LEGAL PACK · VERSION 1.0

Privacy Policy and Data Security Statement

How SCLOE (Pvt) Ltd, trading as Navia, collects, uses, stores and protects Personal Data through Navia, together with the Scloe Suite Data Security Statement that describes the controls applied across all five applications.

VERSION	EFFECTIVE DATE	LAST UPDATED	GOVERNING LAW
1.0	1 September 2026	1 September 2026	Sri Lanka

TWO TOPICS IN THIS PACK

Topic	What it is	Applies to
1. Privacy Policy	How SCLOE collects, uses, shares and retains Personal Data in Navia, including AI processing and data-subject rights.	Navia platform and Personal Data processed through it
2. Data Security Statement	How SCLOE protects Customer Data. It is informational and does not extend obligations beyond the Privacy Policy and the Data Processing Agreement.	All five Scloe Suite applications: My Quote, Navia Tracking, AI Assistant, AIDOC+ and Proq

PLEASE READ

This pack explains how Personal Data is handled when you use Navia. It is issued with the Data Security Statement so that the privacy rules and the security controls can be read together. Where SCLOE processes Personal Data on a Customer's behalf, the Navia Data Processing Agreement also applies.

CONTENTS

How this pack is organised

TOPIC 1

Privacy Policy

This Privacy Policy explains how **SCLOE (Pvt) Ltd** (“SCLOE”, “we”, “us” or “our”) collects, uses, stores, protects and processes Personal Data through Navia.

VERSION	EFFECTIVE DATE	LAST UPDATED	HOSTING
1.0	1 September 2026	1 September 2026	Hetzner dedicated

01 Our commitment

SCLOE is committed to responsible, transparent and secure processing of Personal Data. Navia is designed for business and logistics use and may process information relating to customers, employees, suppliers, service providers, shipment contacts and other individuals involved in supply-chain operations.

SCLOE aims to process Personal Data lawfully, fairly and transparently and in accordance with applicable data-protection requirements.

02 Information we collect

We may collect the categories of information described below.

2.1 Account information

Name; business email; telephone number; job title; company; username; and other account information.

2.2 Shipment information

Container numbers; booking numbers; bill of lading numbers; vessel information; port information; shipment references; shipment milestones; delivery information; and tracking information.

2.3 Documents and communications

Users may upload documents, comments, notes, instructions and other information.

2.4 Technical information

IP address; device information; browser information; login information; access logs; security information; and usage information.

03 Information from third parties

Navia may receive information from shipping lines, carriers, airlines, freight forwarders, ports, terminals, tracking sources, APIs, logistics providers and Service Providers.

Carrier and tracking information is obtained from each carrier, shipping line, airline, terminal or tracking source the Customer uses. Navia does not rely on a single named tracking provider.

04 How we use Personal Data

We may use Personal Data to:

provide Navia and manage accounts;

track shipments and provide notifications;

facilitate collaboration and document exchange;

provide support and maintain security;

detect fraud and misuse;

improve Navia and develop new features; and

meet legal obligations.

05 AI and automated processing

Navia may use artificial intelligence and machine-learning technologies. AI may be used for shipment predictions, ETA analysis, exception detection, document processing, information extraction, classification, search, recommendations, summarisation, operational insights, customer support and other Navia functionality. AI processing may involve Personal Data where necessary to provide the relevant functionality.

5.1 Human oversight

AI-generated information is intended to support users and does not necessarily represent a final or authoritative decision. Where legally required, individuals may have rights relating to decisions based solely on automated processing. Navia should not be used as the sole basis for legally or similarly significant decisions about individuals unless the applicable legal requirements and safeguards have been satisfied.

5.2 AI accuracy

AI outputs can contain errors. Users should verify important information before relying on AI-generated information for material decisions.

5.3 AI model training

SCLOE does not use identifiable Customer Data to train general-purpose AI models unless the Customer has expressly authorised such use. Customer Data may nevertheless be processed by AI systems where necessary to provide an AI Feature requested or enabled by the Customer. SCLOE may use appropriately aggregated or anonymised information for analytics, product development and service improvement where individuals and Customers are not reasonably identifiable.

