



Enterprise Data Protection Policy

Field	Value
Policy Title	Enterprise Data Protection Policy
Document Identifier	OG-DPP-001
Version	1.0
Status	Approved
Effective Date	October 8, 2025
Next Review Date	October 8, 2026
Review Frequency	At least annually; ad hoc upon material regulatory, technological, or operational change
Policy Owner	Oleksii Bondarenko, Chief Security & Compliance Officer (CSCO)
Executive Sponsors	Preston Obadan, Chief Executive Officer; Kevin Kegera, Chief Commercial Officer
Applicability	All personnel, systems, subsidiaries, and third parties processing data on behalf of Orchestrate Global Inc.
Geographic Scope	Global, with mandatory localization for African customer data



Legal Foundations

GDPR, CCPA/CPRA, Nigeria Data Protection Regulation (NDPR),
Rwanda Law No. 058/2021, PCI DSS v4.0, Wyoming Corporate
Code


Contact for Clarification

compliance@orchestrate.global

This policy constitutes a binding directive issued under the authority of the Chief Executive Officer and derives its enforceability from Section 4.1 of the Orchestrate Global Inc. Enterprise Compliance Policy (OG-ECP-001).

Approvals

By signing below, the undersigned confirm that they have reviewed and approved this policy and that Orchestrate Global Inc. commits to implementing and maintaining effective data protection controls aligned with all applicable legal, regulatory, and contractual obligations.

Position	Name	Signature	Date
Chief Executive Officer (CEO)	Preston Obadan		October 8, 2025
Chief Commercial Officer (CCO)	Kevin Kegera		October 8, 2025

Document Owner Acknowledgement

I, Oleksii Bondarenko, in my capacity as Chief Security & Compliance Officer of Orchestrate Global Inc., acknowledge full responsibility for the maintenance, communication, enforcement, and periodic review of this policy, and affirm that it aligns with all applicable legal, regulatory, and contractual obligations across the Company's jurisdictions of operation.

Role	Name	Signature	Date
Chief Security & Compliance Officer (CSCO)	Oleksii Bondarenko		October 8, 2025



Document Control Log

Version	Effective Date	Description of Change	Prepared By	Reviewed By	Approved By
1.0	October 8, 2025	Initial issuance of enterprise-wide data protection policy	CSCO	CEO, CCO	CEO, CCO

Table of Contents

Clause 1. Interpretation and Definitions

1.1 Defined Terms

1.2 Construction and Interpretation Rules

Clause 2. Policy Statement and Foundational Principles

2.1 Recognition of Privacy as a Fundamental Human Right

2.2 Commitment to Lawful, Fair, and Transparent Processing

2.3 Binding Nature and Consequences of Non-Compliance

Clause 3. Purpose and Scope of Application

3.1 Purpose of the Policy

3.2 Covered Legal Entities

3.3 Covered Data Subjects

3.4 Covered Data Types

3.5 Covered Systems, Environments, and Infrastructure

3.6 Covered Personnel and Third Parties

Clause 4. Data Protection Governance and Accountability Framework

4.1 Three Lines of Defense Model

4.2 Role-Based Accountability Matrix

4.3 Incident Escalation and Reporting Protocol

4.4 Board and Executive Oversight Responsibilities

Clause 5. Legal and Regulatory Foundations

5.1 Applicable Statutory and Sectoral Frameworks

5.2 Contractual and Standard-Based Obligations

5.3 Regulatory Obligations Register (ROR)

Clause 6. Data Classification Framework

6.1 Classification Levels and Criteria

6.2 Handling Requirements by Classification Tier

6.3 Misclassification Remediation and Enforcement



Clause 7. Data Inventory, Mapping, and Ownership

7.1 Centralized Data Inventory Requirements

7.2 Data Flow Mapping Standards

7.3 Data Ownership and Stewardship Model

Clause 8. Core Principles of Lawful Data Processing

8.1 Lawfulness, Fairness, and Transparency

8.2 Purpose Limitation

8.3 Data Minimization

8.4 Accuracy and Integrity

8.5 Storage Limitation

8.6 Integrity and Confidentiality

8.7 Accountability

Clause 9. Lawful Bases for Processing

9.1 Consent as a Lawful Basis

9.2 Contractual Necessity

9.3 Legal Obligation

9.4 Vital Interests

9.5 Public Interest and Official Authority

9.6 Legitimate Interests and Balancing Tests

Clause 10. Consent Management Framework

10.1 Standards for Valid Consent

10.2 Consent Capture, Storage, and Versioning

10.3 Withdrawal Mechanisms and Automation

10.4 Special Provisions for Sensitive and Child Data

Clause 11. Data Subject Rights (DSR) Framework

11.1 Right of Access

11.2 Right to Rectification

11.3 Right to Erasure (“Right to Be Forgotten”)

11.4 Right to Restriction of Processing

11.5 Right to Data Portability

11.6 Right to Object

11.7 Rights Related to Automated Decision-Making and Profiling

11.8 Verification, Fulfillment Timelines, and Recordkeeping

Clause 12. Data Accuracy and Integrity Controls

12.1 Self-Service Correction Mechanisms

12.2 Correction Timelines and Audit Trails

12.3 Prevention of Inaccurate Data Propagation

Clause 13. Data Retention and Secure Disposal

13.1 Retention Periods by Data Type

13.2 Secure Disposal Methods



13.3 Legal Hold Suspension Protocol

13.4 Certification of Physical and Cryptographic Destruction

Clause 14. Data Storage Limitations and Localization Requirements

14.1 Geographic Data Residency Rules

14.2 Prohibition on Unauthorized Cross-Region Replication

14.3 Infrastructure Mapping by Jurisdiction

Clause 15. Cross-Border Data Transfers

15.1 Permitted Transfer Mechanisms

15.2 Transfer Impact Assessments (TIAs)

15.3 Supplementary Measures and Encryption Requirements

15.4 Prohibition on Transfers to High-Risk Jurisdictions Without Approval

Clause 16. Third-Party Data Processing and Vendor Management

16.1 Due Diligence Requirements

16.2 Mandatory Contractual Clauses

16.3 Sub-Processor Restrictions

16.4 Continuous Monitoring and Audit Rights

Clause 17. Cryptographic Controls and Key Lifecycle Management

17.1 Encryption Standards for Data at Rest and in Transit

17.2 Key Generation within FIPS 140-2 Level 3 HSMs

17.3 Key Storage, Distribution, and Usage Policies

17.4 Key Rotation Schedules and Emergency Procedures

17.5 Cryptographic Key Destruction and Certification

Clause 18. Access Control and Authentication

18.1 Unique User Identification and Password Complexity

18.2 Multi-Factor Authentication (MFA) Requirements

18.3 Session Timeouts and Inactive Account Management

18.4 Role-Based Access Control (RBAC) and Quarterly Certifications

18.5 Privileged Access Management (PAM) Framework

Clause 19. Logging, Monitoring, and Auditability

19.1 Mandatory Logging of Critical Events

19.2 Immutable Log Storage and Retention

19.3 Cryptographic Hashing and Transparency Logs

19.4 Access Restrictions and Dual Control for Log Modification

Clause 20. Endpoint and Device Security

20.1 Full-Disk Encryption and Screen Lock Policies

20.2 Portable Media Encryption and Wipe Rules

20.3 BYOD Enrollment and Containerization Requirements

20.4 Endpoint Detection and Response (EDR) Deployment



Clause 21. Network and Infrastructure Security

- 21.1 Zero Trust Network Architecture (ZTNA)
- 21.2 Default-Deny Firewall Rules and Microsegmentation
- 21.3 Isolation of Cardholder Data Environment (CDE)
- 21.4 Protocol Restrictions and TLS Enforcement

Clause 22. Application and API Security

- 22.1 Secure Software Development Lifecycle (SDLC)
- 22.2 Threat Modeling and STRIDE Analysis
- 22.3 Automated Code Scanning (SAST/DAST/SCA)
- 22.4 Penetration Testing and Vulnerability Remediation

Clause 23. Incident Response and Breach Notification

- 23.1 Detection, Containment, and Investigation Timelines
- 23.2 Regulatory Notification Windows by Jurisdiction
- 23.3 Customer Notification Content and Channels
- 23.4 Breach Register and Root Cause Analysis

Clause 24. Data Protection Impact Assessments (DPIAs)

- 24.1 DPIA Triggers and Thresholds
- 24.2 Risk Scoring Methodology
- 24.3 Mitigation Requirements by Risk Score
- 24.4 Pre-Launch Approval Workflow

Clause 25. Training, Awareness, and Culture

- 25.1 Mandatory Annual Training Requirements
- 25.2 Role-Specific Curriculum and Validation
- 25.3 New Hire Onboarding and Completion Enforcement
- 25.4 Phishing Simulations and Data Protection Week

Clause 26. Policy Compliance, Monitoring, and Enforcement

- 26.1 Audit Protocols and File Review Schedules
- 26.2 Evidence Repository and Metadata Standards
- 26.3 Violation Logging and Disciplinary Action

Clause 27. Exceptions and Risk Acceptance

- 27.1 Exception Submission and Risk Assessment
- 27.2 Approval Authorities by Risk Level
- 27.3 Maximum Duration and Prohibited Exceptions

Clause 28. Global and Local Alignment

- 28.1 Jurisdictional Addendum Process
- 28.2 Conflict Resolution Rule (Stricter Law Prevails)
- 28.3 Active Addenda for Rwanda, Nigeria, and California



Clause 29. Records Management Integration

29.1 Alignment with Enterprise Retention Policy

29.2 Legal Hold Issuance and Execution Protocol

Clause 30. Metrics, Reporting, and Board Oversight

30.1 Key Performance Indicators (KPIs)

30.2 Executive and Board Reporting Cadence

30.3 Dashboard Design and RAG Status

Clause 31. Policy Review and Continuous Improvement

31.1 Annual and Event-Driven Review Triggers

31.2 Feedback Channels and Stakeholder Input

31.3 Policy Update Workflow and Version Control

Appendices

Appendix A: Data Protection Control Matrix (Mapped to PCI DSS, GDPR, NDPR)

Appendix B: Data Classification Schema

Appendix C: Data Subject Request (DSR) Workflow

Appendix D: Cryptographic Standards Specification

Appendix E: Data Flow Mapping Template

Appendix F: Data Protection Impact Assessment (DPIA) Template

Appendix G: Glossary of Defined Terms



Clause 1. Interpretation and Definitions

1.1 Defined Terms

In this Policy, unless the context otherwise requires, the following terms shall bear the meanings ascribed below:

“Personal Data” means any information relating to an identified or identifiable natural person (“Data Subject”); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier, device fingerprint, behavioral pattern, or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural, or social identity of that natural person. Personal Data includes pseudonymised data but excludes anonymised data where re-identification is not reasonably possible.

“Special Categories of Personal Data” means Personal Data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data processed for the purpose of uniquely identifying a natural person, data concerning health, data concerning a natural person’s sex life or sexual orientation, financial vulnerability indicators, property details, marital status, family composition including names of children or spouse, or any other category designated as sensitive under applicable law, including Nigeria’s NDPR or Rwanda’s Law No. 058/2021.

“Child” means any natural person below the age of thirteen (13) years, or such a higher age as may be prescribed by local law in the jurisdiction of processing, including but not limited to sixteen (16) years under certain interpretations of GDPR Article 8.

“Restricted Data” means the highest sensitivity classification tier, encompassing unmasked Primary Account Numbers (PANs), full track data, cryptographic keys, system root credentials, biometric templates, and any other data element that, if compromised, would constitute a material breach of PCI DSS v4.0 or equivalent regulatory framework.



“Processing” means any operation or set of operations which is performed on Personal Data or on sets of Personal Data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure, or destruction.

“Data Controller” means **Orchestrate Global Inc.**, a corporation duly registered under the laws of the State of Wyoming, United States of America, with principal operational offices at 539 W. Commerce St #1714, Dallas, TX; 14 KG 601 St Remera Gasebo, Kigali, Rwanda; and Kwaji Cl, Maitama, Abuja, Nigeria, which alone or jointly with others determines the purposes and means of the Processing of Personal Data.

“Data Processor” means any natural or legal person, public authority, agency, or other body which processes Personal Data on behalf of the Data Controller pursuant to a binding Data Processing Agreement that incorporates Standard Contractual Clauses (SCCs) or equivalent safeguards.

“Personal Data Breach” means a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, Personal Data transmitted, stored, or otherwise processed, including incidents involving ransomware, insider threat, misconfigured cloud storage, or loss of portable media containing unencrypted Personal Data.

“Consent” means any freely given, specific, informed, and unambiguous indication of the Data Subject’s wishes by which they, by a statement or by a clear affirmative action, signify agreement to the Processing of Personal Data relating to them. Consent must be distinguishable from other matters, presented in intelligible and easily accessible form, using clear and plain language, and must not be bundled with general terms and conditions.

“Data Protection Impact Assessment (DPIA)” means a systematic process to identify, assess, and mitigate risks to the rights and freedoms of Data Subjects arising from Processing activities that are likely to result in a high risk, particularly those involving new technologies, large-scale processing of Special Categories of Personal Data, or systematic monitoring of publicly accessible areas.



“Encryption at Rest” means the application of AES-256 in Galois/Counter Mode (GCM) or equivalent authenticated encryption to all stored Restricted and Confidential Data, ensuring both confidentiality and integrity.

“Encryption in Transit” means the enforcement of TLS 1.3 with Perfect Forward Secrecy (PFS), certificate pinning, and mutual authentication for all internal and external communications involving Personal Data.

1.2 Construction and Interpretation Rules

The headings and subheadings in this Policy are inserted for convenience of reference only and shall not affect its interpretation. Words importing the singular include the plural and vice versa. References to “persons” include individuals, corporations, partnerships, and other legal entities.

References to statutes, regulations, or standards include such instruments as amended, re-enacted, replaced, or supplemented from time to time. Where any provision of this Policy conflicts with a mandatory provision of local law, the stricter requirement shall prevail, provided that such local law does not prohibit compliance with the core principles herein.

Clause 2. Policy Statement and Foundational Principles

2.1 Recognition of Privacy as a Fundamental Human Right

Orchestrate Global Inc. affirms that the right to privacy is a fundamental human right enshrined in international instruments, including the Universal Declaration of Human Rights, and codified in national constitutions and data protection statutes across its jurisdictions of operation.

The Company recognises that every Data Subject retains inherent dignity and autonomy over their Personal Data, and that the exercise of data subject rights—including access, rectification, erasure, and objection—is not a privilege but a non-derogable entitlement.

This Policy operationalises that recognition through enforceable controls, transparent practices, and a culture of accountability that permeates all levels of the organisation.

2.2 Commitment to Lawful, Fair, and Transparent Processing

The Company commits to Processing Personal Data in a manner that is lawful, fair, and transparent to the Data Subject. Lawfulness requires a valid legal basis for each Processing activity, documented in the Lawful Basis Register maintained by the Compliance function.

Fairness demands that Processing does not cause unjustified adverse effects on the Data Subject. Transparency obliges the Company to provide clear, timely, and accessible information about how



Personal Data is used, with whom it is shared, and how rights may be exercised. These obligations apply equally to customers, merchants, employees, contractors, and partners, regardless of geography or relationship duration.

2.3 Binding Nature and Consequences of Non-Compliance

This Policy constitutes a mandatory directive binding upon all directors, officers, employees, interns, contractors, agents, and third parties acting on behalf of Orchestrate Global Inc.

Non-compliance with any provision herein shall constitute a material breach of employment or contractual obligations and may result in disciplinary action up to and including termination of employment or contract, revocation of system access, referral to regulatory authorities, civil litigation, or criminal prosecution.

Local procedures may introduce jurisdiction-specific enhancements but shall not dilute, contradict, or circumvent the core obligations established herein.

Clause 3. Purpose and Scope of Application

3.1 Purpose of the Policy

The purpose of this Policy is to establish a globally consistent, technically robust, and legally compliant framework for the protection of all Personal Data processed by Orchestrate Global Inc. in the course of its business operations across the United States (Dallas, TX), Rwanda (Kigali), and Nigeria (Abuja).

It serves as the authoritative standard for embedding data protection by design and by default into product development, vendor management, incident response, and daily operations. The Policy ensures that the Company meets its obligations under PCI DSS v4.0, GDPR, CCPA/CPRA, NDPR, Rwanda Law No. 058/2021, and other applicable frameworks, while fostering a culture of trust, accountability, and ethical data stewardship.

3.2 Covered Legal Entities

This Policy applies to Orchestrate Global Inc., incorporated under the laws of the State of Wyoming, and all controlled operational subsidiaries and branches in Kigali, Rwanda, and Abuja, Nigeria. It does not extend to non-controlled affiliates, minority investments, or entities operating outside the Company's brand, risk umbrella, or management control.

3.3 Covered Data Subjects

The Policy covers all natural persons whose Personal Data is processed by the Company, including but not limited to customers, merchants, business partners, employees, contractors, vendors, and website users. It does not apply to anonymised data where re-identification is not reasonably possible using available technology, resources, and time.

3.4 Covered Data Types

The Policy governs all categories of data classified as Public, Internal, Confidential, or Restricted, including but not limited to full PANs, KYC/AML documentation, government-issued identification, biometric templates, transaction logs, system credentials, cryptographic keys, API payloads, and



internal business records. Public marketing content and fully anonymised analytics datasets fall outside the scope.

3.5 Covered Systems and Environments

All systems used to store, process, or transmit Personal Data are in scope, including cloud environments (AWS, Oracle Cloud, Azure), on-premise infrastructure, virtual machines, containers, mobile applications, APIs, backups, SIEM platforms, and log repositories. Personal devices not enrolled in the Company's Mobile Device Management (MDM) or Bring Your Own Device (BYOD) programme are excluded.

3.6 Covered Personnel and Third Parties

The Policy binds all personnel with access to Company systems or data, including executives, employees, interns, contractors, and third parties operating under a non-disclosure agreement (NDA). External parties with no contractual relationship to the Company are not covered.

Clause 4. Data Protection Governance and Accountability Framework

4.1 Three Lines of Defense Model

Orchestrate Global Inc. implements a globally recognised Three Lines of Defense model to ensure clear separation of duties, independent assurance, and effective risk management across all data protection activities.

