SNOWFLAKE SECURITY

SNOWFLAKE SECURITY AT A GLANCE

Enterprise for Sensitive Data (ESD)

Access

- All communication secured & encrypted
- TLS 1.2 encryption for all client communications
- Option for encryption in both trusted and untrusted networks
- IP Whitelisting
- Integration with AWS PrivateLink

Authentication

- Password Policy enforcement
- Multifactor Authentication
- SSO using SAML 2.0 Federated Authentication
- Key Pair (link to snowsql, but supported all over)

Authorization

- Flexible & granular authorization controls
- RBAC for data and actions
- OAuth2.0 delegation
- Secure views and UDFs to protect information access

Data

- Encrypted at rest
- Hierarchical key model rooted in HSM
- Automatic key rotation
- Yearly re-keying
- Tri-Secret Secure (BYOK)
- Time Travel 1-90 days
- Fail Safe

Infrastructure

- AWS, Azure Physical Security
- AWS, Azure Redundancy
- Regional Data Centers
  - US
  - EU
  - AP

Snowflake Operational Controls

- FedRAMP / NIST 800-53
- HIPAA
- PCI
- ISO/IEC 27001
- SOC2 Type 2
- SOC1 Type 2

SNOWFLAKE SECURITY AT A GLANCE

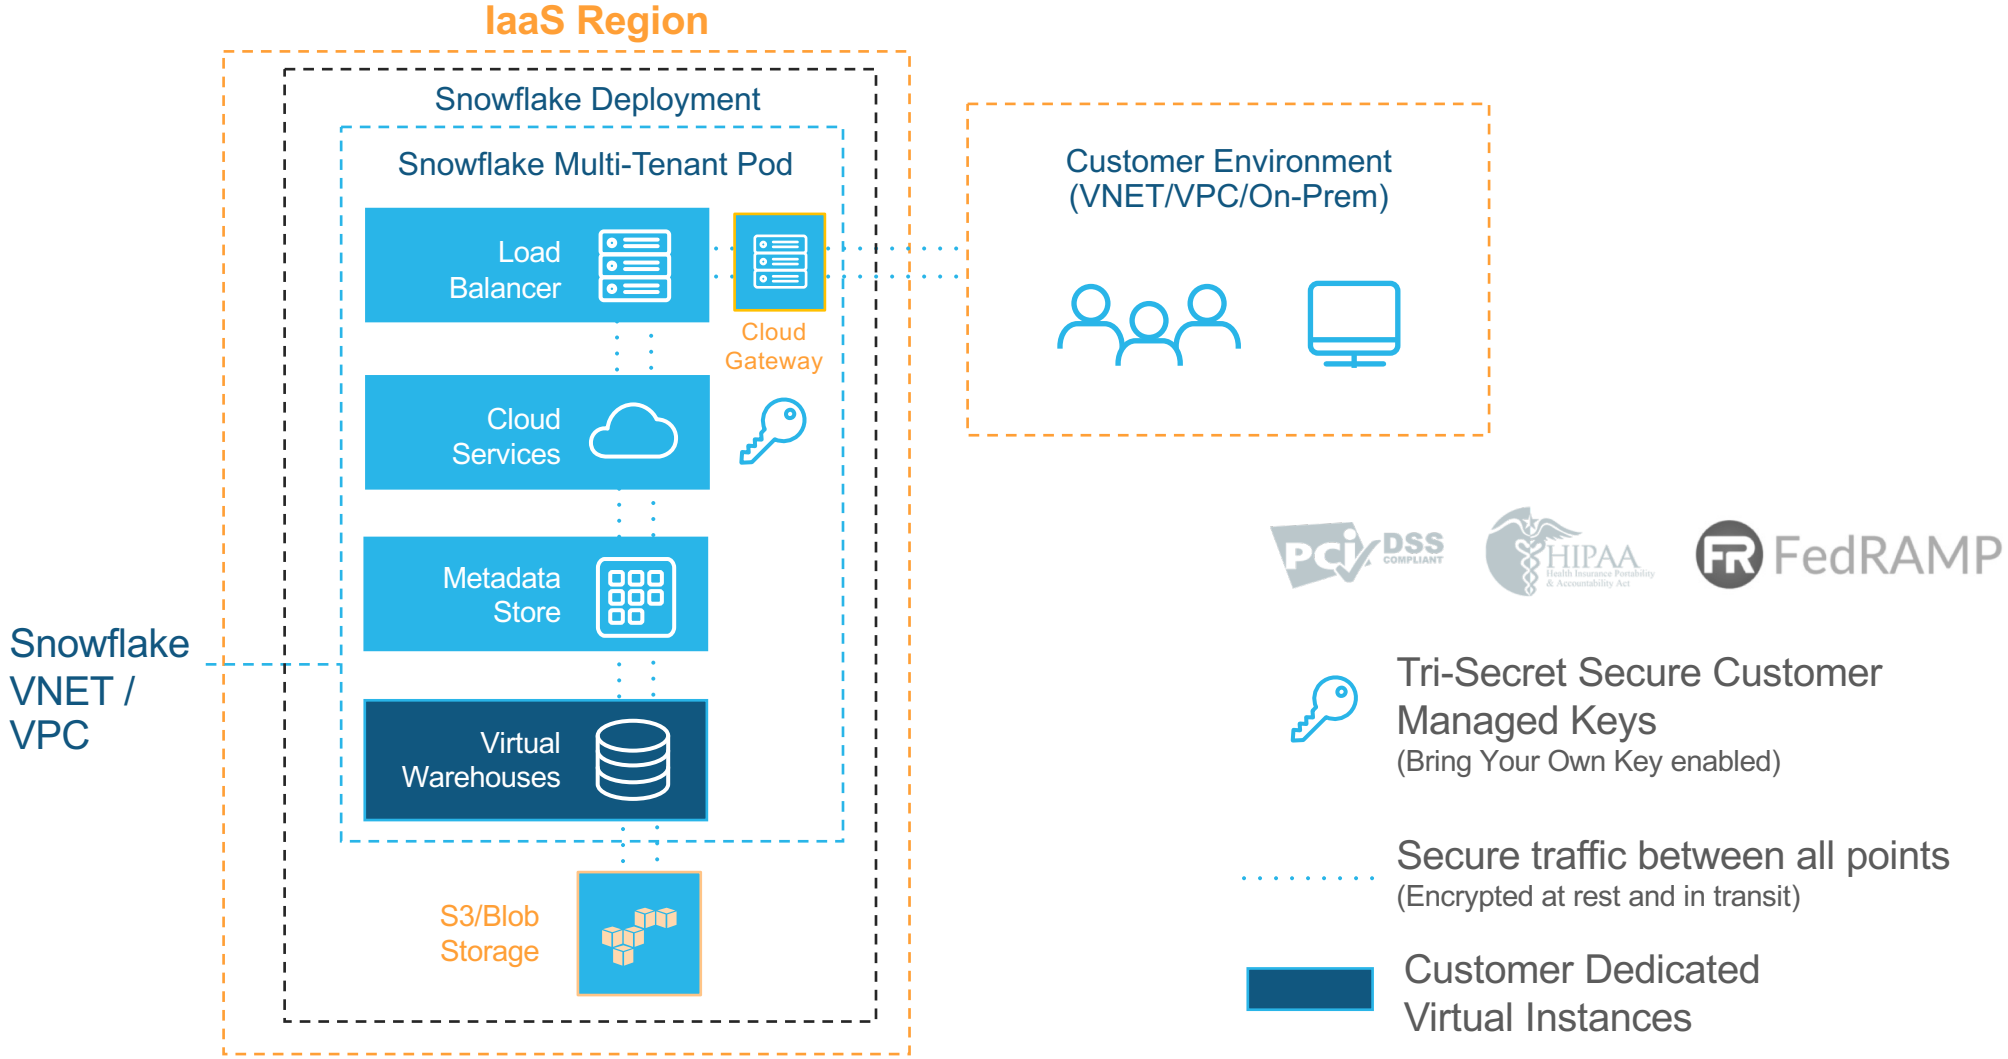

ACCESS – SECURE COMMUNICATION

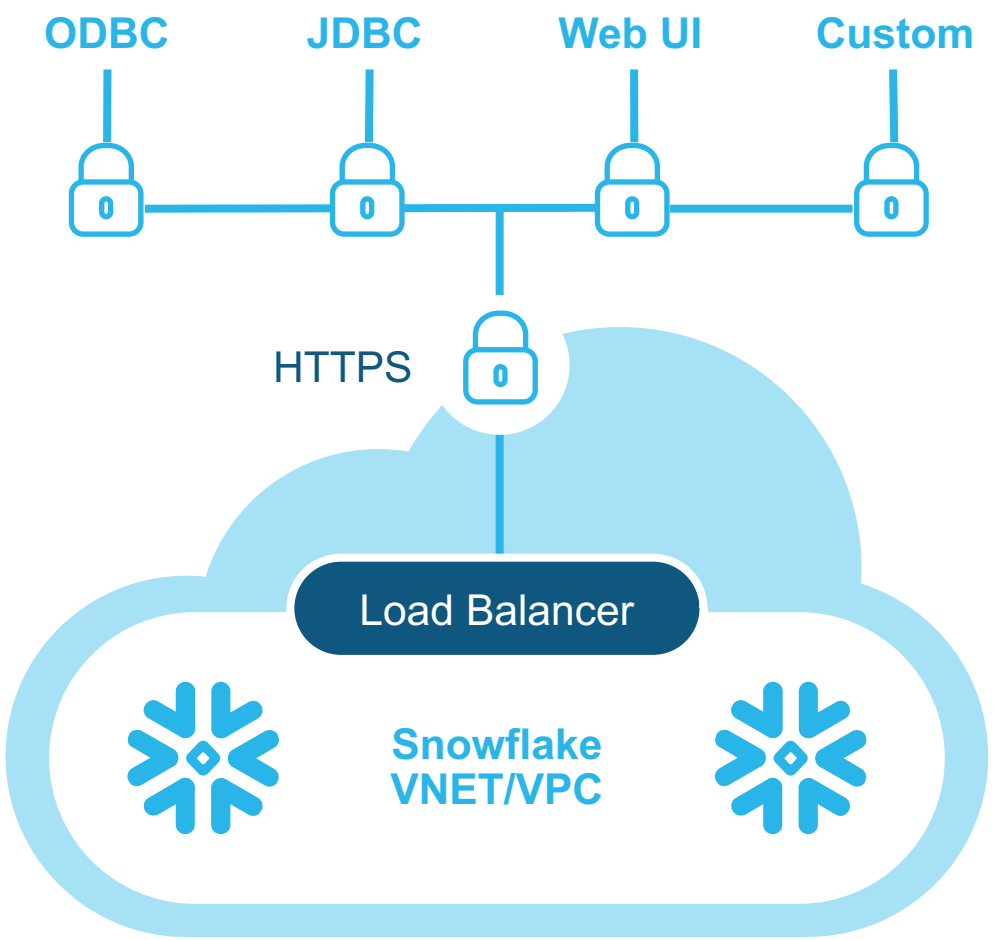

All communication encrypted end-to-end

- Web UI, command line client, and drivers communicate solely over HTTPS
- Connections encrypted using TLS 1.2 from client through to Snowflake Service
- Data encrypted at rest

All access controlled

- IP whitelisting available to restrict client communication to specified IP addresses
- Authentication required for all connections

Customer-Configured Network Policy

APPLICATION SECURITY

Authorization Control

- Role-based authorization
- Authorization for access to all database objects—databases, schemas, tables…
- Authorization for operations in Snowflake—create, stop & start virtual warehouses
- DAC and RBAC info