5.4 NVIDIA-hosted GPU

Inference runs on NVIDIA-hosted GPU instances allocated to SCLOE, not through a multi-model router. Requests are made server-side only. An AI credential is never sent to the browser, and no AI call is made directly from a User's device. Where that hosting processes Personal Data on SCLOE's behalf, SCLOE applies contractual and data-protection safeguards, and the provider is treated as a sub-processor under the Navia Data Processing Agreement.

5.5 AI-generated information

AI-generated information should be treated as decision-support information unless expressly identified as authoritative information obtained from a verified source.

06 Customer Data

Where a Customer uses Navia to process information concerning its customers, employees, suppliers or other individuals, SCLOE may process such information on behalf of the Customer. The Customer generally determines what Personal Data is submitted, why it is processed, who receives access, and how long it should be retained.

The Customer may therefore act as Controller while SCLOE acts as Processor for such processing. These responsibilities are further described in the Navia Data Processing Agreement.

07 Sharing of Personal Data

Personal Data may be shared with:

hosting providers (Hetzner dedicated servers operated by SCLOE);

AI infrastructure (NVIDIA-hosted GPU instances allocated to SCLOE);

transactional-email and map providers listed in Schedule A;

carriers, freight forwarders and Service Providers the Customer uses or invites;

professional advisers; and

regulators or authorities where legally required.

SCLOE does not sell identifiable Customer Personal Data for unrelated commercial purposes.

08 Customer-controlled sharing

Navia may allow a Customer to invite a Service Provider to access shipment information. The Customer determines what information is shared with the Service Provider and is responsible for ensuring that the Service Provider is appropriately authorised.

09 Document processing

Documents uploaded to Navia may contain Personal Data or confidential business information. Navia may use automated or AI-assisted technologies to extract, classify, search, summarise or process information from such documents. The Customer remains responsible for ensuring that it has the necessary authority to upload and process the documents.

10 Data security

SCLOE uses reasonable technical and organisational measures designed to protect Personal Data. These may include access controls, authentication, role-based permissions, encryption where appropriate, monitoring, audit logging, backup procedures and other security controls. The Data Security Statement in Topic 2 describes those controls as implemented at the effective date.

No internet-based service can guarantee absolute security.

11 International processing

Navia is hosted on Hetzner dedicated servers operated by SCLOE, and uses NVIDIA-hosted GPU instances and other technology providers. Those providers may be located outside Sri Lanka. Where Personal Data is transferred across jurisdictions, SCLOE will implement appropriate safeguards where required by applicable law.

12 Data retention

SCLOE retains Personal Data only for as long as reasonably necessary for:

service delivery;

legal compliance;

security;

dispute resolution;

contract enforcement; and

legitimate operational requirements.

Customer-specific retention requirements may be established through the Data Processing Agreement or a commercial agreement.

13 Data-subject rights

Subject to applicable law, individuals may have rights to:

access Personal Data;

request correction;

request deletion;

object to processing;

request restriction; and

request review of certain automated decisions.

Where SCLOE processes Personal Data on behalf of a Customer, requests may need to be directed to that Customer. SCLOE will provide reasonable assistance where required.

14 Cookies

Navia may use cookies and similar technologies for authentication, security, session management, analytics, performance and platform functionality.

15 Children

Navia is intended for business and professional use and is not directed toward children. SCLOE does not knowingly seek to collect children's Personal Data through Navia.

16 Changes to this policy

SCLOE may update this Privacy Policy from time to time. Material changes may be communicated through Navia or by email. The version and effective date are stated at the front of this pack.

17 Contact

Privacy enquiries may be sent to **info@scloe.com** or raised through www.scloe.com. SCLOE (Pvt) Ltd, Sri Lanka. Hotline +94 77 898 9155.

SCHEDULE A

Service providers

The third parties Navia relies on for the functions named below. Where a provider processes Personal Data on SCLOE's behalf it is a sub-processor under the Navia Data Processing Agreement. SCLOE may change a provider and will notify the Customer of a change affecting personal-data processing.