The First Line comprises business units, product teams, and engineering functions that own and operate data processing activities; they are responsible for embedding data protection controls into daily operations, maintaining data inventories, and executing Data Subject Rights requests within mandated timelines.

The Second Line includes the Compliance, Legal, and Security functions, which establish policy standards, provide advisory support, monitor adherence, and manage regulatory relationships; the Chief Security & Compliance Officer (CSCO) serves as the global Data Protection Officer-equivalent and reports directly to the Chief Executive Officer on matters of strategic risk.

The Third Line consists of Internal Audit, which provides independent, objective assurance over the design and operating effectiveness of data protection controls through annual audits, thematic reviews, and ad hoc investigations. This model ensures that no single function holds unchecked authority over data governance, thereby reinforcing organisational resilience and regulatory credibility.

4.2 Role-Based Accountability Matrix

Accountability for data protection is explicitly assigned to defined roles across the organisation, with responsibilities codified in job descriptions, performance objectives, and contractual agreements. The Chief Executive Officer (CEO) bears ultimate accountability for establishing a culture of compliance, allocating sufficient resources to data protection initiatives, and ensuring that the Board receives timely, accurate reporting on material risks.



The Chief Commercial Officer (CCO) is responsible for ensuring that all customer-facing contracts, marketing practices, and partner engagements align with this Policy and do not introduce unauthorised data processing obligations.

The Chief Technology Officer (CTO) oversees the implementation of secure architecture, cryptographic controls, access management, and monitoring systems that enforce technical compliance with data protection requirements.

The Chief Security & Compliance Officer (CSCO) acts as the central authority for policy interpretation, regulatory liaison, DPIA approval, breach coordination, and training programme oversight.

Business Unit Heads serve as Data Owners for the Personal Data under their purview, with responsibility for data accuracy, minimisation, lawful basis documentation, and retention compliance. Engineering Leads function as Data Stewards, implementing controls at the system level and ensuring that privacy by design principles are integrated into the software development lifecycle.

All personnel, regardless of role, are accountable for completing mandatory training, reporting suspected breaches within one hour of discovery, and adhering to access restrictions commensurate with their duties. Failure to discharge these responsibilities constitutes grounds for disciplinary action under the Company's Human Resources Policy.

4.3 Incident Escalation and Reporting Protocol

All suspected or confirmed Personal Data Breaches must be reported immediately—no later than one hour after discovery—to the Chief Security & Compliance Officer via the designated incident hotline or secure reporting portal.

Upon receipt, the CSCO shall activate the Incident Response Team, which includes representatives from Security, Legal, Engineering, Communications, and Customer Success. The team shall execute containment measures within twenty-four (24) hours, initiate forensic investigation, and assess the scope, impact, and root cause of the incident.

Based on severity scoring—defined by data sensitivity, volume affected, jurisdictional exposure, and potential harm—incidents are classified into four tiers: Level 1 (low impact, internal only), Level 2 (moderate impact, limited external notification), Level 3 (high impact, regulatory reporting required), and Level 4 (critical impact, executive and board escalation). Level 4 incidents, such as unauthorised access to unmasked PANs or systemic compromise of the Cardholder Data Environment (CDE), require verbal notification to the CEO and CCO within twelve (12) hours and written briefing to the Board Risk Committee within twenty-four (24) hours.

Regulatory notifications shall be submitted within seventy-two (72) hours where mandated by GDPR, NDPR, or Rwanda Law No. 058/2021, and within forty-five (45) days under CCPA/CPRA if a significant risk to consumer rights exists.

Customer notifications shall be issued within five (5) business days of confirmation, using empathetic, plain-language templates approved by Legal and reviewed for cultural appropriateness in each jurisdiction. All incidents are logged in the central Breach Register, which includes timestamp,



classification, response actions, regulatory filings, customer communications, and post-incident review outcomes.

4.4 Board and Executive Oversight Responsibilities

The Board of Directors exercises strategic oversight of the Company's data protection posture through its Risk and Compliance Committee, which meets quarterly to review key risk indicators, audit findings, breach trends, and regulatory developments.

The Committee receives a standardised dashboard containing metrics on Data Subject Rights fulfilment rates, training completion percentages, open high-risk findings, third-party audit status, and DPIA pipeline volume. The CEO presents an annual attestation confirming that the organisation maintains adequate resources, skilled personnel, and technical capabilities to meet its data protection obligations.

The Executive Committee (ExCo), comprising the CEO, CCO, CTO, and CSCO, convenes monthly to review operational performance against data protection KPIs, approve exceptions, and resolve cross-functional conflicts. ExCo also reviews all Level 3 and Level 4 incidents within seven (7) days of resolution to validate root cause analysis and ensure corrective actions are implemented. This layered governance structure ensures that data protection remains a standing agenda item at the highest levels of leadership, reinforcing its status as a core business imperative rather than a peripheral compliance function.

Clause 5. Legal and Regulatory Foundations

5.1 Applicable Statutory and Sectoral Frameworks

Orchestrate Global Inc. acknowledges that its data processing activities are subject to a complex and evolving landscape of statutory, regulatory, and sectoral obligations across its operational jurisdictions. In the United States, the Company complies with the California Consumer Privacy Act (CCPA), as amended by the California Privacy Rights Act (CPRA), which grants California residents specific rights regarding the collection, use, and sharing of their Personal Information, including the right to opt out of "sales" or "sharing" for cross-context behavioural advertising.

The Company also adheres to sector-specific requirements under federal law, including the Gramm-Leach-Bliley Act (GLBA) Safeguards Rule where applicable, and maintains alignment with emerging state privacy laws in Virginia, Colorado, Connecticut, Utah, and Texas through a dynamic compliance mapping process. In Nigeria, the Company fully implements the Nigeria Data Protection Regulation (NDPR) issued by the National Information Technology Development Agency (NITDA), including mandatory registration as a Data Controller, appointment of a Data Protection Officer (functionally fulfilled by the CSCO), implementation of technical and organisational measures proportional to risk, and adherence to data localisation principles for Nigerian citizens' data.

In Rwanda, the Company operates in strict accordance with Law No. 058/2021 of 13/10/2021 Relating to the Protection of Personal Data, which prohibits the transfer of Rwandan citizens' Personal Data outside the Republic without prior authorisation from the Rwanda Utilities Regulatory Authority (RURA), unless explicit consent is obtained or the transfer is necessary for the performance of a contract. The Company further recognises the extraterritorial reach of the European Union's



General Data Protection Regulation (GDPR), which applies to Processing activities related to offering goods or services to, or monitoring the behaviour of, individuals in the EU, and accordingly implements GDPR-equivalent safeguards—including lawful basis documentation, Data Subject Rights fulfilment within thirty (30) days, and mandatory breach notification to supervisory authorities within seventy-two (72) hours of awareness—for all such activities.

These frameworks are not treated as isolated checklists but as interlocking layers of obligation that collectively define the Company's global data protection baseline.

5.2 Contractual and Standard-Based Obligations

Beyond statutory requirements, Orchestrate Global Inc. is bound by a range of contractual and industry-standard obligations that impose additional or more stringent data protection duties. Chief among these is the Payment Card Industry Data Security Standard (PCI DSS) v4.0, to which the Company is contractually obligated as a service provider handling cardholder data on behalf of acquiring banks and merchants.

PCI DSS v4.0 mandates a comprehensive set of security controls across twelve requirement domains, including the maintenance of a secure network, protection of stored cardholder data through strong cryptography, implementation of robust access control measures, regular testing of security systems, and maintenance of an information security policy. The Company treats PCI DSS not as a point-in-time audit exercise but as a continuous operational discipline, with automated evidence collection, quarterly internal assessments, and annual validation by a Qualified Security Assessor (QSA).

Additionally, all customer contracts include binding Data Processing Addenda that incorporate the European Commission's Standard Contractual Clauses (SCCs) for international data transfers, updated to reflect the 2021 recommendations and Schrems II jurisprudence. These addenda impose specific obligations on the Company as Data Processor, including restrictions on sub-processing, assistance with Data Subject Rights, confidentiality of personnel, and prompt breach notification within twenty-four (24) hours of confirmation.

Cloud service agreements with providers such as AWS, Oracle Cloud, and Microsoft Azure further require adherence to shared responsibility models, with the Company retaining full accountability for data classification, encryption key management, access governance, and incident response coordination. These contractual commitments are catalogued in the Third-Party Risk Register and reviewed during vendor due diligence and contract renewal cycles.

5.3 Regulatory Obligations Register (ROR)

To ensure systematic and auditable compliance across this multi-jurisdictional landscape, the Company maintains a centralized Regulatory Obligations Register (ROR), housed within the Governance, Risk, and Compliance (GRC) platform. The ROR is a living document that maps each legal, regulatory, or contractual requirement to a specific control, responsible owner, implementation status, evidence artifact, automation flag, and review date. For example, GDPR Article 30 (Record of Processing Activities) is mapped to Control OG-DPP-001-C087, owned by the Data Governance Lead, implemented via the central Data Inventory, evidenced by quarterly snapshots, and marked as "automated" due to API integration with the CMDB. Similarly, PCI DSS Requirement 3.5 (Key



Management) is linked to Control OG-DPP-001-C142, owned by the Cryptography Architect, implemented through HSM-based key lifecycle management, evidenced by KMS audit logs, and flagged for manual validation during QSA assessments.

The ROR is updated within five (5) business days of any new law enactment, regulatory guidance publication, or material contract amendment. It is reviewed quarterly by the Compliance function and presented to the Risk and Compliance Committee as part of the Board oversight package. This register serves as the single source of truth for demonstrating compliance to regulators, auditors, and customers, and ensures that no obligation falls through the cracks due to jurisdictional complexity or organisational change.

Clause 6. Data Classification Framework

6.1 Classification Levels and Criteria

Orchestrate Global Inc. implements a four-tier data classification framework to ensure that protective measures are commensurate with the sensitivity, criticality, and regulatory implications of the information in question.

All data created, received, stored, or transmitted by the Company must be assigned to one—and only one—of the following classification levels at the point of creation or ingestion, with reclassification occurring whenever the nature or context of the data changes materially. The classification decision must be documented in the central Data Inventory and reviewed annually by the designated Data Owner.

Public: Information that has been formally approved for unrestricted disclosure to the general public, such as press releases, marketing brochures, public API documentation, or anonymised statistical reports where re-identification risk has been demonstrably eliminated through irreversible aggregation and suppression techniques. Public data may be shared without restriction but must not contain embedded metadata, watermarks, or hidden fields that reveal non-public information.

Internal: Information intended for use within Orchestrate Global Inc. or its controlled subsidiaries, which, if disclosed externally, could cause minor reputational impact, operational inconvenience, or competitive disadvantage but would not constitute a legal or regulatory violation. Examples include organisational charts, internal meeting notes, non-sensitive project plans, and generic training materials. Internal data must not be shared with external parties without prior approval from the Data Owner and must be marked with an “Internal Use Only” footer.

Confidential: Information that is sensitive in nature and whose unauthorised disclosure, alteration, or destruction could result in material harm to individuals, the Company, or its partners, including legal liability, regulatory penalty, financial loss, or significant reputational damage.

This category encompasses all Personal Data as defined under GDPR, CCPA, NDPR, and Rwanda Law No. 058/2021, including names, contact details, government identification numbers, KYC/AML documentation, transaction histories, employment records, and commercial contracts. Confidential data must be protected through role-based access controls, multi-factor authentication, audit logging, and encryption both in transit and at rest.



Restricted: The highest sensitivity tier, reserved for data elements whose compromise would constitute a critical breach of trust, regulatory mandate, or system integrity. Restricted data includes unmasked Primary Account Numbers (PANs), full magnetic stripe data, CVV2/CVC2 codes, cryptographic keys (both data encryption keys and key encryption keys), system root credentials, biometric templates used for authentication, and any other data element explicitly designated as “high risk” under PCI DSS v4.0 Requirement 3 or equivalent frameworks.

Processing of Restricted data is permitted only within the Cardholder Data Environment (CDE) or equivalent secure enclave, with access limited to personnel who have undergone enhanced background checks, signed supplemental confidentiality agreements, and completed specialised security training. Storage of Restricted data outside FIPS 140-2 Level 3 validated Hardware Security Modules (HSMs) or cloud-native equivalents (e.g., AWS CloudHSM, Azure Dedicated HSM) is strictly prohibited.

6.2 Handling Requirements by Classification Tier

Each classification level mandates specific technical and organisational controls that must be implemented without exception. For *Public* data, no access restrictions apply, but content must undergo pre-publication review by the Communications and Legal teams to ensure compliance with brand guidelines and absence of embedded non-public information.

For *Internal* data, access is restricted to authenticated employees and contractors via corporate identity providers; sharing via personal email, consumer-grade file-sharing services, or unapproved collaboration platforms is forbidden.

For *Confidential* data, the following controls are mandatory:

- (a) storage only in systems that enforce encryption at rest using AES-256-GCM or equivalent;
- (b) transmission only over TLS 1.3 with Perfect Forward Secrecy and certificate pinning;
- (c) access granted solely on a need-to-know basis through role-based access control (RBAC) policies reviewed quarterly;
- (d) all access and modification events logged to an immutable SIEM with retention of at least 365 days; and
- (e) inclusion in the central Data Inventory with accurate metadata on purpose, lawful basis, retention period, and third-party recipients.

For *Restricted* data, additional controls apply:

- (a) cryptographic keys must never be exposed to application layers and must be managed exclusively within HSMs
- (b) PANs must be tokenised or truncated immediately after authorisation, with full PANs retained only for the minimum duration necessary for dispute resolution or regulatory reporting



(c) all access requires just-in-time privileged session approval, dual control for critical operations, and real-time anomaly detection

(d) network segmentation must isolate the CDE from all other environments, with firewall rules defaulting to deny-all; and

(e) physical and logical access logs must be preserved for a minimum of seven (7) years to satisfy PCI DSS and financial regulatory requirements.

6.3 Misclassification Remediation and Enforcement

Data Owners are accountable for the initial and ongoing accuracy of classification assignments. Any misclassification—whether intentional or inadvertent—must be corrected within five (5) business days of discovery.

Repeated or systemic misclassification by a Business Unit Head constitutes a reportable control deficiency and triggers a formal review by the Compliance function. Automated data discovery and classification tools are deployed across cloud storage, email systems, and endpoint devices to detect and flag potential misclassifications based on content patterns, metadata, and context. False negatives (e.g., Restricted data labelled as Internal) are treated as high-risk findings and escalated to the Incident Response Team if exposure is suspected.

Personnel who knowingly process data at a lower classification than warranted, thereby circumventing security controls, are subject to disciplinary action up to and including termination, in accordance with the Company's Code of Conduct. Quarterly access certifications include validation of data classification accuracy, and audit trails of classification changes are maintained for forensic review. This framework ensures that data protection is not left to individual discretion but is enforced through architecture, policy, and accountability.

Clause 7. Data Inventory, Mapping, and Ownership

7.1 Centralized Data Inventory Requirements

Orchestrate Global Inc. maintains a single, authoritative, and dynamically updated Data Inventory that serves as the cornerstone of its data protection governance framework. This inventory is not a static spreadsheet but an integrated component of the Company's Governance, Risk, and Compliance (GRC) platform, populated through automated discovery agents, API integrations with cloud service providers, and manual input from designated Data Owners.

Every recordable instance of Personal Data Processing—whether conducted internally or by a third-party processor on the Company's behalf—must be documented in the inventory within ten (10) business days of initiation.

Each entry must include, at a minimum, the following metadata fields:

(a) unique system or process identifier;

(b) Data Owner (Business Unit Head);



- (c) Data Steward (Engineering Lead);
- (d) categories of Data Subjects affected;
- (e) types of Personal Data processed, mapped to the classification framework under Clause 6;
- (f) lawful basis for Processing under Clause 9;
- (g) purpose of Processing, stated with specificity and granularity;
- (h) retention period, aligned with Clause 13;
- (i) geographic locations of storage and processing, including cloud regions and physical data centres;
- (j) list of internal and external recipients, including sub-processors;
- (k) cross-border transfer mechanisms, if applicable;
- (l) encryption status at rest and in transit;
- (m) link to associated Data Protection Impact Assessment (DPIA), if required under Clause 24; and
- (n) last review date and reviewer identity.

The inventory is reconciled monthly against network traffic logs, access control lists, and vendor contracts to ensure completeness and accuracy. Any discrepancy between actual data flows and inventory records constitutes a reportable control gap and must be remediated within fifteen (15) business days.

The Chief Security & Compliance Officer (CSCO) retains ultimate authority over the structure, integrity, and auditability of the Data Inventory, and may suspend Processing activities that remain undocumented beyond the mandated timeline.

7.2 Data Flow Mapping Standards

Complementing the structured inventory, the Company requires detailed Data Flow Maps for all systems that process Confidential or Restricted data, particularly those interfacing with customers, merchants, or financial institutions.

These maps are visual and narrative artefacts that depict the end-to-end journey of Personal Data from collection through disposal, including all touchpoints, transformations, storage repositories, transmission channels, and third-party interfaces.

Each map must be created using the standard template provided in Appendix E and must identify:

- (a) data entry points (e.g., web forms, mobile SDKs, batch uploads);
- (b) validation and transformation logic applied upon ingestion;
- (c) internal system components involved in Processing (e.g., authentication services, fraud engines, payment gateways);



- (d) data stores, including database schemas, field-level encryption status, and backup locations;
- (e) outbound data sharing, including APIs, file transfers, and real-time streams;
- (f) security controls at each boundary (e.g., WAF rules, API gateways, TLS termination points); and
- (g) data deletion or anonymisation triggers. Maps must be version-controlled, signed off by both the Data Owner and Data Steward, and updated within thirty (30) days of any architectural change, such as the introduction of a new microservice, migration to a different cloud region, or integration with a third-party analytics provider.

For PCI DSS compliance, separate Cardholder Data Flow Diagrams must be maintained for the Cardholder Data Environment (CDE), showing all connected systems, segmentation controls, and scope boundaries. These diagrams are subject to annual validation by the Internal Audit function and must be made available to Qualified Security Assessors (QSAs) during PCI assessments.

The purpose of these maps is not merely documentation but risk visibility: they enable the Company to trace the provenance of data in breach scenarios, assess the impact of vendor failures, and demonstrate to regulators that Processing is purpose-limited and technically constrained.