Application Auditing

- All actions are logged
- Audit Logs available through Snowflake Service

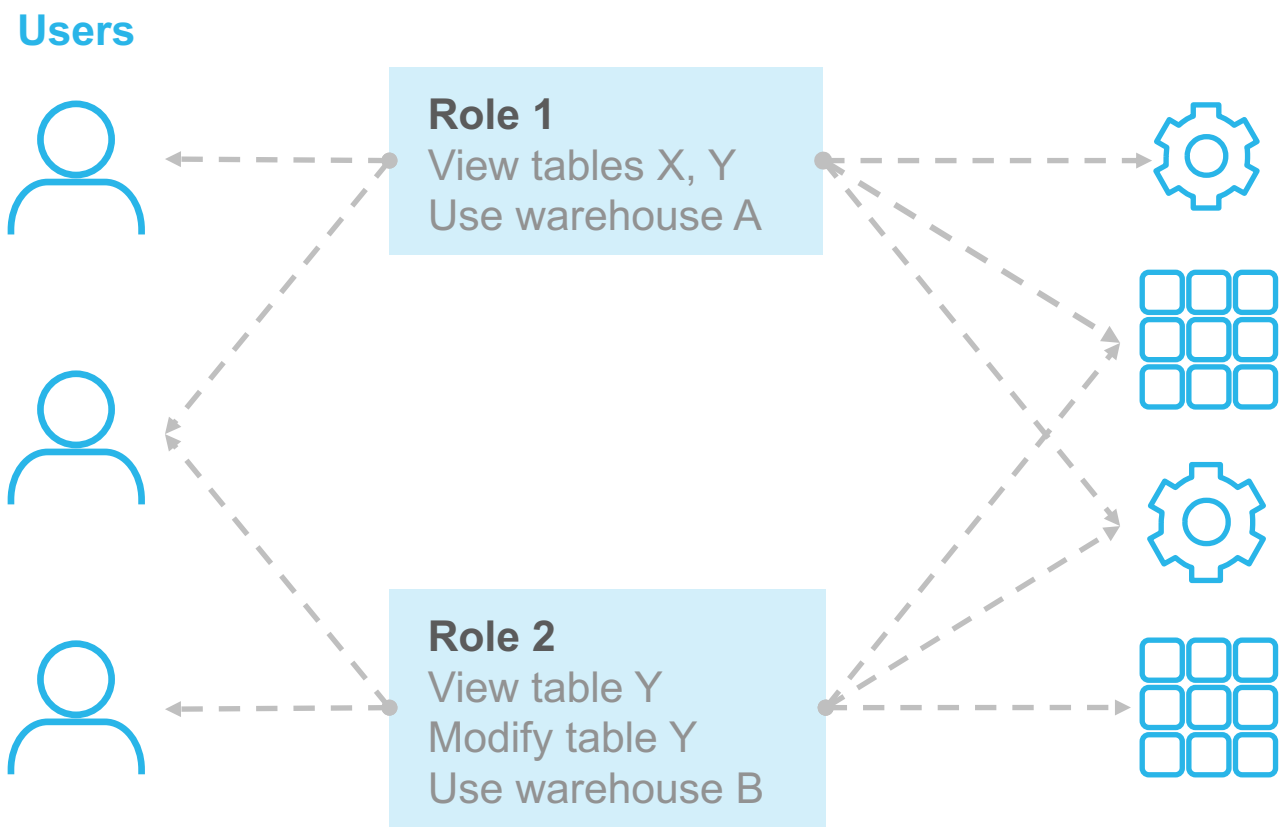

HIERARCHICAL ENCRYPTION FOR DATA AT REST

Hierarchical Key Model using Tri-Secret Secure

- Data is encrypted by default,
  - *no configuration required*
- Hierarchical key model rooted in HSM