Function	Provider	Data reaching the provider
Platform hosting	Hetzner dedicated servers, operated by SCLOE	All Customer Data and documents
Carrier tracking data	Every carrier, shipping line, airline, terminal and tracking source the Customer uses — not a single named provider	Container, booking and bill-of-lading references needed to retrieve events
Transactional email	SendGrid (primary); Resend (fallback)	Recipient name and email, message content
AI model access	NVIDIA-hosted GPU instances allocated to SCLOE	Prompt content: shipment records and document text submitted to an AI Feature
Map tiles	MapTiler	Map viewport requests only; no Customer Data
Single sign-on	Scloe identity service	User identifier and email, sign-in events

Carriers, shipping lines, airlines, terminals, forwarders and other Service Providers the Customer invites to a shipment are not sub-processors of SCLOE. They receive information because the Customer has chosen to share it, and act on their own account.

TOPIC 2

Scloe Suite Data Security Statement

How SCLOE (Pvt) Ltd protects Customer Data across all five applications in the Scloe suite (My Quote, Navia Tracking, AI Assistant, AIDOC+ and Proq), together with the single sign-on, staff console and integrations they share. This statement accompanies the Privacy Policy and describes controls as implemented at the effective date.

Applications	Data location	Backup cadence	Model training
Five, one sign-in	Hetzner dedicated servers	Nightly, verified	Not without consent

01 Scope: every application, and what each holds

Application	Customer Data it holds
My Quote	Quotation requests, provider responses, rates, comparison working and the contact details of the people involved
Navia Tracking	Shipments, containers and air waybills, carrier events, shipping documents, free-time and demurrage contract terms
AI Assistant	Sailing-schedule queries, spreadsheets the Customer uploads and their contents, and conversation history
AIDOC+	Uploaded shipping, customs and invoice documents, and the field values read out of them
Proq	Procurement requests, approvals, supplier records, prices and purchase orders
Shared services	Scloe single sign-on identities, the SCLOE staff console, the upstream tracking interface, notification mail, and public link surfaces such as share links and invitations

Internal engineering and quality-assurance utilities operated by SCLOE are in scope for the personnel and access controls in clause 17, and hold no production Customer Data other than transiently during authorised support work.

02 One sign-in, separated data

A single Scloe identity opens every application the Customer is entitled to. Shared identity is not shared authority: entitlement is evaluated per application, and a User who may quote is not thereby a User who may release a purchase order. Signing out, revoking a session or disabling an identity takes effect across the Suite.

Data is separated by Workspace in every application. One company’s quotations, shipments, documents, spreadsheets and purchase orders are never returned to another company’s Workspace, and cross-application features operate only within one Workspace at a time.

03 Where data lives

The Suite runs on Hetzner dedicated servers operated by SCLOE. Traffic reaches each application through a reverse proxy that terminates TLS on the host; the applications themselves are not exposed directly to the internet. Databases, document stores and nightly backups are held on those same dedicated servers, under SCLOE control.

Data class	Held in	Protection at rest
Business records	Application databases, workspace-scoped rows	Host and volume access control; file permissions limited to the service account

Data class	Held in	Protection at rest
Uploaded documents and spreadsheets	Object store, keyed per workspace	Signed, expiring access; content digest recorded on write
Passwords	Never stored in clear	scrypt, per-user salt, parameters stored with the verifier
Session tokens, workspace keys, link tokens	Only the digest is stored	SHA-256; the raw value exists only in the holder's browser or in one email
Carrier, supplier and integration credentials	Sealed records in the database	AES-256-GCM with a key derived from a deployment secret held outside the database
AI provider API keys	Server-side configuration only	Never sent to the browser; restricted file permissions; rotatable without redeployment
Nightly backups	Restricted backup directory on SCLOE storage	Directory and file permissions limited to the operating account; rotated on a fixed window

Data in transit between the User's browser and the Services, and between the Services and their providers, is encrypted using TLS.

04 Identity and authentication

The Suite recognises three kinds of principal, each with its own credential store: **members** (people in a customer workspace), **workspace keys** (shared devices and integrations, authenticated as the Customer rather than as a person), and **administrators** (SCLOE staff, in a separate session store from customer sessions).

Passwords are hashed with scrypt using a per-user salt and compared in constant time; the hashing work is performed even when no matching account exists, so response timing does not reveal whether an address is registered. Passwords are checked against a list of known-breached values on set. Where a password is issued by SCLOE, the User is required to change it at first sign-in.