7.3 Data Ownership and Stewardship Model

Accountability for data is institutionalised through a dual-role model of Data Ownership and Data Stewardship. The Data Owner is always a senior business leader—typically a Business Unit Head or Product Director—who bears ultimate responsibility for the lawfulness, necessity, and proportionality of the Processing activity under their purview.

The Data Owner must ensure that:

- (a) a valid lawful basis exists and is documented before Processing begins;
- (b) data minimisation principles are applied, so that only necessary fields are collected and retained;
- (c) Data Subject Rights requests related to their systems are fulfilled accurately and within statutory timelines;
- (d) retention schedules are enforced through automated purge jobs; and
- (e) any proposed changes to Processing purposes undergo re-evaluation for legal and ethical compliance. The Data Steward is a technical lead—usually a Principal Engineer or DevOps Manager—who is responsible for implementing and maintaining the technical controls that enforce data protection requirements at the system level.

The Data Steward must ensure that:

- (a) data is stored and transmitted in accordance with the classification mandates of Clause 6;
- (b) access controls are configured per least privilege and reviewed quarterly;
- (c) logging and monitoring are comprehensive and tamper-resistant;



(d) cryptographic standards are applied uniformly; and

(e) system changes do not inadvertently expand data scope or weaken security boundaries. Both roles are formally appointed in writing, listed in the GRC system, and included in incident response call trees.

They participate jointly in DPIA reviews, vendor risk assessments, and breach investigations. Failure by either party to discharge their duties constitutes a material control failure and is escalated to the Executive Committee. This model ensures that data protection is not siloed in compliance or engineering but is a shared, visible, and enforceable partnership between business and technology leadership.

Clause 8. Core Principles of Lawful Data Processing

8.1 Lawfulness, Fairness, and Transparency

All Processing of Personal Data by Orchestrate Global Inc. must be grounded in a valid legal basis as defined under Clause 9, documented in the Lawful Basis Register, and demonstrably aligned with the reasonable expectations of the Data Subject.

Lawfulness requires that no Processing occurs in the absence of consent, contractual necessity, legal obligation, vital interest, public task, or legitimate interest supported by a documented balancing test. Fairness demands that Processing does not result in unjustified adverse consequences for the Data Subject, including discriminatory outcomes, financial harm, reputational damage, or loss of autonomy.

The Company prohibits the use of dark patterns, pre-ticked boxes, bundled consents, or misleading language that could impair informed decision-making. Transparency obliges the Company to provide clear, concise, and timely information about how Personal Data is used, through layered privacy notices delivered at the point of collection and updated within thirty (30) days of any material change.

These notices must specify the identity and contact details of the Data Controller, the purposes and legal bases for Processing, the categories of recipients, the existence of cross-border transfers and associated safeguards, the retention period or criteria used to determine it, and the rights available to the Data Subject.

Notices must be written in plain language accessible to the average user, translated into local languages for local markets (e.g. Kinyarwanda, Hausa, English etc), and made available in alternative formats upon request. Automated decision-making that produces legal or similarly significant effects must be accompanied by meaningful human review and an explicit right to contest the outcome.

These obligations apply uniformly across digital interfaces, paper-based forms, call centre scripts, and third-party integrations.

8.2 Purpose Limitation

Personal Data shall be collected only for specified, explicit, and legitimate purposes and shall not be further Processed in a manner that is incompatible with those purposes. At the time of collection, the purpose must be articulated with sufficient granularity to enable informed consent or legitimate interest assessment—for example, “fraud prevention during onboarding” rather than “security.”



Any proposed secondary use of data—such as repurposing KYC documents for marketing analytics or using transaction history to infer creditworthiness—constitutes a new Processing activity that requires independent legal basis validation, Data Subject notification, and, where required, fresh consent.

System architecture enforces purpose limitation through logical separation of data stores: customer authentication logs are segregated from behavioural analytics pipelines, and payment processing databases are isolated from CRM systems.

Data Owners must certify annually that no incompatible Processing has occurred within their domain, and any deviation must be reported as a policy exception under Clause 27. This principle ensures that data is not treated as a fungible asset but as a trust-bound resource tied to a specific, justifiable objective.

8.3 Data Minimisation

The Company shall collect and retain only such Personal Data as is adequate, relevant, and limited to what is necessary in relation to the purposes for which it is Processed. This principle is enforced at multiple layers:

- (a) at the user interface, optional fields are suppressed by default and require active opt-in;
- (b) at the API layer, schema validation rejects payloads containing non-essential attributes;
- (c) at the database level, tables are normalised to eliminate redundant columns; and
- (d) at the business process level, annual data rationalisation initiatives identify and purge legacy fields that no longer serve a documented purpose. For example, while date of birth may be required for age verification during onboarding, the full birth certificate is not retained unless mandated by AML regulations—and even then, only the relevant excerpt is stored.

Similarly, geolocation data is collected only when necessary for service delivery (e.g., fraud scoring based on transaction location) and is truncated to city-level precision where possible. Data minimisation is validated during Secure SDLC reviews and DPIAs, and violations are treated as high-risk findings. This discipline reduces attack surface, limits breach impact, and reinforces respect for individual autonomy.

8.4 Accuracy and Integrity

Personal Data must be accurate and, where necessary, kept up to date. Inaccurate data must be erased or rectified without undue delay. The Company provides self-service portals enabling Data Subjects to view and correct their information in real time, with changes propagated across integrated systems within five (5) business days.

For data not amenable to self-correction—such as government-issued identification or transaction records—the Company maintains a formal rectification request workflow under Clause 11.2, with resolution timelines not exceeding fifteen (15) business days.

All corrections are logged with immutable audit trails recording the requester, timestamp, original value, new value, and authorising agent. Preventive controls include input validation rules, duplicate



detection algorithms, and reconciliation jobs that flag inconsistencies between source systems. The integrity of data is further protected through cryptographic hashing of critical records and write-once-read-many (WORM) storage for immutable logs.

This principle ensures that decisions based on Personal Data—whether automated or human—are grounded in truth, not error.

8.5 Storage Limitation

Personal Data shall be kept in a form which permits identification of Data Subjects for no longer than is necessary for the purposes for which the data is Processed. Retention periods are defined in Clause 13 and encoded into automated lifecycle management policies across all systems. Upon expiry, data is either securely deleted or anonymised using irreversible techniques that prevent re-identification.

Legal holds issued by the Legal department suspend automated disposal until the hold is formally lifted, with manual certification required for each extension. The Company prohibits indefinite retention under vague justifications such as “future research” or “business intelligence.” Instead, retention schedules are tied to specific regulatory mandates (e.g., seven years for KYC under AML laws), contractual obligations (e.g., duration of merchant agreement plus three years), or demonstrable business needs (e.g., active customer status).

Quarterly audits verify compliance with retention rules, and orphaned data—records with no linked active relationship—is purged automatically. This principle balances operational necessity with the right to be forgotten and reduces long-term liability.

8.6 Integrity and Confidentiality

Personal Data must be processed in a manner that ensures appropriate security, including protection against unauthorised or unlawful processing and against accidental loss, destruction, or damage, using appropriate technical and organisational measures.

These measures include, but are not limited to: encryption at rest (AES-256-GCM) and in transit (TLS 1.3 with PFS); multi-factor authentication for all remote and privileged access; role-based access controls reviewed quarterly; network segmentation isolating sensitive environments; endpoint detection and response (EDR) on all corporate devices; secure software development lifecycle with automated code scanning; and physical security controls at co-location facilities.

The strength of these measures is calibrated to the classification level of the data under Clause 6, with Restricted data receiving the highest assurance. Annual penetration testing, vulnerability scanning, and red team exercises validate the effectiveness of these controls. This principle recognises that privacy cannot exist without security, and that trust is earned through demonstrable resilience.

8.7 Accountability

Orchestrate Global Inc. bears full responsibility for demonstrating compliance with all principles set forth in this Clause. Accountability is not a passive state but an active practice, evidenced through documented policies, trained personnel, auditable logs, risk assessments, and executive oversight.

The Company maintains a comprehensive control framework mapped to each principle, with ownership assigned, implementation verified, and effectiveness measured through key performance



indicators. Regulatory inquiries, Data Subject complaints, and internal audits are treated as opportunities to validate and improve accountability mechanisms.

The Chief Executive Officer attests annually to the Board that the organisation has implemented appropriate measures to comply with its data protection obligations. This final principle binds all others together, transforming ethical commitments into operational reality.

Clause 9. Lawful Bases for Processing

9.1 Consent as a Lawful Basis

Consent may serve as a lawful basis for Processing only when it is freely given, specific, informed, and unambiguous, and when the Data Subject has the genuine ability to refuse or withdraw without detriment.

Consent must be obtained through a clear affirmative action—such as ticking an unchecked box, clicking an “I Agree” button, or signing a physical form—and must never be inferred from silence, pre-ticked boxes, or general terms of service. The request for consent must be presented separately from other matters, in intelligible and easily accessible form, using clear and plain language that explains the purpose, data types, recipients, retention period, and right to withdraw.

For Special Categories of Personal Data—including biometrics, financial vulnerability indicators, or health-related information—explicit consent is required, meaning a written or electronic statement that leaves no room for doubt as to the Data Subject’s intention.

In the case of Children under the age of thirteen (13), verifiable parental or guardian consent must be obtained through secure, auditable means such as government-issued ID verification coupled with two-factor authentication; marketing or profiling based on children’s data is categorically prohibited.

All consents are captured, versioned, and stored in the central Consent Management Platform, which logs the timestamp, IP address, user agent, privacy notice version, and granular purposes accepted. Withdrawal of consent must be as easy as giving it and must trigger automated suppression workflows that halt Processing within twenty-four (24) hours.

Consent records are retained for the duration of the relationship plus three (3) years to support audit and dispute resolution. Consent shall not be used as a lawful basis where a more appropriate basis exists—such as contract or legal obligation—as doing so would undermine its integrity as a voluntary choice.

9.2 Contractual Necessity

Processing is lawful where it is necessary for the performance of a contract to which the Data Subject is party or in order to take steps at the request of the Data Subject prior to entering into a contract.

This basis applies narrowly to data that is objectively essential to delivering the core service—for example, collecting name, contact details, and payment instrument information to process a transaction; verifying identity documents to onboard a merchant; or transmitting shipping addresses to a logistics partner. It does not extend to ancillary activities such as marketing, analytics, or fraud prevention beyond what is strictly required to fulfil the contractual obligation.



If a processing activity could reasonably be carried out without accessing certain Personal Data, then that data falls outside the scope of contractual necessity and requires an alternative lawful basis. All systems invoking this basis must document the specific contractual clause or service requirement that necessitates the Processing, and Data Owners must certify annually that no overreach has occurred. This basis is particularly relevant in the United States and Nigeria, where it provides a stable foundation for commercial operations without reliance on consent.

9.3 Legal Obligation

Processing is lawful where it is necessary for compliance with a legal obligation to which the Controller is subject under Union, Member State, or national law. This includes obligations arising from anti-money laundering (AML) and know-your-customer (KYC) regulations, tax reporting requirements, court orders, regulatory subpoenas, and statutory retention mandates.

The legal obligation must be specific, binding, and enforceable—not aspirational or discretionary. For example, retaining KYC documentation for seven (7) years post-relationship termination is mandated under AML frameworks in multiple jurisdictions and constitutes a valid basis for continued storage.

However, voluntarily sharing customer data with a government agency absent a warrant or statutory compulsion does not qualify. Legal counsel must validate all invocations of this basis, and the underlying statute or regulation must be cited in the Lawful Basis Register. Data collected under legal obligation must be segregated from other datasets and accessed only by authorised personnel for the stated purpose. This basis ensures that the Company meets its civic duties while preventing mission creep under the guise of compliance.

9.4 Vital Interests

Processing is lawful where it is necessary to protect the vital interests of the Data Subject or of another natural person. This basis is interpreted restrictively and applies only in life-or-death situations—such as disclosing medical or location data to emergency services during a cardiac event, or sharing contact information with law enforcement to prevent imminent physical harm. It does not extend to financial, reputational, or psychological harms, however severe.

Due to its exceptional nature, any Processing under this basis must be reviewed and approved in real time by the Chief Security & Compliance Officer or designated deputy, with a written justification recorded in the incident log within one (1) hour. Post-event, the Data Subject must be notified unless doing so would impede the emergency response. This basis is rarely invoked but remains a critical safeguard in extreme circumstances.

9.5 Public Interest and Official Authority

Processing is lawful where it is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the Controller. Orchestrate Global Inc., as a private corporation, does not typically exercise official authority.

However, this basis may apply when the Company acts as a delegated agent of a public body—for example, operating a government-mandated payment infrastructure or administering a state-sponsored financial inclusion programme under formal agreement. In such cases, the public mandate must be documented in a binding contract or legislative instrument, and Processing must be limited strictly to



the scope of that delegation. Absent such a mandate, this basis is not available to the Company. Any proposed use must be pre-approved by Legal and the CSCO, with ongoing oversight by the Risk and Compliance Committee.

9.6 Legitimate Interests and Balancing Tests

Processing is lawful where it is necessary for the purposes of the legitimate interests pursued by the Controller or by a third party, except where such interests are overridden by the fundamental rights and freedoms of the Data Subject. Legitimate interests may include fraud prevention, network security, internal auditing, direct marketing (subject to opt-out rights under CCPA), product improvement, and group-wide reporting.

However, this basis is not a default fallback; it requires a structured, documented assessment that weighs the Company's interest against the Data Subject's reasonable expectations, privacy impact, and potential harm. The assessment must consider:

- (a) the nature of the data (e.g., Special Categories receive heightened scrutiny);
- (b) the context of collection (e.g., B2B vs. B2C);
- (c) the safeguards in place (e.g., pseudonymisation, access controls); and
- (d) the availability of less intrusive alternatives. A Legitimate Interest Assessment (LIA) template, maintained in Appendix F, must be completed and signed off by the Data Owner, Legal, and CSCO before Processing commences. LIAs are reviewed annually or upon material change in context.

Data Subjects must be informed of the legitimate interest basis in the privacy notice and provided a clear mechanism to object, which must be honoured without undue delay. This basis demands humility: the Company's interests do not automatically prevail, and the burden of proof rests entirely on the Controller.

Clause 10. Consent Management Framework

10.1 Standards for Valid Consent

Consent, when relied upon as a lawful basis for Processing under Clause 9.1, must meet the highest threshold of voluntariness, specificity, and informed understanding. It shall not be considered freely given if the Data Subject has no genuine or free choice, or if they are unable to refuse or withdraw consent without detriment, such as denial of core service functionality, imposition of fees, or degradation of user experience.

Consent requests must be unbundled from general terms and conditions and presented in a manner that is clearly distinguishable, using layered notices that allow the Data Subject to drill down into granular purposes without overwhelming cognitive load. Each purpose must be stated with precision—for example, “to send promotional offers about payment products via email” rather than “for marketing”—and each category of third-party recipient must be identified by function or name where feasible.



The language used must be plain, jargon-free, and translated into the primary languages of the jurisdictions served: English for Nigeria and the United States, Kinyarwanda for Rwanda, and French where required by Rwandan regulatory guidance, etc.

Pre-ticked boxes, opt-out mechanisms, continued browsing, or silence shall never constitute valid consent. For electronic consent, the interface must log technical metadata including timestamp, IP address, user agent, screen resolution, and interaction sequence to support forensic validation in the event of dispute. Consent obtained under duress, deception, or material misrepresentation is void ab initio and must be purged from all systems upon discovery.

10.2 Consent Capture, Storage, and Versioning

All consents must be captured through the Company's central Consent Management Platform (CMP), a tamper-evident, access-controlled system integrated with customer identity and onboarding workflows.

At the moment of capture, the CMP shall record:

- (a) the unique identifier of the Data Subject;
- (b) the specific purposes accepted, each mapped to a machine-readable code;
- (c) the version of the privacy notice displayed, stored as an immutable snapshot;
- (d) the method of capture (e.g., web form, mobile SDK, call centre script);
- (e) the identity of any agent involved in verbal consent collection; and
- (f) cryptographic hash of the session transcript where applicable. Consent records are classified as confidential data under Clause 6 and protected with AES-256-GCM encryption at rest and TLS 1.3 in transit.

The CMP maintains a complete historical ledger of all consent events, including grants, withdrawals, and updates, enabling point-in-time reconstruction of a Data Subject's preference state. Privacy notice versions are archived indefinitely, ensuring that future audits can verify the exact information provided at the time of consent. Manual overrides or bulk uploads of consent data are prohibited unless pre-approved by the Chief Security & Compliance Officer and accompanied by a compensating control log.

This architecture ensures that consent is not a one-time checkbox but a persistent, verifiable state.

10.3 Withdrawal Mechanisms and Automation

Data Subjects have the absolute right to withdraw consent at any time, and such withdrawal must be as easy as giving consent.

The Company provides multiple, low-friction channels for withdrawal:

- (a) a dedicated "Privacy Preferences" portal accessible from all customer dashboards;



- (b) an unsubscribe link in every marketing email, compliant with CAN-SPAM and GDPR;
- (c) a toll-free number staffed by trained agents for voice-based requests; and
- (d) in-person assistance at designated service centres in Abuja and Kigali. Upon receipt of a withdrawal request—whether partial (e.g., revoke marketing consent only) or total—the CMP triggers an automated workflow that propagates the instruction across all relevant systems within twenty-four (24) hours.

This includes halting email campaigns, disabling analytics tracking, revoking API access tokens, and flagging records for suppression in batch processes. Confirmation of withdrawal is sent to the Data Subject within forty-eight (48) hours, detailing the actions taken and any residual Processing that continues under an alternative lawful basis (e.g., contractual necessity).

Withdrawal does not affect the lawfulness of Processing based on consent prior to its withdrawal, but it does terminate all future Processing tied exclusively to that consent. Logs of withdrawal requests and execution outcomes are retained for three (3) years to support audit and regulatory inquiry.

10.4 Special Provisions for Sensitive and Child Data

Consent for the Processing of Special Categories of Personal Data—as defined in Clause

1.1—requires explicit, unambiguous affirmation that is separate from any other consent request. This may take the form of a standalone declaration, digital signature, or two-step verification process that confirms the Data Subject’s awareness of the sensitivity involved.