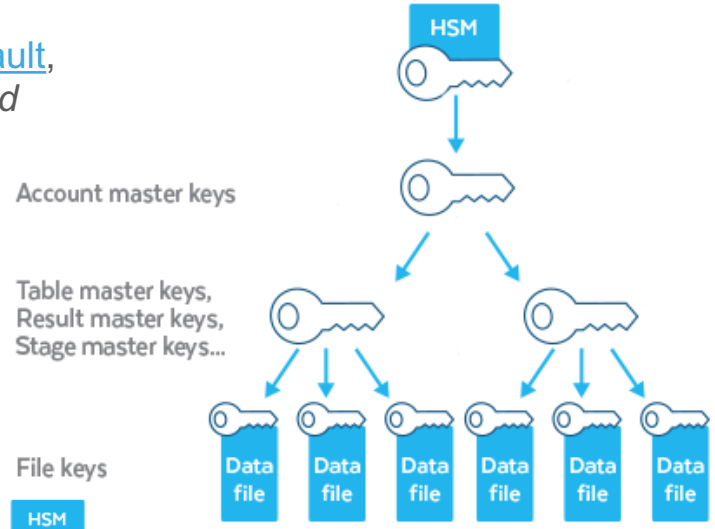

- Tri-Secret Secure

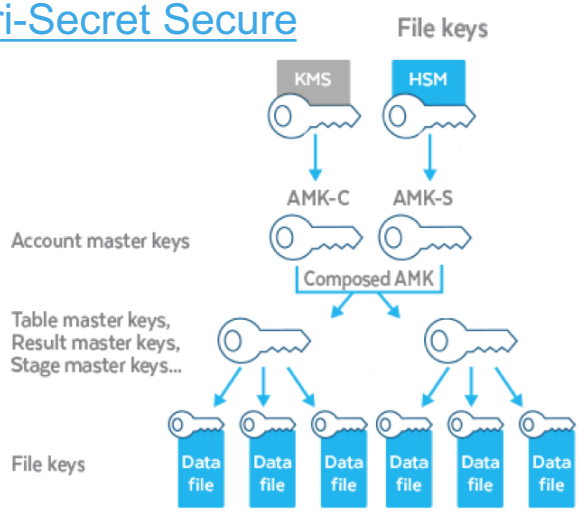

Key Rotation & Re-Keying

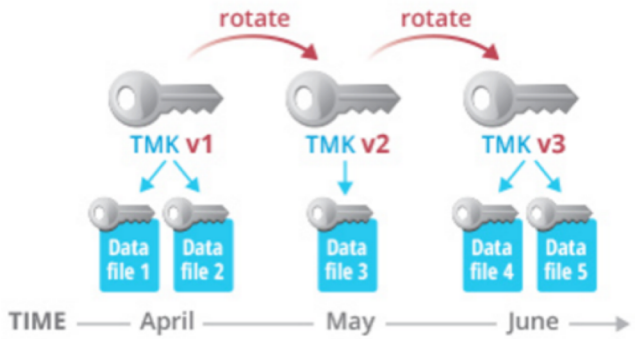

- Keys rotated programmatically by Snowflake
- All data re-encrypted with annual re-keying
- Key Rotation transparent to customer and queries

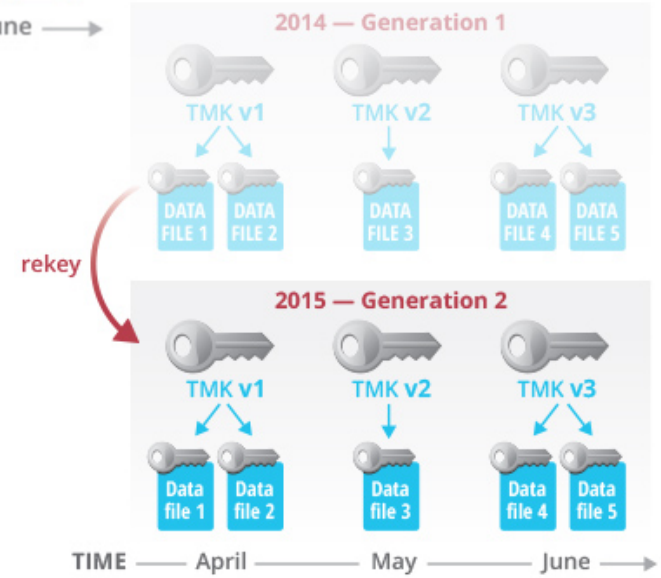

More resources on Key Management

ACCOUNT LOGGING & MONITORING

Customers may audit users, access, and query activity related to their data.

- LOGIN_HISTORY
- QUERY_HISTORY
- Availability for up to 7 days in functions; 365 days in views in the past
- Export through JDBC or as JSON for use in SIEM

