches a public bucket at pull-request time, which is an order of magnitude cheaper than catching it in production and immeasurably cheaper than catching it in a disclosure.

04 If data was exposed

Close the exposure, then establish scope before anything else: what objects were present, for how long, and what access logs exist. Server access logging is often not enabled, which makes it impossible to say whether anything was retrieved — an answer that is itself material to the notification assessment and should be stated honestly rather than assumed either way.

If personal data was involved, the Data Protection Act's notification duties engage, with the seventy-two-hour clock running from awareness. Document the timeline carefully: when the exposure began, when it was detected, when it was closed, and what the logs do and do not show. That record is what a regulator will ask for.

About this publication

This document is published by SCIAT AFRICA, a cybersecurity and digital forensics agency based in Nairobi, Kenya. It is free to read, print and circulate inside your organisation. The web edition, which carries the same content, is at <https://sciatafrica.com/research/advisory-cloud-storage-exposure>.

It is general guidance, not legal advice and not a substitute for an assessment of your own environment. Where it refers to a law, regulation or standard, read the instrument itself before relying on it – and talk to us if you want help applying it to your estate.

SCIAT AFRICA

info@sciatafrica.com · +254 794 946633 · <https://sciatafrica.com>