Examples include biometric enrolment for authentication, disclosure of financial hardship indicators for hardship programmes, or submission of health documentation for insurance-linked services. In all such cases, the privacy notice must include a prominent warning about the nature of the data and the heightened risks of misuse.

For Children under the age of thirteen (13), the Company prohibits the collection of Personal Data unless verifiable parental or guardian consent is obtained through a robust, multi-factor process.

This process requires:

- (a) submission of government-issued identification for both the parent/guardian and the child;
- (b) liveness detection and facial matching to confirm co-presence;
- (c) knowledge-based authentication questions derived from public records; and
- (d) a confirmation email or SMS sent to a pre-verified contact channel. The resulting consent is recorded with dual identifiers (child and parent) and flagged in the CMP for enhanced monitoring.

Any processing of children’s data is subject to quarterly review by the Compliance function, and automated alerts are triggered by anomalous access patterns. Marketing, profiling, behavioural advertising, or social sharing features are categorically disabled for child accounts. These measures



reflect the Company's recognition that children lack the capacity to fully comprehend digital risks and therefore deserve the highest standard of protection.

Clause 11. Data Subject Rights (DSR) Framework

11.1 Right of Access

Every Data Subject has the right to obtain confirmation as to whether or not Personal Data concerning them is being Processed, and, where that is the case, access to that data and supplementary information.

Upon receipt of a valid request, the Company shall provide, without undue delay and in any event within thirty (30) days of verification, a complete copy of all Personal Data undergoing Processing, including:

- (a) the purposes of the Processing;
- (b) the categories of Personal Data concerned;
- (c) the recipients or categories of recipients to whom the data have been or will be disclosed;
- (d) where possible, the envisaged retention period or criteria used to determine it;
- (e) the existence of the right to request rectification, erasure, restriction, or objection;
- (f) the right to lodge a complaint with a supervisory authority;
- (g) the source of the data if not collected directly from the Data Subject; and
- (h) the existence of automated decision-making, including profiling, and meaningful information about the logic involved.

The response shall be provided in a commonly used electronic format (e.g., JSON, PDF, or CSV), structured to enable portability, and delivered through a secure, authenticated channel. Where the request is manifestly unfounded or excessive—particularly due to repetitive character—the Company may charge a reasonable fee or refuse to act, provided it notifies the Data Subject within one (1) week with justification. All access requests are logged in the central DSR Register, and fulfillment is tracked against service-level objectives by jurisdiction.

11.2 Right to Rectification

Data Subjects have the right to obtain, without undue delay, the rectification of inaccurate Personal Data concerning them and, taking into account the purposes of the Processing, the right to have incomplete Personal Data completed. Requests for rectification must specify the data elements in question and, where possible, provide supporting evidence.

The Company shall verify the identity of the requester using multi-factor KYC protocols and validate the accuracy of the proposed correction against authoritative sources. If the disputed data originates from a third party (e.g., credit bureau, government registry), the Company shall inform the Data



Subject and assist in contacting the source. Rectification must be completed within fifteen (15) business days of verification, with changes propagated across all integrated systems and logs maintained showing the original value, new value, timestamp, and authorising agent. Where rectification is refused, the Data Subject must be notified within ten (10) business days with a clear explanation and recourse options.

11.3 Right to Erasure (“Right to Be Forgotten”)

Data Subjects have the right to obtain erasure of Personal Data without undue delay where:

- (a) the data is no longer necessary for the purposes for which it was collected;
- (b) consent is withdrawn and no other lawful basis exists;
- (c) the Data Subject objects to Processing under Clause 11.6 and no overriding legitimate grounds exist;
- (d) the data has been unlawfully processed; or
- (e) erasure is required to comply with a legal obligation. The Company shall honour such requests within thirty (30) days of verification, unless retention is mandated by law (e.g., AML/KYC records) or necessary for the establishment, exercise, or defence of legal claims.

In cases involving PCI DSS-regulated data, unmasked PANs may only be erased after the statutory retention period for transaction disputes has expired; until then, tokenisation or truncation serves as a proportionate alternative. Erasure includes deletion from primary databases, backups, caches, and analytics sandboxes, with cryptographic erasure used where physical deletion is impractical.

A certificate of erasure, signed by the Data Steward, is issued upon completion. Requests that would infringe on the rights of others (e.g., co-signers on joint accounts) or compromise fraud prevention systems are evaluated on a case-by-case basis by Legal and Compliance.

11.4 Right to Restriction of Processing

Data Subjects may request restriction of Processing where:

- (a) the accuracy of the data is contested, pending verification;
- (b) the Processing is unlawful but the Data Subject opposes erasure;
- (c) the Company no longer needs the data but the Data Subject requires it for legal claims; or
- (d) the Data Subject has objected to Processing pending verification of legitimate interests.

During restriction, the data may be stored but not otherwise Processed except with the Data Subject’s consent, for legal claims, or to protect the rights of another natural or legal person.

The Company shall implement technical controls to enforce restriction—such as flagging records in databases, disabling API endpoints, or isolating data in read-only partitions—and notify the Data



Subject before lifting any restriction. Fulfilment occurs within fifteen (15) business days of verification.

11.5 Right to Data Portability

Where Processing is based on consent or contractual necessity and carried out by automated means, Data Subjects have the right to receive their Personal Data in a structured, commonly used, and machine-readable format and to transmit that data to another controller without hindrance.

The Company provides portable data via secure download or direct API transfer, using open standards such as JSON or XML, and includes only data actively provided by the Data Subject or generated through their activity (e.g., transaction history, profile settings). Derived data (e.g., risk scores, inferred preferences) is excluded unless it is essential to service continuity. Portability requests are fulfilled within thirty (30) days, and the Company does not impose fees unless the request is manifestly excessive. This right supports user autonomy and market competition while respecting system integrity.

11.6 Right to Object

Data Subjects have the right to object, on grounds relating to their particular situation, to Processing based on legitimate interests or public task, including profiling. They also have an absolute right to object to direct marketing at any time.

Upon receipt of an objection, the Company shall cease the contested Processing unless it demonstrates compelling legitimate grounds that override the Data Subject's interests, rights, and freedoms, or if Processing is necessary for legal claims. Objections to marketing result in immediate suppression across all channels. All objections are reviewed by the Compliance function within ten (10) business days, and the Data Subject is notified of the outcome with appeal instructions.

11.7 Rights Related to Automated Decision-Making and Profiling

Data Subjects have the right not to be subject to a decision based solely on automated processing—including profiling—that produces legal effects or similarly significant effects. The Company prohibits fully automated decisions that deny financial services, terminate contracts, or assign risk scores without human review.

Where automated tools are used (e.g., fraud scoring, creditworthiness assessment), a qualified employee must validate the output before action is taken. Data Subjects may request meaningful information about the logic, significance, and envisaged consequences of such processing and may contest the decision through a dedicated appeals channel. These safeguards ensure that algorithms serve humans—not the reverse.

11.8 Verification, Fulfilment Timelines, and Recordkeeping

All DSR requests must be submitted through approved channels: online portal, email to privacy@orchestrate.global, toll-free number, or in-person at designated offices in Dallas, Abuja, or Kigali. Identity verification uses risk-based authentication: for low-risk requests (e.g., marketing opt-out), email confirmation suffices; for high-risk requests (e.g., erasure, access to financial data), multi-factor KYC with government ID and liveness check is required. Once verified, timelines commence as follows:



- GDPR jurisdictions: 30 days (extendable by 60 days for complexity)
- California (CCPA/CPRA): 45 days (with one 45-day extension)
- Nigeria (NDPR): 30 days
- Rwanda (Law No. 058/2021): 30 days

All requests, responses, extensions, and justifications are logged in the immutable DSR Register, retained for seven (7) years, and reported quarterly to the Board Risk Committee. Failure to meet timelines constitutes a reportable incident under Clause 23.

Clause 12. Data Accuracy and Integrity Controls

12.1 Self-Service Correction Mechanisms

Orchestrate Global Inc. recognises that Data Subjects are often the most reliable source of truth regarding their own Personal Data.

To this end, the Company provides secure, intuitive self-service portals accessible via web and mobile applications that enable authenticated users to view, verify, and correct a wide range of Personal Data fields in real time. These include, but are not limited to: name, contact details (email, phone, physical address), date of birth, occupation, and business registration information. Corrections are validated against format rules (e.g., email syntax, phone number length) and, where applicable, cross-referenced with authoritative external sources such as national identity databases or business registries through secure API integrations.

For fields that cannot be self-corrected due to regulatory or security constraints—such as government-issued identification numbers, biometric templates, or transaction records—the portal provides a structured request workflow that routes the submission to the appropriate Data Steward for manual review under Clause 11.2.

All self-service interfaces are localised into English, Kinyarwanda, and Hausa, with accessibility features compliant with WCAG 2.1 AA standards to ensure equitable access for persons with disabilities. The availability of these mechanisms is prominently disclosed in all privacy notices and onboarding flows, reinforcing the Company's commitment to collaborative data stewardship.

12.2 Correction Timelines and Audit Trails

All corrections to Personal Data—whether initiated through self-service or formal Data Subject Rights requests—must be processed without undue delay and in any event within fifteen (15) business days of verification.

Upon completion, the updated data must be synchronised across all integrated systems, including customer relationship management (CRM) platforms, payment processing engines, fraud detection models, and third-party service providers acting as sub-processors, to prevent propagation of outdated or conflicting information.

Every correction event is logged in an immutable audit trail that captures:

- (a) the unique identifier of the Data Subject;



- (b) the specific data field(s) modified;
- (c) the original value and new value;
- (d) the timestamp of the change;
- (e) the identity of the user or system initiating the change;
- (f) the method of initiation (self-service, agent-assisted, automated reconciliation); and
- (g) the verification method employed. These logs are classified as Confidential under Clause 6, stored in write-once-read-many (WORM) storage for a minimum of seven (7) years, and protected by cryptographic hashing to ensure tamper resistance. Access to correction logs is restricted to authorised personnel in Compliance, Internal Audit, and Security, and all queries are themselves logged for oversight.

This granular traceability enables forensic reconstruction in the event of disputes, supports regulatory audits, and deters unauthorized alterations.

12.3 Prevention of Inaccurate Data Propagation

To proactively safeguard data integrity, the Company implements preventive controls at multiple layers of the data lifecycle. At the point of ingestion, all data entry interfaces enforce strict input validation rules that reject malformed, incomplete, or implausible values—for example, future dates of birth, non-existent postal codes, or PANs failing Luhn algorithm checks.

During processing, automated reconciliation jobs compare data across authoritative systems to detect and flag discrepancies; for instance, if a customer's legal name in the KYC repository diverges from their profile name in the transaction system, an alert is generated for manual resolution. In analytics and machine learning pipelines, data quality metrics—including completeness, consistency, validity, and timeliness—are monitored in real time, and models are retrained only when quality thresholds are met.

Furthermore, the Company prohibits the use of stale or unverified data for high-impact decisions such as credit approval, fraud blocking, or account termination; such decisions require either real-time validation against source systems or explicit human override with documented justification. Data Owners are accountable for conducting quarterly data quality reviews within their domains, and persistent inaccuracies trigger root cause analysis and process improvement initiatives. These measures ensure that inaccuracy is not merely corrected after the fact but is systematically prevented at its source, thereby upholding the dignity and autonomy of every Data Subject whose information the Company stewards.

Clause 13. Data Retention and Secure Disposal

13.1 Retention Periods by Data Type

Orchestrate Global Inc. maintains a granular, risk-based data retention schedule that aligns with statutory mandates, contractual obligations, and legitimate business needs across all jurisdictions of operation. Retention periods are defined at the data element level and encoded into automated



lifecycle management policies within all systems that store Personal Data. The following minimum retention periods apply unless superseded by a longer legal requirement or active legal hold:

- KYC/AML Documentation (including government-issued ID, proof of address, beneficial ownership records): Seven (7) years from the date of termination of the customer or merchant relationship, in accordance with anti-money laundering regulations in the United States, Nigeria, and Rwanda.
- Transaction Logs (including payment authorisations, settlements, refunds, and chargeback records): Seven (7) years from the date of transaction, to satisfy PCI DSS Requirement 10.7, financial audit standards, and dispute resolution timelines under card network rules.
- Unmasked Primary Account Numbers (PANs) and Full Track Data: Retained only for the duration strictly necessary to complete authorisation, settlement, and reconciliation; thereafter, immediately tokenised or truncated. Full PANs may be retained in encrypted form within the Cardholder Data Environment (CDE) for up to eighteen (18) months solely for chargeback defence, after which they must be cryptographically erased.
- Employee Personnel Records (including contracts, performance reviews, payroll data): Duration of employment plus six (6) years, in compliance with U.S. federal and state labour laws and Nigerian Labour Act provisions.
- System Access Logs and Authentication Events: Three hundred sixty-five (365) days in immutable SIEM storage, as required by PCI DSS Requirement 10.7 and cybersecurity best practices.
- Customer Support Interactions (call recordings, chat transcripts, email correspondence): Three (3) years from the date of last interaction, unless relevant to an open dispute or regulatory inquiry.
- Marketing Consent and Preference Records: Duration of the customer relationship plus three (3) years, to support audit of consent validity and withdrawal history.
- Data Subject Rights Request Logs: Seven (7) years from the date of closure, to demonstrate compliance with GDPR Article 12(4), NDPR Section 2.1(2), and Rwanda Law No. 058/2021 Article 23.
- Backup Tapes and Snapshots Containing Personal Data: Retained only for the duration of the primary data's retention period; incremental backups are purged automatically upon expiry of the underlying dataset.

Retention periods are reviewed annually by the Compliance function and updated within thirty (30) days of any change in law, regulation, or material contract term. Data Owners must certify quarterly that no data under their purview exceeds its scheduled retention window.

13.2 Secure Disposal Methods

Upon expiry of the applicable retention period—and absent a valid legal hold—Personal Data must be disposed of using methods that ensure irrecoverability and prevent unauthorised reconstruction. The choice of disposal method is calibrated to the classification level of the data under Clause 6 and the medium on which it resides:

- Digital Data (Confidential or Restricted): Cryptographic erasure is the default method, achieved by destroying the encryption keys that protect the data at rest, rendering the ciphertext permanently unreadable. For systems not encrypted at rest, secure deletion utilities



compliant with NIST SP 800-88 Revision 1 (Clear and Purge standards) must be used, overwriting storage sectors multiple times to prevent forensic recovery. Cloud-based data must be deleted using provider-native secure delete APIs that guarantee immediate and irreversible removal from all replicas and backups.

- Physical Media (hard drives, SSDs, backup tapes): Must be degaussed, shredded, or physically destroyed using certified third-party vendors. A Chain of Custody log must accompany all media from decommissioning through destruction, and a Certificate of Destruction—bearing serial numbers, destruction date, method, and vendor attestation—must be filed in the Records Management Repository.
- Paper Records: Must be cross-cut shredded on-site or transported in locked containers to NAID-certified destruction facilities. Shredding logs must include box identifiers, destruction date, and witness signatures.
- Archival or Immutable Storage (e.g., WORM, blockchain ledgers): Where technical constraints prevent deletion, data must be rendered anonymous through irreversible pseudonymisation or redaction, such that re-identification is not reasonably possible using available technology, resources, and time.

All disposal activities are logged in the central Data Lifecycle Register, with metadata including data type, volume, disposal method, timestamp, responsible personnel, and evidence reference. Automated purge jobs are tested quarterly to validate efficacy.

13.3 Legal Hold Suspension Protocol

The automatic disposal of data is suspended immediately upon issuance of a legal hold by the Legal department in anticipation of litigation, regulatory investigation, audit, or formal dispute. Legal holds are issued in writing, specifying the data categories, custodians, systems, and timeframes covered.

The Records Management team implements technical controls to exclude held data from automated deletion workflows and notifies all relevant Data Stewards. Legal holds remain in effect until formally released by the General Counsel or designated deputy, with periodic reviews conducted every ninety (90) days to assess continued necessity. Premature disposal of data under legal hold constitutes a serious breach of policy and may result in spoliation sanctions; therefore, system-level safeguards prevent override of hold flags without dual approval from Legal and Compliance.

13.4 Certification of Physical and Cryptographic Destruction

To provide auditable assurance of compliance, the Company requires formal certification for all high-sensitivity disposal activities. For physical destruction, vendors must provide a signed Certificate of Destruction that complies with ISO/IEC 27001 Annex A.8.3 and includes photographic evidence of the destruction process where feasible.

For cryptographic erasure, the Key Management System (KMS) must generate an immutable event log confirming key deletion, including key identifier, deletion timestamp, authorising user, and cryptographic proof of zeroisation. These certificates and logs are stored in the GRC platform for a minimum of ten (10) years and made available to regulators, auditors, and customers upon verified request. This practice transforms disposal from an operational task into a demonstrable act of accountability.



Clause 14. Data Storage Limitations and Localization Requirements

14.1 Geographic Data Residency Rules

Orchestrate Global Inc. adheres to strict data residency requirements that reflect both legal mandates and ethical commitments to data sovereignty in the jurisdictions where it operates. Personal Data relating to natural persons residing in the Republic of Rwanda must be stored and processed exclusively within the territory of Rwanda, unless one of the following conditions is met:

- (a) the Data Subject has provided explicit, informed, and verifiable consent to a specific cross-border transfer;
- (b) the transfer is necessary for the performance or conclusion of a contract between the Data Subject and the Company; or
- (c) the transfer is required by Rwandan law or a court order issued under Law No. 058/2021.

Similarly, Personal Data of Nigerian citizens and residents must be hosted within Nigeria or in cloud regions designated as compliant with the Nigeria Data Protection Regulation (NDPR), and any export of such data outside Nigeria requires prior notification to the National Information Technology Development Agency (NITDA) and, where applicable, documented consent.

For customers in the United States, data may be stored in U.S.-based cloud regions (e.g., AWS us-east-1, Azure East US), but transfers to non-U.S. jurisdictions for processing are prohibited without a valid transfer mechanism under Clause 15. These residency rules apply to all categories of Personal Data, including metadata, logs, backups, and analytics derivatives, and extend to third-party subprocessors engaged by the Company. Violation of data residency obligations constitutes a critical policy breach and triggers immediate incident response protocols under Clause 23.