Column Name	Data Type	Description
QUERY_ID	TEXT	The statement's unique id.
QUERY_TEXT	TEXT	Text of the SQL statement.
DATABASE_NAME	TEXT	Database that was in use at the time of the query.
SCHEMA_NAME	TEXT	Schema that was in use at the time of the query.
QUERY_TYPE	TEXT	DML, query, etc. If the query is currently running, or the query failed, then the query type may be UNKNOWN.
SESSION_ID	NUMBER	Session that executed the statement.
USER_NAME	TEXT	User who issued the query.
ROLE_NAME	TEXT	Role that was active in the session at the time of the query.
WAREHOUSE_NAME	TEXT	Warehouse that the query executed on, if any.
WAREHOUSE_SIZE	TEXT	Size of the warehouse when this statement executed.
WAREHOUSE_TYPE	TEXT	Type of the warehouse when this statement executed.
CLUSTER_NUMBER	NUMBER	The cluster (in a multi-cluster warehouse) that this statement executed on.
QUERY_TAG	TEXT	Query tag set for this statement through the QUERY_TAG session parameter.
EXECUTION_STATUS	TEXT	Executio
ERROR_CODE	NUMBER	Error cod
ERROR_MESSAGE	TEXT	Error me
START_TIME	TIMESTAMP_LTZ	Statemer
END_TIME	TIMESTAMP_LTZ	Statemer
TOTAL_ELAPSED_TIME	NUMBER	Elapsed
BYTES_SCANNED	NUMBER	Number
ROWS_PRODUCED	NUMBER	Number
COMPILATION_TIME	NUMBER	Compilat
EXECUTION_TIME	NUMBER	Executio
QUEUED_PROVISIONING_TIME	NUMBER	Time (in resize.
QUEUED_REPAIR_TIME	NUMBER	Time (in
QUEUED_OVERLOAD_TIME	NUMBER	Time (in
TRANSACTION_BLOCKED_TIME	NUMBER	Time (in
OUTBOUND_DATA_TRANSFER_CLOUD	TEXT	Target cl
OUTBOUND_DATA_TRANSFER_REGION	TEXT	Target re
OUTBOUND_DATA_TRANSFER_BYTES	NUMBER	Number of bytes transferred in statements that unload data to another region and/or cloud.
INBOUND_DATA_TRANSFER_CLOUD	TEXT	Source cloud provider for statements that load data from another region and/or cloud.
INBOUND_DATA_TRANSFER_REGION	TEXT	Source region for statements that load data from another region and/or cloud.
INBOUND_DATA_TRANSFER_BYTES	NUMBER	Number of bytes transferred in statements that load data from another region and/or cloud.

Column Name	Data Type	Description
EVENT_TIMESTAMP	TIMESTAMP_LTZ	Time of the event occurrence.
EVENT_ID	NUMBER	Event's unique id.
EVENT_TYPE	TEXT	Event type, such as LOGIN for authentication events.
USER_NAME	TEXT	User associated with this event.
CLIENT_IP	TEXT	IP address where the request originated from.
REPORTED_CLIENT_TYPE	TEXT	Reported type of the client software, such as JDBC_DRIVER, ODBC_DRIVER, etc. This information is not authenticated.
REPORTED_CLIENT_VERSION	TEXT	Reported version of the client software. This information is not authenticated.
FIRST_AUTHENTICATION_FACTOR	TEXT	Method used to authenticate the user (the first factor, if using multi factor authentication).
SECOND_AUTHENTICATION_FACTOR	TEXT	The second factor, if using multi factor authentication, or NULL otherwise.
IS_SUCCESS	TEXT	Whether the user's request was successful or not.
ERROR_CODE	NUMBER	Error code, if the request was not successful.
ERROR_MESSAGE	TEXT	Error message returned to the user, if the request was not successful.
RELATED_EVENT_ID	NUMBER	Reserved for future use.

https://www.snowflake.com/use-cases/monitoring-security-analytics/

INFRASTRUCTURE LOGGING & MONITORING

Snowflake uses advanced threat detection tools to monitor production infrastructure

- Failed logins
- File integrity monitoring
- Unauthorized system modifications

Snowflake also uses behavioral monitoring tools to monitor a baseline of production infrastructure behavior

- Network traffic
- User activity
- Binaries

Additional Testing

- 7-10 Penetration Tests / year
- Weekly Vulnerability Scanning
- Automated Web Application Scanning