Sign-in attempts are rate-limited per address and per source network, with lockout after repeated failure. Time-based one-time-password authentication is available where enabled for the Workspace, and single sign-on uses a signed handoff from Scloe with the authorisation-code flow and PKCE where OIDC is in use.

SESSIONS

Session tokens are delivered in cookies marked HttpOnly and SameSite, and served as Secure in production, so they are not readable by page scripts and are not sent on cross-site requests. Only the SHA-256 digest of a token is stored. Revocation is enforced at lookup: a revoked session stops working immediately, and an administrator can revoke a single session or every session belonging to a User. Changing an administrator's password ends that administrator's existing sessions.

05 Authorisation and tenant isolation

Every request that reads or writes Customer Data resolves its caller to exactly one account before any query runs, and queries are scoped to that account. Write authority is checked separately from read authority, so a read-only User cannot alter records. Where an upstream provider issues SCLOE a single credential covering more than one customer, that provider has no concept of workspaces: the Service performs the separation itself by matching each record to the workspace that owns it and returning only owned records, cached reads included.

Automated checks run as part of the release process assert that workspace-scoped endpoints refuse cross-workspace reads. A finding in that class is treated as a security defect, not a functional one.

06 Application-layer protections

Cross-site request forgery. Every state-changing request is checked for same-origin provenance. Absent headers are treated as refusal, not permission.

Rate limiting. Sliding-window limits are applied to sign-in, password reset, verification resend, feedback, public link access and other unauthenticated surfaces. Throttled requests answer in a way that does not confirm whether a record exists.

Enumeration resistance. Sign-up, reset and verification responses are deliberately identical whether or not the address is known.

Bounded input. Uploads and imports are size-capped and parsed against an asserted structure rather than trusted.

Link tokens. Share links, invitations, verification and reset links carry high-entropy tokens, expire, and are stored only as digests. A reset link is single-use.

Inbound webhooks. Payloads from integrated providers are verified by HMAC signature before being applied, and each delivery is recorded with its digest.

07 Documents, spreadsheets and uploads

Each uploaded file is recorded with its size and SHA-256 digest, and its declared file type is compared against the type detected from its content. Files carry a scan state, and a file that fails scanning is withheld rather than served. Downloads are issued through short-lived, authorised links rather than public paths. Identical re-uploads are recognised by digest, so a document is not silently duplicated. This applies equally to shipping documents in Navia Tracking, invoices in AIDOC+, spreadsheets in AI Assistant and attachments in My Quote and Proq.

08 AI models: what is sent, and what is not

Every application in the Suite uses AI to some degree. Inference runs on NVIDIA-hosted GPU instances allocated to SCLOE. The models in use are recorded in the platform configuration and will be named on request. Requests are made server-side only. An AI credential is never sent to the browser, and no AI call is made directly from a User's device.

Application	Sent to a model	Purpose
My Quote	Quotation and rate lines being compared	Normalising and comparing responses
Navia Tracking	Shipment records and document text for the request in hand	Delay prediction, exception analysis, insights, natural-language search
AI Assistant	The question asked, and the rows of an uploaded spreadsheet needed to answer it	Schedule lookup and answering from the Customer's own data
AIDOC+	The page content of the document being read	Classification and field extraction
Proq	Requisition and supplier-response text	Drafting, matching and summarising a procurement record

What is never sent. Passwords, session tokens, workspace keys, API credentials and whole databases are never included in a prompt. A request carries the record or document in hand, not the Workspace.

Training. SCLOE does not use identifiable Customer Data to train a general-purpose AI model and does not make it available for general model training, in any application in the Suite, without the Customer's prior written authorisation. Aggregated or anonymised information that does not reasonably identify the Customer or an individual may be used for analytics and product development.

Human review. AI output across the Suite is decision support, not an instruction or an authorisation. A predicted arrival, a compared rate, an extracted invoice value and a drafted purchase order are all reviewable by a person before they are acted on, and the Customer remains responsible for that review. Unless expressly agreed in writing, the

Services are not intended to make legally or similarly significant decisions about individuals solely by automated means.