14.2 Prohibition on Unauthorized Cross-Region Replication

To enforce geographic boundaries, the Company prohibits all forms of unauthorised cross-region data replication, whether intentional or accidental. Database replication streams, content delivery networks (CDNs), backup pipelines, disaster recovery failover systems, and analytics exports must be explicitly configured to respect jurisdictional firewalls.

For example, customer records created through the Kigali onboarding portal must never replicate to Dallas-based read replicas, and Nigerian transaction logs must not appear in global log aggregation clusters hosted outside Africa. Network segmentation policies enforce this at the infrastructure layer: virtual private clouds (VPCs) in Rwanda (Oracle Cloud IAD) and Nigeria (Azure Nigeria Central) are isolated from global peering groups, and firewall rules default to deny-all for inter-regional traffic involving Personal Data.

Application code must not contain hard-coded region overrides or “global sync” features that bypass residency controls. All data movement workflows undergo architecture review by the Data Governance team, and automated configuration scanners flag any drift from approved topologies. Personnel who knowingly initiate unauthorized cross-border transfers—such as exporting a customer list to a personal cloud drive or configuring a test environment in a non-compliant region—are subject to disciplinary action up to and including termination.



14.3 Infrastructure Mapping by Jurisdiction

The Company maintains a real-time, authoritative map of its data infrastructure aligned to jurisdictional boundaries, published in the central Configuration Management Database (CMDB) and reviewed quarterly by Internal Audit. The current approved topology is as follows:

- United States (Dallas, TX):
Primary cloud provider: Amazon Web Services (AWS)
Region: us-east-1 (North Virginia)
Systems hosted: U.S. customer onboarding, merchant settlement, corporate HRIS, global executive dashboards
Data types: U.S. citizen/resident Personal Data, corporate records, aggregated analytics
- Rwanda (Kigali):
Primary cloud provider: Oracle Cloud Infrastructure (OCI)
Region: af-johannesburg-1 (with local edge caching in Kigali via partnership with RURA-certified data centre)
Systems hosted: Rwandan customer KYC, mobile wallet transactions, local fraud monitoring
Data types: Rwandan citizen/resident Personal Data, national ID scans, biometric templates, transaction logs
- Nigeria (Abuja):
Primary cloud provider: Microsoft Azure
Region: Nigeria Central (Lagos)
Systems hosted: Nigerian merchant onboarding, payment routing, AML screening
Data types: Nigerian citizen/resident Personal Data, BVN-linked records, business registration documents

All new system deployments must select a jurisdiction-aligned region during provisioning, and legacy systems migrating to cloud infrastructure must complete residency realignment within ninety (90) days. Any deviation requires a formal exception under Clause 27, supported by a Transfer Impact Assessment (TIA) and CSCO approval. This infrastructure discipline ensures that data sovereignty is not an afterthought but an architectural constant.

Clause 15. Cross-Border Data Transfers

15.1 Permitted Transfer Mechanisms

Orchestrate Global Inc. recognises that cross-border data transfers are sometimes necessary for legitimate business operations, regulatory compliance, or service delivery. However, such transfers are permitted only when accompanied by robust, legally valid safeguards that ensure an adequate level of protection for Personal Data equivalent to that guaranteed within the originating jurisdiction. The following mechanisms are exclusively authorised:

(a) Standard Contractual Clauses (SCCs): The European Commission's 2021 SCCs (as updated for Schrems II compliance) must be incorporated into all Data Processing Agreements (DPAs) with third-party processors or controllers located outside jurisdictions deemed "adequate" by the European Commission, the National Information Technology Development Agency (NITDA), or the Rwanda Utilities Regulatory Authority (RURA). These clauses impose binding obligations on the recipient



regarding purpose limitation, security, sub-processing restrictions, audit rights, and breach notification within twenty-four (24) hours.

(b) Explicit Informed Consent: Where no other mechanism applies, transfers may proceed if the Data Subject has provided explicit, granular, and revocable consent after being informed of the specific risks associated with the transfer, including the absence of enforceable data subject rights and the potential for government surveillance in the recipient country. Consent must be documented in the Consent Management Platform and re-validated annually.

(c) Contractual Necessity: Transfers strictly necessary for the performance of a contract between the Data Subject and the Company—or for pre-contractual measures taken at the Data Subject’s request—are permitted without additional safeguards, but only for the minimum data required and for the duration of the contractual relationship. Examples include transmitting shipping addresses to international logistics partners or sharing payment details with cross-border acquirers.

(d) Legal Obligation or Public Interest: Transfers required by law (e.g., court order, regulatory subpoena) or necessary to protect vital interests (e.g., public health emergency) are permitted, provided they are proportionate, documented, and limited to the specific legal mandate.

Transfers based on Binding Corporate Rules (BCRs) are under development and will be implemented upon approval by relevant supervisory authorities. Ad hoc contractual clauses, model agreements not approved by Legal, or reliance on “privacy shields” invalidated by courts are expressly prohibited.

15.2 Transfer Impact Assessments (TIAs)

Prior to initiating any cross-border transfer relying on SCCs or consent, a Transfer Impact Assessment (TIA) must be completed and approved by the Chief Security & Compliance Officer (CSCO). The TIA evaluates whether the laws and practices of the recipient country—particularly those concerning government access to data—undermine the effectiveness of the transfer safeguards.

The assessment must consider:

- (a) the specific categories of data being transferred;
- (b) the technical and organisational measures applied (e.g., end-to-end encryption, pseudonymisation);
- (c) the documented practices of public authorities in the recipient country regarding data access;
- (d) the availability of effective legal remedies for the Data Subject; and
- (e) any supplementary measures that could mitigate residual risks. TIAs must reference authoritative sources such as EDPB Recommendations 01/2020, U.S. Executive Order 14086, and national intelligence oversight reports. The completed TIA, along with the signed SCCs and evidence of supplementary measures, is filed in the GRC platform and reviewed annually or upon material change in the recipient country’s legal landscape.

15.3 Supplementary Measures and Encryption Requirements

Where a TIA identifies residual risks that cannot be eliminated through contractual clauses alone, the



Company must implement supplementary technical measures to restore an essentially equivalent level of protection. These include:

- **End-to-End Encryption:** Data must be encrypted on the Data Subject's device or at the point of collection using keys controlled solely by Orchestrate Global Inc. or the Data Subject, ensuring that the recipient processor never accesses plaintext data.
- **Pseudonymisation:** Direct identifiers must be replaced with tokens or aliases before transfer, with the re-identification key stored separately in a jurisdiction-compliant environment.
- **Split or Dual Processing:** Sensitive operations are divided across multiple jurisdictions so that no single entity possesses sufficient data to re-identify individuals.
- **Strict Access Controls:** Recipient systems must enforce zero-trust architecture, just-in-time privileged access, and real-time anomaly detection.

For Restricted data—including unmasked PANs, cryptographic keys, and biometric templates—cross-border transfer is prohibited unless end-to-end encryption with customer-held keys is implemented, and even then, only with CSCO and CEO approval. All supplementary measures are validated through penetration testing and documented in the TIA.

15.4 Prohibition on Transfers to High-Risk Jurisdictions Without Approval

Transfers of Personal Data to countries designated by the Company's Risk Committee as "high-risk"—based on factors such as absence of data protection laws, systemic government surveillance, weak rule of law, or history of data exfiltration—are categorically prohibited unless an exception is granted under Clause 27. The current high-risk list includes jurisdictions identified in the U.S. Department of Commerce's Country Commercial Guides as lacking adequate privacy frameworks, as well as any country subject to comprehensive sanctions by the United Nations, European Union, or U.S. Office of Foreign Assets Control (OFAC).

Even with SCCs or consent, transfers to these jurisdictions require:

- (a) a Level 9 risk score in the TIA;
- (b) external privacy counsel review;
- (c) Board Risk Committee notification; and
- (d) direct approval by the CEO and CSCO. Personnel who initiate unauthorised transfers to high-risk jurisdictions face immediate suspension and potential termination. This prohibition reflects the Company's commitment to treating data not as a commodity to be routed arbitrarily, but as a trust-bound asset whose movement must always serve the rights and safety of the individual.

Clause 16. Third-Party Data Processing and Vendor Management

16.1 Due Diligence Requirements

No third party may process Personal Data on behalf of Orchestrate Global Inc. without first undergoing a comprehensive due diligence assessment commensurate with the sensitivity of the data and the criticality of the service.



The assessment is conducted by the Procurement and Compliance functions using a standardised questionnaire that evaluates:

- (a) the vendor's data protection governance structure, including appointment of a Data Protection Officer or equivalent;
- (b) certification against recognised standards such as ISO/IEC 27001, ISO/IEC 27701, SOC 2 Type II, or PCI DSS (as applicable);
- (c) technical security controls, including encryption at rest and in transit, access management, logging, and incident response capabilities;
- (d) geographic data handling practices, including sub-processor locations and cross-border transfer mechanisms;
- (e) breach history over the past three (3) years, including root cause analysis and remediation evidence; and
- (f) financial stability and business continuity planning. For vendors processing Restricted data—including unmasked PANs, cryptographic keys, or biometric templates—on-site audits or independent penetration test reports are mandatory. Vendors that fail to meet minimum thresholds, refuse to provide evidence, or operate exclusively in high-risk jurisdictions as defined under Clause 15.4 are disqualified from engagement.

All due diligence records are retained in the Vendor Risk Register for the duration of the relationship plus seven (7) years.

16.2 Mandatory Contractual Clauses

All contracts with third-party processors must include a binding Data Processing Addendum (DPA) that incorporates, at a minimum, the following clauses:

- (a) **Processing Limitations:** The processor shall process Personal Data only on documented instructions from Orchestrate Global Inc., including with regard to transfers to third countries, and shall not process data for its own purposes.
- (b) **Confidentiality:** All personnel authorised to process Personal Data must be bound by confidentiality obligations, and background checks must be performed for those with access to Restricted data.
- (c) **Security Measures:** The processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including those specified in Clauses 17 through 22 of this Policy.
- (d) **Sub-Processing Restrictions:** The processor shall not engage another processor without prior specific or general written authorisation from Orchestrate Global Inc. Where general authorisation is granted, the processor must notify the Company of any intended changes concerning the addition or



replacement of subprocessors, and the Company reserves the right to object within ten (10) business days.

(e) Assistance with Data Subject Rights: The processor shall provide reasonable assistance to enable the Company to fulfil its obligations to respond to Data Subject Rights requests under Clause 11.

(f) Breach Notification: The processor shall notify the Company without undue delay—and in any event within twenty-four (24) hours—after becoming aware of a Personal Data Breach, providing full details of the nature, scope, and impact.

(g) Audit and Inspection Rights: The processor shall make available to the Company all information necessary to demonstrate compliance with these obligations and allow for and contribute to audits, including inspections, conducted by the Company or an auditor mandated by the Company.

(h) Return or Deletion: At the end of the provision of services, the processor shall, at the choice of Orchestrate Global Inc., delete or return all Personal Data and delete existing copies unless Union, Member State, or national law requires storage.

These clauses are non-negotiable for vendors processing confidential or restricted data. Any deviation requires CSCO approval and is logged as a high-risk exception.

16.3 Sub-Processor Restrictions

Orchestrate Global Inc. maintains a central registry of approved sub-processors, published internally and updated quarterly. Vendors must disclose their entire sub-processor chain during due diligence, and any new sub-processor must undergo the same risk assessment as the primary vendor.

Sub-processors handling Restricted data must be located only in jurisdictions compliant with the Company's data residency rules under Clause 14 and must themselves sign DPAs with equivalent obligations.

The Company reserves the right to conduct announced or unannounced audits of critical sub-processors, particularly those involved in payment processing, identity verification, or cloud infrastructure. Vendors that conceal sub-processor relationships or fail to enforce downstream compliance are subject to immediate contract termination.

16.4 Continuous Monitoring and Audit Rights

Third-party risk is not static; therefore, continuous monitoring is required throughout the lifecycle of the relationship. The Company employs automated vendor risk platforms that ingest external threat intelligence, security rating scores, and regulatory enforcement actions to flag emerging risks.

Critical vendors—those processing Restricted data or supporting core payment functions—are subject to annual on-site or remote audits, including review of access logs, configuration settings, and incident response playbooks. All vendors must certify compliance annually via a signed attestation, and failure to respond within thirty (30) days results in suspension of data flows.



Additionally, the Company exercises its contractual audit rights upon:

- (a) material change in vendor ownership or control;
- (b) public disclosure of a significant security incident;
- (c) downgrade in security certification status; or
- (d) initiation of regulatory enforcement action against the vendor. Audit findings are tracked to remediation in the GRC platform, and unresolved high-risk issues trigger contract review or termination. This proactive oversight ensures that third parties remain accountable stewards of the trust placed in them by Orchestrate Global Inc. and its Data Subjects.

Clause 17. Cryptographic Controls and Key Lifecycle Management

17.1 Encryption Standards for Data at Rest and in Transit

Orchestrate Global Inc. mandates the use of strong, industry-standard cryptographic algorithms to protect confidential and restricted data against unauthorized access, disclosure, and tampering. For data at rest, all storage systems—including databases, file shares, backup repositories, and portable media—must employ AES-256 in Galois/Counter Mode (GCM) or an equivalent authenticated encryption scheme that provides both confidentiality and integrity.

Legacy modes such as ECB, CBC without HMAC, or RSA encryption of bulk data are prohibited. Transparent Data Encryption (TDE) alone is insufficient for PCI DSS-regulated environments; application-layer or column-level encryption must be used for PANs and other Restricted data elements.

For data in transit, all communications involving Personal Data—whether over public networks, internal corporate networks, or between microservices—must be protected using Transport Layer Security (TLS) version 1.3 with Perfect Forward Secrecy (PFS), ephemeral key exchange (ECDHE), and certificate pinning. TLS 1.2 is permitted only where TLS 1.3 is technically infeasible, provided it uses strong cipher suites (e.g., TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384) and disables weak protocols (SSLv3, TLS 1.0, TLS 1.1).

Unencrypted protocols such as FTP, Telnet, HTTP, and SMTP without STARTTLS are categorically forbidden. All cryptographic implementations must undergo annual validation by a Qualified Security Assessor (QSA) and be subject to automated configuration scanning to detect drift from approved standards.

17.2 Key Generation within FIPS 140-2 Level 3 HSMs

Cryptographic keys used to protect Restricted data—including Data Encryption Keys (DEKs) and Key Encryption Keys (KEKs)—must be generated exclusively within FIPS 140-2 Level 3 validated Hardware Security Modules (HSMs) or cloud-native equivalents that provide equivalent assurance (e.g., AWS CloudHSM, Azure Dedicated HSM, Oracle Key Vault with FIPS 140-2 Level 3 certification).



These modules must be physically or logically isolated from application servers, with access restricted to authorised cryptographic officers via multi-factor authentication and split knowledge protocols.

Key generation must use cryptographically secure random number generators (CSPRNGs) compliant with NIST SP 800-90A, and output must be verified for entropy and unpredictability. Software-based key generation—even within trusted execution environments—is prohibited for Restricted data. All key generation events are logged with immutable metadata including timestamp, operator identity, key type, algorithm, length, and HSM serial number.

17.3 Key Storage, Distribution, and Usage Policies

The principle of hierarchical key management governs all cryptographic operations. DEKs, which encrypt the actual data, are themselves encrypted under KEKs, which are stored exclusively within HSMs and never exposed to application memory, disk, or network. Under no circumstances may KEKs be stored in configuration files, environment variables, source code repositories, or cloud secret managers lacking HSM backing.

Application systems retrieve DEKs from the HSM on-demand via secure APIs, and keys remain in volatile memory only for the duration of the cryptographic operation. Key distribution to third parties—for example, in hybrid cloud architectures—is prohibited unless mediated through HSM-to-HSM replication with mutual authentication and audit logging. Access to perform cryptographic operations (e.g., encrypt, decrypt, sign) is granted only to service accounts with least-privilege permissions, and all operations are logged with full context (user, data object, operation type, result). Session keys for TLS must be ephemeral and destroyed immediately upon connection termination.

17.4 Key Rotation Schedules and Emergency Procedures

Cryptographic keys must be rotated on a schedule calibrated to their sensitivity and usage volume to limit the impact of potential compromise. The following minimum rotation intervals apply:

- Data Encryption Keys (DEKs): Every ninety (90) days for static data; every thirty (30) days for high-volume transactional data.
- Key Encryption Keys (KEKs): Annually, or immediately upon personnel changes affecting cryptographic officer roles.
- TLS Private Keys: Annually, or upon certificate renewal.
- Root Certificate Authority (CA) Keys: Every three (3) years, with offline storage and multi-person activation.

Automated key rotation workflows must ensure that old keys remain available for decryption of legacy data until the underlying data is purged under Clause 13. In the event of a suspected or confirmed key compromise—such as unauthorised HSM access, malware detection on key management systems, or anomalous decryption requests—an emergency rotation protocol is activated.

This protocol includes:

- (a) immediate revocation of affected keys;
- (b) forensic isolation of impacted systems;



(c) re-encryption of all data protected by the compromised key; and

(d) notification to the Incident Response Team under Clause 23. Emergency rotations override standard change management windows and require CSCO approval.

17.5 Cryptographic Key Destruction and Certification

When keys reach the end of their lifecycle—and all data they protect has been securely disposed of under Clause 13—they must be cryptographically destroyed using methods that render recovery impossible. Within HSMs, this is achieved through secure zeroisation commands that overwrite key material with random data in accordance with FIPS 140-2 requirements.

For cloud KMS, the “delete key” API must be used with permanent deletion enabled, bypassing soft-delete retention periods. All destruction events generate an immutable log entry containing the key identifier, destruction timestamp, authorising user, and cryptographic proof of erasure (e.g., hash of zeroised memory block). A Certificate of Key Destruction, signed by the Cryptography Architect and attested by the HSM audit log, is filed in the GRC platform and retained for ten (10) years.

This final act of key lifecycle management ensures that even if archived ciphertext were exhumed, it would remain permanently unreadable—thereby fulfilling the Company’s enduring obligation to protect data long after its operational life has ended.