TIME TRAVEL & FAIL SAFE

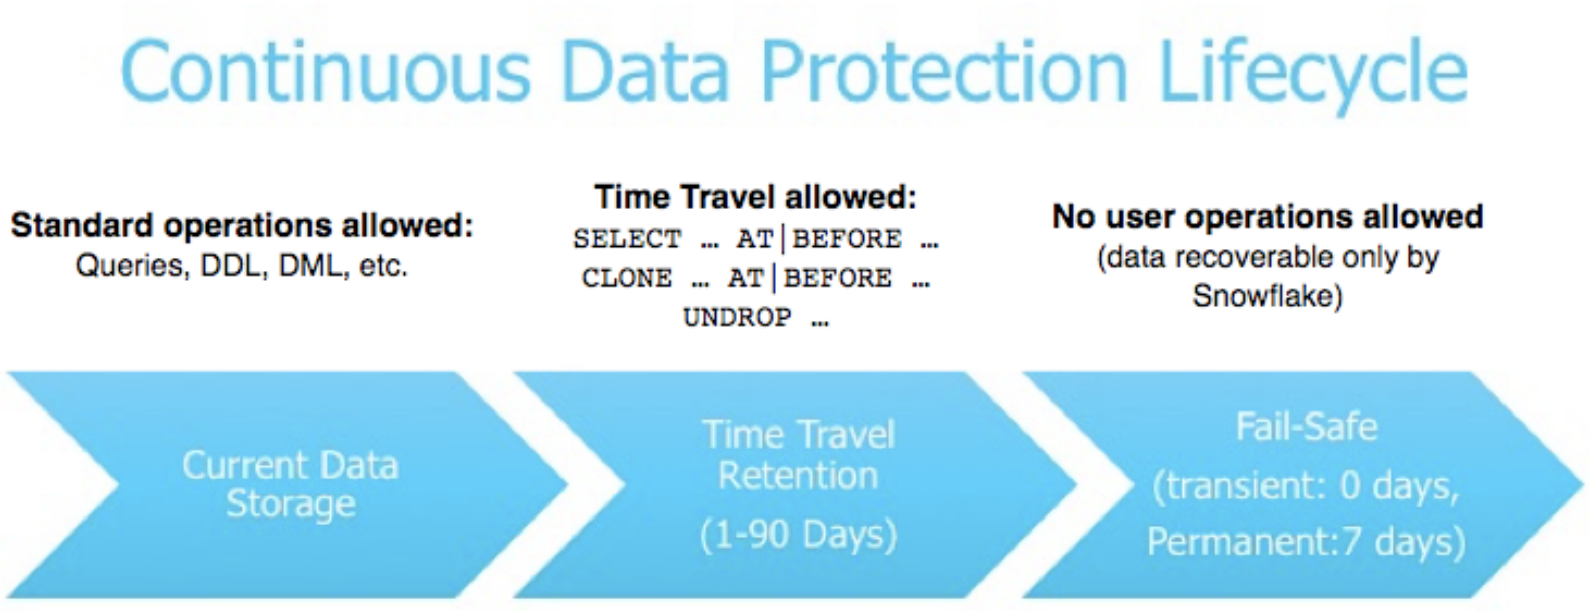

Time Travel

- Select from data as it existed in the past, e.g. before some specific event
- Up to 90 days

Fail Safe

- Request recovery of lost data
- Up to 7 days for most objects

More on Time Travel & Fail Safe

SECURE SDLC

Snowflake has developed a Secure Software Development Lifecycle (SDLC)

- Follows the Microsoft threat modeling SDLC and incorporates elements of OWASP
- All Snowflake development personnel are trained on Snowflake Secure SDLC
- All changes must be documented
- All changes are reviewed for potential security impact
- Major changes require third party penetration test prior to deployment
- All Critical, High, or Medium Findings must be remediated prior to production deployment

Separation of Duties

- Changes must be reviewed and approved prior to deployment
- Developers are prevented from deploying
- Changes are deployed by our Operational team programmatically using Ansible

COMPLIANCE

Third Party Attestations and Certifications

SOC 1 Type II
6 month Coverage Period
SOC 2 Type II
12 Month Coverage Period

PCI-DSS

ISO/IEC 27001

FedRAMP
(Available from OMB/MAX)

HIPAA
(HITRUST in Progress)

Self-Assessments
CAIQ, SIG, Pen Test Results

GDPR – GENERAL DATA PROTECTION REGULATION

What it is

- GDPR is a new EU regulation that becomes effective on May 25, 2018
- Governs the protection and processing of EU personal data

What it means in the context of Snowflake

Different requirements apply to different types of entities

- Controller – Snowflake Customers are responsible for complying with GDPR independently from Snowflake
- Processor – Snowflake is responsible for the following:
    - Putting data processing addendums in place with our customers and our vendors
    - Only using our customers' EU personal data to provide our service to them
    - Being transparent about how we handle and process our customers' EU personal data on their behalf and keeping accurate records
    - Securing customers' EU personal data in our service
    - Facilitating our customers' compliance with data subject requests
    - Notifying customers about changes to our list of subcontractors

Snowflake responsibilities are documented in a Data Processing Addendum (DPA)

Available for signature now

ADDITIONAL COLLATERAL

[Snowflake Security Product Documentation]

The above link provides information on how to configure:

- Network Policies
- MFA, IP Whitelisting
- Federated Authentication / SSO
- Access Control (DAC, RBAC)
- Best Practices
- Audit Logs

THANK YOU