09 Data leaving the platform

Customer Data leaves SCLOE infrastructure only for the functions listed in Schedule A of the Privacy Policy: carrier and supplier lookups, transactional email, AI model access, map tiles, vessel and aircraft position feeds, and single sign-on. Map tile, vessel-position and flight-position requests carry no Customer Data, only a location viewport, a vessel identifier or a flight callsign. Carriers, terminals, forwarders, suppliers and other Service Providers the Customer invites are not sub-processors of SCLOE: they receive information because the Customer chose to share it, and act on their own account.

10 Logging, audit and monitoring

Administrative and security-relevant actions are recorded to an audit trail that identifies the actor, the action, the affected record and the time. Live sessions are visible to SCLOE administrators, and the platform reports authentication failures, job failures, integration health, dead-lettered events and integrity findings on operational dashboards.

Where an IP address is retained for abuse control, it is stored as a salted digest and truncated to its network rather than kept in full, except on the session screen where the full address is needed to answer “was this me”.

Support sessions. Where SCLOE staff need to enter a customer workspace to resolve a fault, a reason is required and kept on the Customer’s record, a banner naming the Customer stays visible while the session is active, the grant expires by itself, and the session is recorded with the staff member, the workspace and both timestamps. Staff credentials are separate from customer credentials and are never shared.

11 Backup and recovery

A backup of each primary database is taken nightly using the database engine’s online backup mechanism, so the copy is internally consistent rather than a file snapshot taken mid-write. Each backup is compressed and then verified before it is kept: the copy is re-opened, an integrity check is run, and a witness count is taken across core tables. A backup that does not verify is deleted rather than retained, and every run, successful or failed, is recorded with its outcome, size and duration.

Cadence	Recovery point	Retained copies
Nightly, on a fixed schedule	Up to 24 hours of data loss in a total-loss scenario	A rolling fourteen-night window, then rotated out

These figures are the platform’s engineering targets. They are not a contractual recovery-time commitment unless an executed service-level appendix states one.

12 Retention, legal holds and deletion

Records are retained under the retention schedule applying to the Workspace. When a period elapses, records enter a deletion queue and deletion is carried out against that queue, so what was removed and when is recorded rather than inferred. A legal hold suspends scheduled deletion for the records it covers; a hold is placed and released only by the party holding the legal reason for it, and the release is recorded with its reason. Post-termination export and retention are described in the Privacy Policy (retention) and any applicable Data Processing Agreement.

13 Encryption and key management

Three different cryptographic jobs are done with three different tools, deliberately. **Passwords** are stretched with scrypt, whose cost parameters are stored alongside each verifier so they can be raised later without invalidating

existing credentials. **Tokens** (sessions, workspace keys, share links, reset and verification links) are stored as SHA-256 digests, which are one-way and fast because the token is already high-entropy and cannot be guessed. **Stored third-party credentials** are encrypted with AES-256-GCM, an authenticated cipher: a record altered by hand fails to decrypt rather than decrypting to something plausible.

The encryption key is derived from a deployment secret that lives outside the database, in a configuration file owned by the service account with restricted permissions. It is never committed to source control and never reaches the browser. Loss of that secret makes sealed credentials unrecoverable, which is the intended property and the reason they are re-enterable rather than merely re-decryptable.

Rotation. Session-signing secrets, console secrets, workspace keys and provider API keys can each be rotated independently and without redeployment of the applications. Rotation has a known, documented effect. Rotating a session-signing secret signs every User out; rotating a workspace key requires the integration holding it to be re-issued. It is therefore performed deliberately rather than on a blind schedule. A workspace key is displayed once at issue and stored only as a digest with a short non-secret prefix for identification; it cannot be recovered, only replaced.

14 Infrastructure and network security

The applications are not exposed directly to the internet. A reverse proxy terminates TLS on the host and is the only public listener; each application runs behind it under a dedicated unprivileged service account, supervised by the host service manager and restarted automatically on failure.

Supporting components are bound to the loopback interface rather than published: the databases are files readable only by their service account, the response cache is password-protected and local, and the scheduled job runner accepts calls only over loopback and only with a runner token. The staff console additionally answers as though it did not exist until a browser presents a valid unlock path, so it is not discoverable by scanning.

Configuration files carrying secrets are owned by the service account with restricted file modes, and host access is limited to named administrators. The physical infrastructure is Hetzner dedicated servers; Hetzner is responsible for data-centre physical security, and SCLOE is responsible for everything above the host.