Clause 18. Access Control and Authentication

18.1 Unique User Identification and Password Complexity

All users granted access to systems that store, process, or transmit Personal Data must be assigned a unique user identifier that is never shared, reused, or recycled. Generic, group, or shared accounts are prohibited except in rare, justified cases (e.g., emergency break-glass accounts under Clause 18.5), and even then, only with compensating controls.

Passwords serving as primary authentication factors must comply with stringent complexity requirements: a minimum length of twelve (12) characters, inclusion of at least three of the following four character classes (uppercase letters, lowercase letters, digits, special symbols), prohibition of dictionary words, personal information, or sequential patterns (e.g., “12345”, “qwerty”), and prevention of reuse of any of the last four (4) passwords. Passwords must be stored using adaptive, salted cryptographic hashing algorithms such as Argon2id, scrypt, or PBKDF2 with a minimum iteration count of 100,000—never in plaintext, reversible encryption, or weak hashes like MD5 or SHA-1.

Systems must enforce password changes only upon compromise or expiration of the ninety (90)-day maximum lifetime; arbitrary periodic resets without evidence of risk are discouraged per NIST guidance but retained for PCI DSS compliance in the Cardholder Data Environment (CDE). All authentication attempts—successful or failed—are logged with source IP, timestamp, user ID, and outcome for forensic analysis.



18.2 Multi-Factor Authentication (MFA) Requirements

Multi-Factor Authentication (MFA) is mandatory for all access to Confidential or Restricted data, regardless of network location. MFA must combine at least two of the following factor types:

- (a) something you know (e.g., password, PIN);
- (b) something you have (e.g., FIDO2 security key, hardware token, authenticator app); or
- (c) something you are (e.g., biometric verification). SMS-based one-time passwords (OTPs) are permitted only as a temporary fallback during onboarding or device loss, and must be phased out within fourteen (14) days. For remote access—including VPN, cloud consoles, and SSH sessions—MFA is required without exception.

Within the CDE, MFA is enforced at every login boundary, including administrative interfaces, database query tools, and API gateways. Customer-facing applications offering access to Personal Data (e.g., merchant dashboards, customer portals) must provide MFA as an opt-in feature at minimum, and as mandatory for high-risk actions such as changing payment methods or viewing full PANs. MFA bypass mechanisms are prohibited unless pre-approved under a formal exception and protected by dual control.

18.3 Session Timeouts and Inactive Account Management

To mitigate the risk of unattended sessions, all interactive systems must enforce automatic session termination after a period of inactivity proportionate to the sensitivity of the data accessed. The following maximum timeout thresholds apply:

- Restricted data environments (e.g., CDE, HSM management consoles): Fifteen (15) minutes of inactivity.
- Confidential data systems (e.g., CRM, HRIS, KYC platforms): Thirty (30) minutes of inactivity.
- Internal-use systems with no Personal Data: Sixty (60) minutes of inactivity.



Upon timeout, users must re-authenticate with full credentials, including MFA where applicable. Additionally, user accounts that remain inactive for ninety (90) consecutive days must be automatically disabled, and after one hundred eighty (180) days, permanently deleted unless formally recertified by the user's manager and Data Owner.

Service accounts are exempt from inactivity timeouts but must undergo quarterly access reviews under Clause 18.4. Session tokens must be cryptographically random, bound to the user's IP address and user agent where feasible, and invalidated immediately upon logout or password change.

18.4 Role-Based Access Control (RBAC) and Quarterly Certifications

Access to systems and data is granted strictly on a need-to-know basis through a formal Role-Based Access Control (RBAC) model. Roles are defined by business function—not individual—and mapped to granular permissions that align with job responsibilities.

For example, a “Fraud Analyst” role may view transaction metadata but not full PANs, while a “Settlement Officer” may initiate fund transfers but not modify customer profiles. Role definitions are documented in the central Identity Governance system and reviewed annually by Data Owners.

All user-role assignments undergo automated quarterly access certifications: managers must attest to the continued necessity of each permission for their direct reports, and discrepancies trigger deprovisioning workflows. Privileged roles—such as System Administrator, Database Owner, or Security Auditor—are subject to additional scrutiny, including background checks, supplemental training, and dual approval for assignment. Access reviews are logged, and failure to complete certifications within ten (10) business days results in automatic revocation of access pending validation.

18.5 Privileged Access Management (PAM) Framework

Privileged accounts—those with elevated rights to configure systems, access sensitive data, or bypass standard controls—are managed under a strict Just-in-Time (JIT) and Just-Enough-Access (JEA) model. Standing privileged access is prohibited; instead, users request time-bound sessions (maximum four (4) hours) through a secure PAM platform that enforces MFA, records all activity via session playback, and requires business justification.



Break-glass accounts for emergency use are stored in dual-custody physical safes or digital vaults requiring two authorised personnel to unlock, and their use triggers immediate alerting to the Security Operations Center and mandatory post-incident review within twenty-four (24) hours.

All privileged commands—especially those involving data export, user creation, or firewall modification—are logged to immutable storage with full command-line detail and correlated to user identity. Annual penetration tests include attempts to escalate privileges, and findings are treated as critical vulnerabilities. This framework ensures that the most powerful access rights are treated not as entitlements, but as temporary, audited, and accountable privileges.

Clause 19. Logging, Monitoring, and Auditability

19.1 Mandatory Logging of Critical Events

Orchestrate Global Inc. requires the systematic capture of audit-relevant events across all systems that store, process, or transmit Personal Data, with particular emphasis on Confidential and Restricted data environments. The following event categories must be logged without exception:

- (a) **Authentication and Access Events:** All login attempts (successful and failed), logout events, password changes, MFA activations, session creations and terminations, and privilege escalations. Logs must include user identifier, source IP address, timestamp (in UTC), device fingerprint, and authentication method used.
- (b) **Data Access and Modification:** All queries, views, exports, creations, updates, deletions, and transfers of Personal Data. Logs must specify the data object (e.g., customer ID, transaction reference), fields accessed or altered, before-and-after values for modifications, and the business purpose or workflow triggering the action.
- (c) **System and Configuration Changes:** All modifications to security-critical system components, including firewall rules, access control lists, encryption settings, user role assignments, API permissions, and HSM policies. Logs must record the operator identity, change description, configuration snapshot before and after, and approval ticket reference.
- (d) **Privileged and Administrative Activity:** All commands executed by privileged users, including shell sessions, database queries, cloud console actions, and break-glass account usage. Full command-line input and output must be captured where technically feasible.
- (e) **Security and Anomaly Detection Alerts:** All alerts generated by intrusion detection systems (IDS), endpoint detection and response (EDR) platforms, SIEM correlation rules, and user behaviour analytics (UBA) engines, including severity level, affected assets, and automated response actions taken.



(f) Data Subject Rights Fulfilment: All steps in the execution of Data Subject Rights requests under Clause 11, including verification actions, data retrieval, redaction, erasure, and communication timestamps.

These logs are classified as Confidential under Clause 6 and must be generated in real time—no later than five (5) minutes after the event occurs—to support timely incident detection and response.

19.2 Immutable Log Storage and Retention

All audit logs must be transmitted securely to a centralised Security Information and Event Management (SIEM) platform or equivalent immutable storage system that prevents alteration, deletion, or backdating of records. Logs are written to Write-Once-Read-Many (WORM) storage or blockchain-backed ledgers that enforce cryptographic chaining of log entries, ensuring that any attempt to modify historical data breaks the chain and is immediately detectable. The following retention periods apply:

- General audit logs (authentication, access, configuration): Three hundred sixty-five (365) days, in compliance with PCI DSS v4.0 Requirement 10.7 and cybersecurity best practices.
- Logs involving Restricted data (e.g., PAN access, HSM operations): Seven (7) years, to satisfy financial regulatory requirements and long-term forensic needs.
- Incident-related logs: Retained indefinitely until the associated incident is fully resolved and all legal, regulatory, and audit obligations are discharged.

Log storage systems must implement strict capacity planning to prevent rollover or loss due to disk exhaustion, and daily integrity checks must validate the continuity and cryptographic consistency of the log chain. Backups of log data are encrypted and stored in geographically separate locations, but restoration must preserve immutability and chronological order.

19.3 Cryptographic Hashing and Transparency Logs

To provide independent verifiability of log integrity, the Company implements a cryptographic transparency mechanism. Each day at midnight UTC, the SIEM platform computes a SHA-384 hash of all log entries ingested during the preceding 24-hour period. This daily hash is then published to a public, append-only transparency log hosted on a permissioned blockchain or distributed ledger accessible to regulators, auditors, and designated third parties.

Any discrepancy between the locally stored logs and the published hash constitutes evidence of tampering and triggers an immediate investigation by the Security and Internal Audit functions. Additionally, critical events—such as PAN access or privileged command execution—are hashed individually at creation time, and these hashes are stored alongside the event record to enable granular integrity validation during forensic reviews. This layered approach ensures that logs are not merely records of activity but provably trustworthy evidence.

19.4 Access Restrictions and Dual Control for Log Modification

Access to audit logs is strictly controlled and segregated from operational roles to prevent conflicts of interest. Only personnel in the Security Operations Center (SOC), Internal Audit, and designated Compliance investigators may query or analyse logs, and even then, only through read-only interfaces that prohibit export, filtering, or annotation without explicit approval. Requests to access logs for



investigative purposes must be submitted via a formal ticketing system, specifying the scope, justification, and duration, and approved by a manager who is not involved in the underlying system or incident. Under no circumstances may system administrators, developers, or business users modify, delete, or suppress log data.

The SIEM platform enforces dual control for any administrative action that could affect log integrity—such as changing retention policies, disabling log sources, or restoring from backup—requiring approval from both the Chief Security Officer and the Head of Internal Audit. All queries against the log repository are themselves logged, creating a meta-audit trail that deters unauthorized surveillance or cover-up attempts. This rigorous access model ensures that audit logs remain an objective, incorruptible source of truth for accountability, incident response, and regulatory assurance.

Clause 20. Endpoint and Device Security

20.1 Full-Disk Encryption and Screen Lock Policies

All corporate-owned endpoints—including laptops, desktops, tablets, and mobile phones—must enforce full-disk encryption (FDE) using FIPS 140-2 validated cryptographic modules or equivalent industry-standard solutions such as BitLocker (Windows), FileVault (macOS), or dm-crypt (Linux). Encryption must be enabled by default at provisioning and protected by a pre-boot authentication mechanism that requires a password, PIN, or hardware token before the operating system loads.

Decryption keys must never be stored on the same device in plaintext and must be escrowed in a centralised key management system accessible only to authorised IT personnel under dual control. Additionally, all devices must enforce automatic screen locking after a maximum of five (5) minutes of inactivity, requiring re-authentication via password, biometric verification, or smart card to resume use. These controls apply regardless of whether the device is connected to corporate networks, operating remotely, or in sleep mode.

Devices that fail to comply—due to disabled encryption, weak passwords, or excessive lock timeouts—are automatically quarantined from network access until remediated. This ensures that lost or stolen devices cannot be used to extract Confidential or Restricted data, thereby upholding the Company's obligation to protect data at rest even in physical compromise scenarios.

20.2 Portable Media Encryption and Wipe Rules

All portable storage media—including USB drives, external hard disks, SD cards, and optical media—used to store or transfer Personal Data must be encrypted using AES-256 or equivalent strong cryptography, with encryption managed through centrally administered tools that enforce password complexity and prevent bypass.

Unencrypted portable media is prohibited for storing Confidential or Restricted data; any such use constitutes a reportable security incident under Clause 23. Encrypted media must be registered in the Asset Management System with custodian assignment, and usage logs must track connection events, file transfers, and user activity. Portable media that remains inactive for thirty (30) consecutive days must be automatically wiped of all data, and upon decommissioning, must undergo secure cryptographic erasure or physical destruction in accordance with Clause 13.4. Cloud-based file



sharing services (e.g., OneDrive, Google Drive) are permitted only when configured with end-to-end encryption, access logging, and data loss prevention (DLP) policies that block unauthorized sharing of sensitive files. Employees are prohibited from using personal cloud accounts or consumer-grade file transfer tools for business data under any circumstances.

20.3 BYOD Enrollment and Containerization Requirements

The Company permits limited Bring Your Own Device (BYOD) usage for email, calendar, and approved collaboration applications, provided that the device is enrolled in the Mobile Device Management (MDM) platform and complies with strict security policies. Enrollment requires acceptance of a BYOD agreement that grants the Company the right to remotely wipe corporate data—and only corporate data—in the event of loss, theft, or termination of employment.

All corporate applications and data must be isolated within a secure container that enforces encryption, prevents copying to personal apps, blocks screenshot functionality for sensitive screens, and disables jailbroken or rooted devices. The MDM platform enforces minimum OS versions, patch levels, and passcode requirements, and automatically revokes access if the device falls out of compliance.

Personal data outside the container remains under the user's control and is not accessed by the Company. However, devices that process Restricted data—including unmasked PANs or cryptographic keys—are categorically excluded from BYOD; such activities must occur only on corporate-owned, fully managed endpoints. This balanced approach enables workforce flexibility while maintaining a clear boundary between personal privacy and corporate data stewardship.

20.4 Endpoint Detection and Response (EDR) Deployment

All corporate endpoints and BYOD devices with access to Confidential or Restricted data must have Endpoint Detection and Response (EDR) agents installed and actively reporting to the Security Operations Center (SOC). These agents provide continuous monitoring for malicious activity, including malware execution, ransomware behaviour, suspicious process injection, lateral movement attempts, and unauthorised data exfiltration.

EDR solutions must support real-time threat intelligence feeds, behavioural analytics, memory forensics, and automated response capabilities such as process termination, network isolation, and file quarantine. Agents must be configured to prevent uninstallation or tampering without administrative override, and integrity checks must run hourly to detect rootkit interference. EDR telemetry is retained for a minimum of three hundred sixty-five (365) days and integrated with the SIEM platform for cross-correlation with network and log data.

Devices that disable, uninstall, or interfere with EDR agents are immediately blocked from accessing corporate resources until restored to compliance. This proactive defence layer ensures that endpoints—often the initial vector in cyberattacks—are not passive targets but active participants in the Company's security posture.



Clause 21. Network and Infrastructure Security

21.1 Zero Trust Network Architecture (ZTNA)

Orchestrate Global Inc. implements a Zero Trust Network Architecture (ZTNA) predicated on the principle of “never trust, always verify.” Traditional perimeter-based defences are insufficient in a cloud-native, remote-work environment; therefore, all access requests—whether originating from inside or outside the corporate network—are treated as untrusted until explicitly authenticated, authorised, and encrypted.

Access is granted on a least-privilege, just-in-time basis to specific applications or services, not broad network segments. Identity is the new perimeter: every user, device, and service must present verifiable credentials before any connection is established.

The Company employs a cloud-delivered ZTNA platform that enforces strict identity-aware proxying, continuous session validation, and adaptive access policies based on real-time risk signals such as device health, location anomalies, and behavioural baselines. Legacy VPNs are deprecated in favour of this application-centric model, which eliminates implicit trust and reduces the attack surface by orders of magnitude. All ZTNA policies are reviewed quarterly and logged for audit, ensuring that network access remains tightly coupled to business needs and user context.

21.2 Default-Deny Firewall Rules and Microsegmentation

All network traffic—both ingress and egress—is governed by stateful firewall rules that default to “deny all” and permit only explicitly authorised flows. Firewall rulebases are maintained in a centralised configuration management system, version-controlled, and subject to change control workflows that require peer review and CISO approval for modifications. Rules must specify source and destination IP ranges, ports, protocols, and business justification; overly permissive rules (e.g., “any/any”) are prohibited.

Within data centres and cloud environments, microsegmentation is enforced at the workload level using software-defined networking (SDN) or next-generation firewalls to isolate systems based on function and sensitivity. For example, web application servers are segmented from database servers, and development environments are isolated from production. Critical systems—including those processing Restricted data—are placed in dedicated security zones with no direct internet exposure and only tightly controlled, encrypted pathways to upstream/downstream components. Network flow logs are retained for 365 days and analysed for anomalous patterns indicative of lateral movement or data exfiltration.

This granular segmentation ensures that a breach in one component does not cascade across the infrastructure.

21.3 Isolation of Cardholder Data Environment (CDE)

The Cardholder Data Environment (CDE)—defined as any system component that stores, processes, or transmits cardholder data or sensitive authentication data—must be logically and physically isolated from all other networks and systems. This isolation is achieved through dedicated VLANs, separate virtual private clouds (VPCs), or on-premise air-gapped networks, with no shared services (e.g., DNS, DHCP, logging) unless mediated through secure, audited gateways. Systems connected to



the CDE are in scope for PCI DSS and must undergo annual scoping validation by a Qualified Security Assessor (QSA).

Wireless networks are prohibited within the CDE, and remote access is allowed only via MFA-protected jump hosts with session recording. All traffic entering or leaving the CDE is inspected by an intrusion prevention system (IPS) and decrypted for deep packet inspection where TLS termination is performed. Quarterly internal vulnerability scans and biannual external penetration tests specifically target CDE boundaries to validate segmentation effectiveness.

This rigorous isolation ensures that payment card data remains protected even if adjacent systems are compromised.

21.4 Protocol Restrictions and TLS Enforcement

To minimise attack surface and prevent exploitation of legacy vulnerabilities, the Company enforces strict protocol restrictions across all network communications. The following protocols are categorically prohibited due to inherent security weaknesses: FTP, Telnet, rsh, rlogin, HTTP (unencrypted), SMTP without STARTTLS, SNMPv1/v2c, and SMBv1.

All administrative access must occur over SSHv2 or equivalent encrypted channels, with key-based authentication preferred over passwords. For web traffic, HTTP Strict Transport Security (HSTS) is enforced with a minimum max-age of one year, and all public-facing services must support TLS 1.3 with Perfect Forward Secrecy. Internal service-to-service communication uses mutual TLS (mTLS) with short-lived certificates issued by a private Certificate Authority, ensuring that only trusted workloads can communicate.

Certificate pinning is implemented for mobile applications and critical APIs to prevent man-in-the-middle attacks. Network devices—routers, switches, firewalls—must disable unused services (e.g., SNMP, HTTP management) and enforce strong administrative credentials. These measures collectively eliminate entire classes of network-based attacks and ensure that data in motion remains confidential and intact across all transmission paths.

