pravidha

Pramana

Release Notes

Governance evidence for Snowflake — classification, policy coverage, drift, audit packs

Document	Pramana — Release Notes
Version	v1.0
Date	17 November 2026
Author	Pravidha — Snowflake Select Service Partner
Status	Released
Classification	Public

© 2026 Pravidha. This document and the software it describes are Pravidha intellectual property.

Contents

Pramana 1.0
- What is in it
- What you need
- Known limits in this version
- Upgrades

© 2026 Pravidha · Public

What each version of the Pramana application contains, what changed, and what to know before relying on it. Newest first.

Pramana 1.0

The first release.

What is in it

- **Thirteen controls** across classification, policy, access, retention and data quality, each with a fixed identifier that can be cited in an audit report.
- **Evidence packs** for UK GDPR, DORA and ISO 27001, produced on every run, kept as history, and sealed with a SHA-256 that a copy can be checked against. A framework of your own is a few rows of configuration.
- **Classification proposals.** Columns that Snowflake's classifier or your naming rules recognise as sensitive and that carry no label are proposed to a data steward, who approves or rejects each one.
- **Governed tags.** Retention is checked out of the box: every table holding sensitive data carries a retention class, and each class is one you defined. Any other table-level tag you govern by, such as residency or ownership, is one more row.
- **Drift.** Every night is compared with the one before. Changes that weaken your position are flagged and stay open until someone acknowledges them.
- **Remediation as statements.** Every fixable finding comes with the statement that closes it, in an order a change can follow, including templates for policies that do not exist yet.
- **Three consoles**, for administrators, stewards and read-only users, in Snowsight's light and dark modes.
- **A built-in sample estate** with planted defects, which loads without any grant and shows every screen with something in it.
- **A view for GRC tooling**, with a role that can read that view and nothing else.
- **Optional link to Vigil**, to carry reconciliation results in every pack.

What you need

- A Snowflake account on Enterprise Edition or higher, for scanning your own account. The built-in sample runs on any edition.
- `ACCOUNTADMIN` once, for three grants.

Known limits in this version

- **Results follow Snowflake's account metadata**, which typically runs up to two or three hours behind the account. A fix made today is confirmed by a later run.
- **The first policies an account attaches can take many hours to be listed** by Snowflake. Until they are, Pramana reports those columns and tables as unprotected, and notes the likely cause where it can tell. When the listing catches up the findings clear, and the attachments are not reported as drift.
- **Pramana proves that a policy is attached, not what its rule returns.** Reviewing the logic of a masking or row access policy remains yours.
- **Name rules recognise names.** A sensitive column with an uninformative name is found only where Snowflake's classifier has examined its table.

© 2026 Pravidha · Public

- **Very large estates.** This version writes each snapshot in batches. If your estate runs to hundreds of thousands of columns, talk to us before scheduling a nightly run.
- **No alerts of its own.** A Native App cannot be given a notification integration, so being told about weakening drift is an alert you create in your own schema. A ready-made one is in Operations and Troubleshooting.

Upgrades

Snowflake applies new versions in place. Your policy pack, every snapshot and every evidence pack are kept, and a schedule that was running keeps running.

Because Snowflake upgrades every installation in place, the current version is the supported one. Each release is described on this page.

© 2026 Pravidha · Public