15 Secure development, testing and change management

Changes are reviewed before release and must pass an automated verification suite that runs against a live instance, not a mock. The suite covers authentication flows, access control, tenant isolation, public link behaviour, data-layer integrity and backup verification.

Those checks are written to assert refusals, not only successes: a test that would still pass if the guard were removed is treated as no test at all. Where a check cannot yet be met, it reports its own gap rather than reporting green: an unverified control is shown as unverified.

Secrets are not committed to source control, and any data file containing credential material is excluded from it. Schema changes are forward-only and recorded; operations that would destroy customer records are not performed as part of routine change. Dependencies are pinned and updated on review.

16 Vulnerability management and service continuity

SCLOE maintains an internal register of security and integrity findings, each with a severity, an owner and a resolution state, and reviews it as part of ongoing platform work. Findings are also produced automatically: the platform runs its own integrity and cross-tenant checks and raises a finding when one fails. Dependency and platform patches are applied on review, and urgent patches out of cycle.

Planned maintenance is scheduled outside Sri Lanka business hours where practicable. Service continuity rests on the nightly verified backups described in clause 11 together with a documented rebuild path for the application host.

SCLOE states plainly what is not yet in place: the Suite runs in a single region without automatic cross-region failover, and recovery from a total host loss is a rebuild-and-restore operation rather than a switchover. Recovery time is therefore a best-effort target, not a guarantee, unless an executed service-level appendix states one.

17 Incident response, personnel and change

On becoming aware of a security incident affecting Customer Data in any application, SCLOE will contain the incident, preserve relevant logs, assess what data and which workspaces are affected, and remediate the cause. Because one identity opens the whole Suite, a credential compromise is assessed against every application, not only the one where it was noticed. Where a personal-data breach is involved, SCLOE will notify the affected Customer without undue delay and provide the information the Customer needs to meet its own notification duties.

Access to production systems is limited to the SCLOE personnel who require it, under individually issued credentials, and is removed when a role ends. Staff are bound by confidentiality obligations. Administrative capability is permission-gated per action rather than granted wholesale, and administrative actions are audited. Changes are reviewed and pass an automated verification suite before release, including access-control and tenant-isolation checks.

Suspected vulnerabilities and incidents may be reported to info@scloe.com. SCLOE will acknowledge a report and will not pursue a good-faith reporter who does not access, alter or exfiltrate data beyond what is necessary to demonstrate the issue.

18 What the Customer is responsible for

- who is invited into the Workspace, which applications each User may open, and what permissions they hold in each;
- removing Users who have left, which is one removal across all five applications, and reviewing membership periodically;
- the confidentiality of passwords, workspace keys and invitation links, and the physical security of enrolled shared devices;
- the accuracy and lawfulness of what it uploads, and its authority to share each document or spreadsheet;
- reviewing AI output before acting on it, in every application;
- reporting suspected unauthorised access, or any case where a Service appears to show another party’s data, without delay.

19 Security programme

SCLOE reviews the platform’s security posture on a continuing basis and maintains an internal register of findings with owners and priorities. Work currently in the programme includes an off-host copy of nightly backups, a scheduled restore drill, extending backup coverage to the document store, broader enrolment of two-factor authentication, and periodic rotation of provider API keys.

SCLOE does not currently hold a third-party security certification. Where a Customer requires one, or requires a security questionnaire, penetration-test summary or audit right, that may be agreed in the subscription agreement.

20 Status of this statement and contact

This statement describes controls as implemented at the effective date across all five applications, and is reviewed at least annually or sooner on a material change. It is provided for information and does not extend SCLOE’s obligations beyond the Privacy Policy and the Data Processing Agreement. Where it conflicts with any of those, they prevail.

CONTACT

How to reach SCLOE

Privacy and security enquiries about this pack may be sent to the addresses below.

Issued by	Privacy and security enquiries
SCLOE (Pvt) Ltd Sri Lanka	info@scloe.com · www.scloe.com Hotline +94 77 898 9155

END OF LEGAL PACK

TOPIC 1 PRIVACY POLICY · TOPIC 2 DATA SECURITY STATEMENT

MY QUOTE · NAVIA TRACKING · AI ASSISTANT · AIDOC+ · PROQ