Clause 22. Application and API Security

22.1 Secure Software Development Lifecycle (SDLC)

Orchestrate Global Inc. embeds security into every phase of the software development lifecycle through a formal Secure SDLC framework that integrates people, processes, and technology to prevent, detect, and remediate vulnerabilities before they reach production. The framework applies to all internally developed applications, APIs, microservices, and infrastructure-as-code (IaC) templates that process Personal Data or interface with critical systems. It consists of the following mandatory phases:

(a) **Requirements & Design:** Security and privacy requirements are defined alongside functional specifications, including data classification, encryption needs, access control models, logging obligations, and compliance constraints (e.g., PCI DSS, GDPR). All new features involving Personal Data must undergo a preliminary Privacy by Design assessment.



(b) Threat Modelling: Prior to coding, engineering teams conduct structured threat modelling using the STRIDE methodology (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) to identify potential attack vectors, trust boundaries, and mitigating controls. Threat models are documented, reviewed by Security Architecture, and stored in the central repository.

(c) Secure Coding: Developers adhere to language-specific secure coding standards (e.g., CERT C, OWASP Java, Microsoft .NET Guidelines) that prohibit known anti-patterns such as SQL injection, cross-site scripting (XSS), insecure deserialization, and hardcoded credentials. Code reviews must include explicit validation of security controls, and peer reviewers are trained to spot common vulnerabilities.

(d) Automated Testing: Every code commit triggers automated security scans in the CI/CD pipeline, including Static Application Security Testing (SAST), Software Composition Analysis (SCA) for open-source dependencies, and IaC scanning for cloud misconfigurations. Findings above a defined severity threshold block merge or deployment until remediated.

(e) Dynamic & Interactive Testing: Pre-production environments undergo Dynamic Application Security Testing (DAST) and Interactive Application Security Testing (IAST) to validate runtime behaviour, authentication flows, session management, and API contract enforcement.

(f) Production Deployment & Monitoring: Applications are deployed with runtime protections such as Web Application Firewalls (WAFs), rate limiting, and anomaly detection. Post-deployment, security telemetry is correlated with SIEM alerts to detect exploitation attempts.

All SDLC activities are logged, and evidence is retained in the GRC platform for audit. Projects that bypass any phase require CSCO approval and are treated as high-risk exceptions.

22.2 Threat Modeling and STRIDE Analysis

Threat modelling is not a one-time exercise but an iterative discipline applied at feature inception, major refactoring, and after significant incident learnings. Engineering teams use standardised templates to map data flows, identify assets, define trust boundaries, and enumerate threats per the STRIDE model. For example, an API accepting KYC documents would be assessed for:

- Spoofing: Can an attacker impersonate a legitimate user? → Mitigation: MFA, OAuth 2.0 with PKCE.
- Tampering: Can documents be altered in transit? → Mitigation: TLS 1.3, request signing.
- Repudiation: Can a user deny submitting data? → Mitigation: Immutable audit logs with digital signatures.
- Information Disclosure: Could documents leak via error messages? → Mitigation: Generic error handling, data masking.
- Denial of Service: Could large uploads exhaust resources? → Mitigation: File size limits, queuing.
- Elevation of Privilege: Could a user access another's documents? → Mitigation: RBAC, tenant isolation.



Threat models are reviewed by the Security Architecture team and updated whenever the system design changes. High-risk threats trigger mandatory control implementation; medium risks require compensating controls or acceptance; low risks are monitored. This structured approach ensures that security is designed in—not bolted on.

22.3 Automated Code Scanning (SAST/DAST/SCA)

The Company mandates the use of industry-leading automated scanning tools integrated directly into the developer workflow:

- **Static Application Security Testing (SAST):** Tools such as Checkmarx or SonarQube analyse source code for security flaws without executing it. Scans run on every pull request and must achieve zero critical/high findings before merge. Rulesets are customised to align with OWASP Top 10, CWE Top 25, and PCI DSS Requirement 6.3.
- **Software Composition Analysis (SCA):** Tools like Snyk or Black Duck inventory open-source dependencies, flag known vulnerabilities (CVEs), and enforce license compliance. Any dependency with a CVSS score ≥ 7.0 or active exploit in the wild blocks deployment. Patching SLAs are enforced: critical within 48 hours, high within 7 days.
- **Infrastructure-as-Code (IaC) Scanning:** Tools such as Checkov or tfsec validate Terraform, CloudFormation, and Kubernetes manifests against security benchmarks (e.g., CIS AWS Foundations) to prevent misconfigurations like public S3 buckets or unencrypted databases.
- **Dynamic Application Security Testing (DAST):** Tools like OWASP ZAP or Burp Suite perform black-box testing against running applications to detect runtime issues such as broken authentication, IDOR, and server-side request forgery (SSRF). Scans run weekly in staging and after every major release.

All scan results are triaged by the AppSec team, tracked to resolution in the vulnerability management system, and reported to engineering leadership. False positives must be formally documented and approved by Security.

22.4 Penetration Testing and Vulnerability Remediation

In addition to automated testing, the Company conducts comprehensive manual penetration tests to uncover complex logic flaws, business logic abuse, and chained attack scenarios that tools miss. External pen tests are performed annually by PCI-qualified Approved Scanning Vendors (ASVs) for internet-facing systems and biannually for internal applications. Internal red team exercises simulate advanced persistent threats to test detection and response capabilities. All tests follow PTES (Penetration Testing Execution Standard) methodology and produce detailed reports with risk ratings, proof-of-concept exploits, and remediation guidance.

Vulnerabilities are categorised by severity:

- **Critical (CVSS ≥ 9.0):** Exploitable remotely with no user interaction (e.g., RCE, SQLi in auth flow). Remediation required within 48 hours.
- **High (CVSS 7.0–8.9):** Significant impact but requires some conditions (e.g., XSS leading to account takeover). Remediation within 7 days.
- **Medium (CVSS 4.0–6.9):** Limited impact or complex exploitation. Remediation within 30 days.



- Low (< 4.0): Informational or cosmetic. Remediation encouraged but not mandated.

Unremediated critical/high vulnerabilities block production deployment and trigger executive escalation. Post-remediation, validation scans confirm closure. This rigorous testing regime ensures that applications are not only functional but resilient against real-world adversaries.

Clause 23. Incident Response and Breach Notification

23.1 Detection, Containment, and Investigation Timelines

All personnel of Orchestrate Global Inc. are obligated to report any suspected or confirmed Personal Data Breach to the Chief Security & Compliance Officer (CSCO) or designated delegate within one (1) hour of discovery, via the 24/7 incident hotline, secure web portal, or encrypted email channel. Upon receipt, the CSCO shall activate the Incident Response Team (IRT)—comprising representatives from Security, Legal, Engineering, Communications, and Customer Success—within two (2) hours. The IRT shall execute initial containment measures within twenty-four (24) hours to prevent further data exfiltration, system compromise, or service disruption.

These measures may include network segmentation, account disablement, API key revocation, or temporary service suspension. A formal investigation shall commence immediately thereafter, employing digital forensics, log correlation, and threat intelligence to determine the scope, root cause, data elements affected, number of Data Subjects impacted, and potential harm. The investigation must produce a preliminary report within seventy-two (72) hours and a final root cause analysis within fourteen (14) calendar days of incident closure. All actions, decisions, and communications are logged in the central Incident Register, which serves as the single source of truth for regulatory reporting and internal review.

23.2 Regulatory Notification Windows by Jurisdiction

The Company adheres strictly to jurisdiction-specific breach notification timelines, calculated from the moment of “awareness”—defined as the point at which there is a reasonable degree of certainty that a breach has occurred, not merely suspicion. The following deadlines apply:

- European Union (GDPR): Notification to the relevant Supervisory Authority within seventy-two (72) hours of awareness, unless the breach is unlikely to result in a risk to the rights and freedoms of natural persons. If delayed, a reasoned justification must accompany the submission.
- California (CCPA/CPRA): Notification to the California Attorney General within forty-five (45) days of discovery if the breach affects more than 500 California residents and involves sensitive Personal Information as defined under Cal. Civ. Code § 1798.82.
- Nigeria (NDPR): Notification to the National Information Technology Development Agency (NITDA) within seventy-two (72) hours of becoming aware of a breach that is likely to result in a real risk of serious harm to individuals.
- Rwanda (Law No. 058/2021): Notification to the Rwanda Utilities Regulatory Authority (RURA) within seventy-two (72) hours of confirmation of a breach involving Personal Data.
- Payment Card Industry (PCI DSS v4.0): Immediate notification to the acquiring bank and payment brands upon confirmation of a suspected or confirmed compromise of cardholder data, followed by submission of a detailed Incident Report within five (5) business days.



Notifications must include:

- (a) a description of the nature of the breach, including categories and approximate number of Data Subjects and records concerned;
- (b) the name and contact details of the CSCO or other contact point;
- (c) a description of the likely consequences of the breach; and
- (d) a description of the measures taken or proposed to address the breach, including mitigation of its possible adverse effects. Where information is not immediately available, it may be provided in phases without undue further delay.

23.3 Customer Notification Content and Channels

Where a Personal Data Breach is likely to result in a high risk to the rights and freedoms of Data Subjects—such as identity theft, financial loss, discrimination, or reputational damage—the Company shall communicate directly with affected individuals without undue delay. Notifications must be written in clear, plain language, avoiding technical jargon, and delivered through secure, verifiable channels such as registered email, SMS, postal mail, or in-app alerts. The content must include:

- A description of the breach in non-technical terms;
- The categories of Personal Data involved (e.g., “name, email, and partial payment card number”);
- Contact information for the Company’s privacy team and relevant supervisory authority;
- A description of the likely consequences;
- Mitigation steps the individual can take (e.g., password reset, credit monitoring);
- Measures the Company has implemented to address the breach.

For African customers, notifications must be localised into local languages e.g Kinyarwanda (Rwanda) or Hausa/English (Nigeria), and for vulnerable populations—including children or elderly users—additional support channels (e.g., toll-free helplines, in-person assistance in Abuja or Kigali) must be provided.

Customer notifications must be issued within five (5) business days of breach confirmation, unless law enforcement requests delay for investigative purposes—in which case, documentation of the request must be retained. All customer communications are pre-approved by Legal and reviewed for cultural appropriateness by regional compliance leads.

23.4 Breach Register and Root Cause Analysis

The Company maintains a central, immutable Breach Register that documents every reported incident—whether confirmed or false positive—along with its classification, timeline, response actions, regulatory filings, customer communications, and lessons learned. Each entry includes:

- Unique incident ID and date/time of discovery;
- Classification level (Level 1–4 per Clause 4.3);
- Systems and data types affected;
- Number of Data Subjects impacted;



- Root cause (e.g., misconfigured cloud bucket, phishing, insider threat);
- Containment and remediation steps;
- Regulatory and customer notification status;
- Post-incident review outcomes.

Within thirty (30) days of incident closure, the IRT conducts a blameless post-mortem to identify systemic gaps in people, process, or technology. Findings are translated into actionable improvement items—such as policy updates, control enhancements, or training modules—and assigned to owners with deadlines.

The Risk and Compliance Committee reviews all Level 3 and Level 4 breaches quarterly to ensure organisational learning. This continuous feedback loop transforms incidents from failures into opportunities for resilience, ensuring that trust, once tested, is rebuilt through transparency and action.

Clause 24. Data Protection Impact Assessments (DPIAs)

24.1 DPIA Triggers and Thresholds

A Data Protection Impact Assessment (DPIA) is mandatory prior to the initiation of any Processing activity that is likely to result in a high risk to the rights and freedoms of natural persons. The following activities are presumed to meet this threshold and require a DPIA without exception:

- (a) Systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects or similarly significant effects (e.g., credit scoring, fraud risk assignment, employment eligibility algorithms).
- (b) Large-scale processing of Special Categories of Personal Data as defined in Clause 1.1, including biometric data for identification, health information, financial vulnerability indicators, or data concerning children.
- (c) Systematic monitoring of publicly accessible areas on a large scale (e.g., facial recognition in physical payment kiosks, location tracking via mobile SDKs across millions of users).
- (d) Processing that involves novel technologies or innovative use of data where the potential impact on individuals is not fully understood (e.g., AI-driven behavioural analytics, cross-device fingerprinting, synthetic identity generation).
- (e) Cross-border transfers of Personal Data to jurisdictions not deemed adequate under applicable law, particularly where government surveillance laws may undermine data subject rights.
- (f) Any Processing activity involving Restricted data within the Cardholder Data Environment (CDE), including new tokenisation schemes, PAN storage architectures, or cryptographic key management models.

In addition to these presumptive triggers, a preliminary screening questionnaire must be completed for all new products, features, vendor integrations, or material changes to existing systems. If the screening indicates potential risk—such as collection of sensitive identifiers, sharing with third



parties, or use of opaque algorithms—a full DPIA is required. The Chief Security & Compliance Officer (CSCO) retains final authority to mandate a DPIA even in the absence of a listed trigger if emerging threats or stakeholder concerns warrant it.

24.2 Risk Scoring Methodology

Each DPIA employs a structured risk scoring matrix that evaluates both the likelihood and severity of potential harms to data subjects across multiple dimensions:

- Data Sensitivity: Public (1) → Internal (2) → Confidential (3) → Restricted (4)
- Scale of Processing: <1,000 subjects (1) → 1,000–100,000 (2) → >100,000 (3)
- Vulnerability of Data Subjects: General adult (1) → Elderly, low-literacy (2) → Children (3)
- Transparency & Control: Full user control (1) → Partial opt-out (2) → No meaningful choice (3)
- Potential Harm: Minor inconvenience (1) → Financial/reputational loss (2) → Discrimination, exclusion, or physical harm (3)

The composite risk score is calculated as the product of the highest severity factor and the sum of likelihood factors, yielding a value between 1 and 36. Scores are interpreted as follows:

- 1–5: Low risk – DPIA not required; standard controls suffice.
- 6–11: Medium risk – DPIA recommended; mitigation optional.
- 12–20: High risk – DPIA mandatory; mitigation required.
- 21–36: Critical risk – DPIA mandatory; external privacy counsel review; Board notification.

This quantitative approach ensures consistent, defensible decision-making while allowing for contextual nuance.

24.3 Mitigation Requirements by Risk Score

The DPIA must propose specific, measurable, and technically feasible mitigation measures proportional to the identified risk level:

- High Risk (12–20): Requires implementation of at least three of the following: data minimisation by design, purpose limitation enforcement, end-to-end encryption, human-in-the-loop for automated decisions, granular consent mechanisms, and independent algorithmic auditing. Approval requires sign-off from the CSCO, Legal, and the relevant Data Owner.
- Critical Risk (21–36): In addition to the above, mandates:
 - (a) engagement of external privacy counsel to validate compliance;
 - (b) pre-launch consultation with relevant supervisory authorities where required by law;
 - (c) executive summary briefing to the CEO and CCO;
 - (d) real-time monitoring dashboard for early anomaly detection; and



(e) sunset clause requiring re-assessment after twelve (12) months. No Processing may commence until all mitigations are implemented and verified.

All mitigations are tracked in the GRC platform with ownership, deadlines, and evidence requirements. Residual risk after mitigation must be formally accepted in writing by the Data Owner and CSCO before launch.

24.4 Pre-Launch Approval Workflow

The DPIA process follows a formal, gated workflow integrated into the product development lifecycle:

1. Initiation: Product or Engineering Lead submits a DPIA request via the GRC portal upon feature inception.
2. Assessment: Privacy Engineering conducts interviews, reviews architecture diagrams, and completes the risk scoring template (Appendix F).
3. Mitigation Design: Cross-functional team (Security, Legal, UX) designs controls and documents implementation plan.
4. Review: Draft DPIA is circulated to Data Owner, CTO, and CSCO for comment.
5. Approval: Final DPIA is submitted for electronic sign-off:
 - Medium Risk: CSCO
 - High Risk: CSCO + Legal + CTO
 - Critical Risk: CSCO + Legal + CTO + CEO
6. Implementation Verification: AppSec validates that mitigations are deployed in staging.
7. Launch Gate: Release is blocked in CI/CD pipeline until DPIA status = “Approved & Verified.”

Approved DPIAs are retained for the life of the Processing activity plus seven (7) years and made available to regulators upon request. This disciplined workflow ensures that privacy is not an afterthought but a gatekeeper to innovation—balancing agility with accountability, and ambition with respect for human dignity.

Clause 25. Training, Awareness, and Culture

25.1 Mandatory Annual Training Requirements

All personnel of Orchestrate Global Inc.—including employees, executives, contractors, interns, and third-party agents with access to Company systems or data—must complete mandatory data protection and security awareness training annually, within thirty (30) days of their start date for new joiners.

The core curriculum, developed and maintained by the Compliance and Security functions, covers foundational principles including: the definition and classification of Personal Data; lawful bases for Processing under Clause 9; individual rights under Clause 11; secure handling of Confidential and Restricted data; incident reporting obligations under Clause 23; and consequences of policy violations.



Training is delivered via an interactive e-learning platform that includes scenario-based assessments, knowledge checks, and multilingual support (English, Kinyarwanda, Hausa). Completion is tracked in real time through integration with the Human Resources Information System (HRIS), and non-completion after forty-five (45) days results in automatic suspension of system access until remediated. Records of completion—including user identity, course version, assessment score, and timestamp—are retained for seven (7) years to demonstrate regulatory compliance and individual accountability.

25.2 Role-Specific Curriculum and Validation

Beyond the baseline requirement, personnel in high-risk or high-privilege roles must complete supplemental, role-specific training modules validated by subject matter experts:

- Engineering and DevOps Staff: Secure coding practices (OWASP Top 10), cryptographic implementation, API security, infrastructure hardening, and privacy by design principles. Must pass a hands-on lab validating secure configuration of cloud environments.
- Customer Support and Call Centre Agents: Handling of sensitive customer data, social engineering red flags, verification protocols for Data Subject Rights requests, and secure screen-sharing practices. Must pass a simulated phishing and impersonation test.
- Product and Marketing Teams: Consent management, cookie compliance, data minimisation in feature design, and ethical use of analytics. Must submit a privacy review checklist for each new campaign or product launch.
- Finance and HR Personnel: Protection of payroll, tax, and employee health records; secure file transfer protocols; and recognition of business email compromise (BEC) scams. Must demonstrate correct use of encrypted email and document vaults.
- Executives and Data Owners: Governance responsibilities under Clause 4, DPIA oversight, breach escalation protocols, and board-level risk reporting. Must attend an annual in-person workshop led by the CSCO.

Role-specific training is updated quarterly to reflect emerging threats, regulatory changes, and internal incident learnings. Validation requires either a passing score on knowledge assessments or successful completion of practical exercises, with recertification mandated annually.

25.3 New Hire Onboarding and Completion Enforcement

Data protection training is embedded into the formal onboarding workflow for all new personnel. Before being granted access to any production system, new hires must:

- (a) Acknowledge receipt and understanding of this Policy and the Code of Conduct;
- (b) Complete the core data protection module;
- (c) Sign a supplemental confidentiality agreement if accessing Restricted data;
- (d) Pass a baseline security awareness quiz with a minimum score of 80%.

IT provisioning systems are integrated with the learning management platform via SCIM APIs, ensuring that system access is automatically granted only upon verified training completion. Managers receive weekly dashboards of their team's compliance status, and failure to onboard



personnel within the 30-day window is flagged as a control deficiency in executive reports. This “train-first” approach ensures that security and privacy are foundational expectations—not optional add-ons—for every member of the organisation.

25.4 Phishing Simulations and Data Protection Week

To reinforce vigilance beyond formal training, the Company conducts quarterly, company-wide phishing simulations using realistic, context-aware lures (e.g., fake payroll updates, urgent vendor invoices, or customer support tickets). Simulations are tailored to regional threat landscapes—Nigerian BEC variants, Rwandan smishing campaigns, U.S. credential harvesters—and results are anonymised for aggregate reporting while providing individual feedback to those who click. Personnel who repeatedly fail simulations are enrolled in remedial coaching and may face temporary access restrictions.

Additionally, the Company observes an annual Data Protection Week during the first week of October, featuring:

- Executive-led town halls on privacy strategy;
- Technical deep dives on encryption, anonymisation, and DPIAs;
- “Ask Me Anything” sessions with the CSCO and DPO-equivalent;
- Recognition of teams demonstrating privacy-by-design excellence;
- Updated policy previews and feedback forums.

This sustained, multi-channel approach transforms data protection from a compliance obligation into a shared cultural value—where every employee understands that they are not just a user of data, but a guardian of trust.

Clause 26. Policy Compliance, Monitoring, and Enforcement

26.1 Audit Protocols and File Review Schedules

Orchestrate Global Inc. maintains a rigorous internal audit programme to independently verify the design and operating effectiveness of data protection controls across all business units, systems, and third-party relationships. The Internal Audit function, reporting directly to the Audit Committee of the Board, conducts risk-based audits on a cyclical schedule:

- High-Risk Areas (e.g., Cardholder Data Environment, KYC processing, cross-border transfers): Annual audits.
- Medium-Risk Areas (e.g., HRIS, CRM, marketing platforms): Biennial audits.
- Low-Risk Areas (e.g., public websites, internal wikis): Triennial audits or event-driven.

Audits follow a standard methodology that includes:

- (a) review of policies, procedures, and training records;
- (b) inspection of system configurations, access logs, and encryption settings;
- (c) testing of control execution through transactional sampling; and



(d) interviews with process owners and end users. Special emphasis is placed on validating the seven core principles under Clause 8, the integrity of the Data Inventory under Clause 7, and the timeliness of Data Subject Rights fulfilment under Clause 11.

Findings are classified as Critical, Major, or Minor based on potential impact and likelihood, with remediation deadlines assigned accordingly (30, 60, or 90 days). All audit reports are submitted to the Risk and Compliance Committee and tracked to closure in the GRC platform. External audits—including PCI DSS assessments by Qualified Security Assessors (QSAs) and ISO 27001 certification audits—are coordinated by the CSCO and treated with equal rigor.

26.2 Evidence Repository and Metadata Standards

To demonstrate continuous compliance, the Company maintains a centralized Evidence Repository within the Governance, Risk, and Compliance (GRC) platform. This repository houses artefacts that substantiate the implementation and effectiveness of every control referenced in this Policy. Each evidence item must include the following metadata:

- Control ID (e.g., OG-DPP-001-C187);
- Control Description (linked to policy clause);
- Owner (individual accountable for implementation);
- Evidence Type (e.g., screenshot, log extract, certificate, policy document);
- Collection Date and Review Date;
- Automation Flag (Yes/No – indicating whether collection is API-driven);
- Validity Period (e.g., “Annual” for training records, “Continuous” for SIEM logs).

Examples of required evidence include: quarterly access certification reports, HSM key rotation logs, DPIA approval records, penetration test summaries, vendor due diligence questionnaires, and breach notification templates. Automated integrations pull evidence directly from source systems (e.g., Okta for MFA logs, AWS Config for encryption status) to minimise manual effort and ensure freshness. Evidence older than its validity period is flagged for renewal, and gaps trigger automatic alerts to control owners. This structured approach transforms compliance from a point-in-time exercise into a living, auditable state.

26.3 Violation Logging and Disciplinary Action

All suspected or confirmed violations of this Policy must be documented in the central Policy Violations Register, maintained by the Compliance function. Each entry includes:

- (a) description of the violation;
- (b) personnel or systems involved;
- (c) date of occurrence and discovery;
- (d) root cause (process gap, technical flaw, willful misconduct);
- (e) impact assessment (data exposed, regulatory risk, reputational harm); and
- (f) corrective and disciplinary actions taken. Violations are categorised by severity:



- Minor: Unintentional, low-impact deviations (e.g., late training completion, misclassified non-sensitive file). Addressed via coaching and retraining.
- Major: Repeated or systemic failures (e.g., unauthorised data sharing, bypassing MFA). Results in formal written warning, access revocation, and mandatory remediation.
- Critical: Willful, malicious, or high-harm acts (e.g., exfiltration of PANs, disabling logging, falsifying audit evidence). Leads to immediate suspension, termination of employment or contract, regulatory referral, and potential civil or criminal prosecution.

Disciplinary outcomes are determined in consultation with Human Resources and Legal, consistent with the Company's Code of Conduct and local labour laws in Wyoming, Nigeria, and Rwanda. All violations—regardless of severity—are reviewed quarterly by the Executive Committee to identify trends and drive systemic improvements.

This transparent, consistent enforcement ensures that policy adherence is not optional but an inseparable component of professional conduct at Orchestrate Global Inc.

Clause 27. Exceptions and Risk Acceptance

27.1 Exception Submission and Risk Assessment

Temporary deviations from the mandatory requirements of this Policy may be granted only under exceptional circumstances where strict compliance is demonstrably impractical, technically infeasible, or would result in disproportionate business impact, provided that compensating controls are implemented to maintain an equivalent level of protection. All exception requests must be submitted electronically via the Governance, Risk, and Compliance (GRC) platform using the standard Policy Exception Form, which requires the following elements:

- (a) Description of the Control: Precise citation of the policy clause, control ID, and requirement being deviated from.
- (b) Business Justification: Clear explanation of the operational, technical, or strategic necessity driving the request, including why compliance cannot be achieved within a reasonable timeframe.
- (c) Risk Assessment: Detailed analysis of the increased risk exposure, including likelihood and impact on data confidentiality, integrity, availability, and individual rights.
- (d) Compensating Controls: Specification of alternative measures that mitigate the residual risk to an acceptable level (e.g., enhanced monitoring, manual review, network segmentation, additional logging).
- (e) Proposed Duration: Maximum period for which the exception is required, not to exceed six (6) months without revalidation.
- (f) Stakeholder Endorsement: Pre-approval from the relevant Data Owner, System Owner, and Engineering Lead.

Requests lacking any of these elements will be automatically rejected by the GRC system. The Chief Security & Compliance Officer (CSCO) may require independent validation of the risk assessment by the Internal Audit or Security Architecture teams before proceeding to approval.



27.2 Approval Authorities by Risk Level

Exception approvals are tiered based on the severity of the residual risk after compensating controls are applied, as scored using the methodology in Clause 24.2:

- Low to Medium Risk (Score 1–11): Approved by the Chief Security & Compliance Officer (CSCO). Examples include temporary use of TLS 1.2 for legacy vendor integration with WAF inspection, or delayed staff training due to extended medical leave with supervised access restrictions.
- High Risk (Score 12–20): Requires joint approval by the CSCO and the Chief Executive Officer (CEO). Examples include short-term storage of unmasked PANs beyond standard retention for forensic investigation of a sophisticated fraud ring, with HSM encryption and dual-custody access.
- Critical Risk (Score 21–36): Requires approval by the CEO and formal notification to the Board Risk Committee. Such exceptions are exceedingly rare and permitted only when national security, public safety, or existential business continuity is at stake. External privacy counsel review is mandatory.

No exception may be approved retroactively unless it pertains to an emergency action taken under Clause 23.3 (incident containment), and even then, documentation must be submitted within seventy-two (72) hours. All approved exceptions are published in the central Exceptions Register with status, duration, and oversight conditions.

27.3 Maximum Duration and Prohibited Exceptions

All exceptions are granted for a fixed term not exceeding one hundred eighty (180) days. Extensions require resubmission of a new request with updated justification and risk reassessment; no exception may be extended beyond one (1) year without Board-level review. Automatic expiration occurs at term end, triggering immediate revocation of associated access or functionality unless compliance is restored.

Certain core principles and controls are non-derogable and admit no exceptions under any circumstances. These include:

- Processing of children's data without verifiable parental consent (Clause 10.4);
- Storage of unencrypted Restricted data at rest or in transit (Clause 17.1);
- Failure to report a Personal Data Breach within mandated timelines (Clause 23.2);
- Use of pre-ticked boxes or bundled consent (Clause 10.1);
- Retention of Personal Data beyond statutory or policy-defined periods without legal hold (Clause 13.1);
- Cross-border transfer of African citizen data to non-compliant jurisdictions without explicit consent or legal basis (Clauses 14.1 and 15.1).

Any attempt to circumvent these prohibitions—whether through misclassification, architectural obfuscation, or third-party delegation—is treated as a Critical violation under Clause 26.3 and subject to immediate disciplinary action. This absolute boundary ensures that the foundational ethics of data stewardship are never compromised by expediency or convenience.



Clause 28. Global and Local Alignment

28.1 Jurisdictional Addendum Process

While this Policy establishes a single, globally applicable baseline for data protection, Orchestrate Global Inc. recognises that certain jurisdictions impose additional, stricter, or materially different requirements that cannot be fully addressed by the core text alone.

In such cases, a Jurisdictional Addendum may be issued as a formal supplement to this Policy, approved by the Chief Security & Compliance Officer (CSCO) and Chief Executive Officer (CEO). An Addendum applies only to Processing activities involving Data Subjects resident in, or data regulated by, the specified jurisdiction, and must never dilute, contradict, or circumvent the core principles established herein. Each Addendum must:

- (a) Cite the specific legal provision necessitating the deviation or enhancement;
- (b) Define the scope of applicability (e.g., “all Nigerian citizens,” “California residents”);
- (c) Specify the additional or modified controls required;
- (d) Identify the responsible local compliance lead;
- (e) Include a sunset clause if the requirement is temporary.

Addenda are published in the GRC platform alongside the main Policy and distributed to all relevant personnel in the affected region. They are reviewed annually or upon material change in local law, and automatically retired when the underlying legal basis is repealed or harmonised with the global standard.

28.2 Conflict Resolution Rule (Stricter Law Prevails)

In the event of an actual or apparent conflict between the requirements of this Policy and a mandatory provision of local law, the following hierarchy shall govern:

- Step 1: If local law imposes a stricter obligation (e.g., shorter breach notification window, broader definition of sensitive data, stronger consent requirements), the local law prevails, and the stricter standard must be implemented.
- Step 2: If local law imposes a weaker or contradictory obligation (e.g., permitting indefinite retention, allowing covert surveillance, or requiring disclosure without due process), the Company shall comply only to the extent permitted by international human rights norms and its own ethical commitments, and shall seek legal remedies—including judicial review or regulatory consultation—to challenge unlawful mandates.
- Step 3: Where compliance with one jurisdiction’s law would violate another’s (e.g., U.S. subpoena demanding data stored in Rwanda), the Company shall engage external counsel to pursue lawful alternatives, such as mutual legal assistance treaties (MLATs), and shall notify affected Data Subjects unless prohibited by court order.

This rule ensures that the Company never uses “local law” as a pretext for lowering standards, but instead treats legal diversity as a call to uphold the highest bar of protection available.



28.3 Active Addenda for Rwanda, Nigeria, and California

The following Jurisdictional Addenda are currently in force and form an integral part of this Policy:

- **Rwanda Addendum (Effective 8 October 2025):**
Pursuant to Law No. 058/2021 of 13/10/2021 Relating to the Protection of Personal Data, all Personal Data of Rwandan citizens and residents must be stored and processed exclusively within the territory of Rwanda, unless explicit, informed consent is obtained or the transfer is necessary for contract performance. Data Controllers must register with the Rwanda Utilities Regulatory Authority (RURA) and appoint a local representative. Breach notifications to RURA must be submitted within seventy-two (72) hours and include Kinyarwanda-language summaries for affected individuals. Automated decision-making affecting legal rights is prohibited without prior RURA authorisation.
- **Nigeria Addendum (Effective 8 October 2025):**
In accordance with the Nigeria Data Protection Regulation (NDPR) issued by NITDA, all Processing of Nigerian citizens' data requires registration as a Data Controller, implementation of technical and organisational measures proportional to risk, and adherence to data localisation principles. Consent must be documented in writing or electronic form, and Data Subjects have the right to data portability and compensation for damages. Annual audit reports must be submitted to NITDA, and cross-border transfers require explicit notification and, where applicable, consent. The Addendum designates the Abuja office as the local point of contact for regulatory inquiries.
- **California Addendum (Effective 8 October 2025):**
Under the California Consumer Privacy Act (CCPA), as amended by the California Privacy Rights Act (CPRA), California residents have the right to opt out of the "sharing" of their Personal Information for cross-context behavioural advertising, to correct inaccurate data, and to limit the use of sensitive personal information. The Company maintains a "Do Not Sell or Share My Personal Information" link on its homepage, honours Global Privacy Control (GPC) signals, and provides a toll-free number for rights requests. Sensitive personal information includes government identifiers, financial account details, precise geolocation, and biometric data. Contractors processing California data must certify compliance annually.

These Addenda are not exhaustive; new ones will be issued as the Company expands into additional jurisdictions or as laws evolve. All personnel serving customers in these regions must complete jurisdiction-specific training modules under Clause 25.2.

Clause 29. Records Management Integration

29.1 Alignment with Enterprise Retention Policy

The data retention and disposal requirements established under Clause 13 are not standalone obligations but integral components of Orchestrate Global Inc.'s broader Enterprise Records Management Policy (OG-RMP-001).

All Personal Data, whether structured or unstructured, digital or physical, is classified as a corporate record upon creation and assigned a formal Retention Schedule Code that dictates its lifecycle from ingestion to disposition. These codes are derived from a master taxonomy approved by Legal,



Compliance, and Internal Audit, and map directly to statutory, regulatory, fiscal, and operational requirements across all jurisdictions of operation. For example:

- KYC Documentation is assigned code FIN-KYC-07, reflecting a seven-year retention period post-relationship termination under AML laws in the U.S., Nigeria, and Rwanda.
- Transaction Logs carry code PCI-TXN-07, aligning with PCI DSS v4.0 Requirement 10.7 and card network dispute resolution timelines.
- Employee Personnel Files are tagged HR-EMP-06, satisfying U.S. federal labour law and Nigerian Labour Act mandates.

Retention schedules are embedded into system metadata at the point of creation via automated policy engines, ensuring that disposal actions are triggered without manual intervention. The Records Management function maintains the authoritative register of all schedule codes and reviews them quarterly for legal updates. This integration prevents fragmentation between privacy and records governance, ensuring that “deletion” is not an ad hoc act but a scheduled, auditable event within the corporate information lifecycle.

29.2 Legal Hold Issuance and Execution Protocol

The automatic disposal of records under their assigned retention schedule is suspended immediately upon issuance of a Legal Hold Notice by the Legal department. Legal Holds are issued in writing when litigation, regulatory investigation, audit, or formal dispute is reasonably anticipated, and must specify:

- (a) The precise categories of records covered (e.g., “all communications with Merchant X between 1 January 2024 and 30 June 2025”);
- (b) The custodians or systems involved;
- (c) The triggering legal matter or authority;
- (d) The designated Legal Hold Manager responsible for oversight.

Upon receipt, the Records Management team implements technical controls to exclude affected records from automated deletion workflows across all repositories—including primary databases, backups, email archives, and cloud object stores. Custodians receive automated notifications confirming preservation status, and any attempt to delete held records triggers real-time alerts to Legal and Compliance.

Legal Holds remain in effect until formally released in writing by the General Counsel or designated deputy, with mandatory revalidation every ninety (90) days to assess continued necessity. Premature disposal of records under Legal Hold constitutes spoliation and may result in severe sanctions; therefore, system-level safeguards enforce immutable hold flags that cannot be overridden without dual approval from Legal and the CSCO.

This protocol ensures that the Company’s duty to preserve evidence never conflicts with its obligation to minimise data retention.



Clause 30. Metrics, Reporting, and Board Oversight

30.1 Key Performance Indicators (KPIs)

Orchestrate Global Inc. measures the effectiveness of its data protection programme through a defined set of Key Performance Indicators (KPIs) that reflect both operational execution and strategic risk posture. These KPIs are quantitative, time-bound, and tied to specific policy clauses, ensuring that compliance is not asserted but demonstrated. The core KPIs include:

- Data Subject Rights (DSR) On-Time Fulfilment Rate: Percentage of verified DSR requests fulfilled within statutory deadlines (GDPR: 30 days; CCPA: 45 days; NDPR/Rwanda: 30 days). Target: $\geq 95\%$. Threshold for escalation: $< 95\%$.
- Mandatory Training Completion Rate: Percentage of personnel who have completed annual data protection training within the required window. Target: 100%. Threshold: $< 95\%$ triggers HR intervention.
- Critical & High Vulnerability Remediation Timeliness: Percentage of critical vulnerabilities remediated within 48 hours and high vulnerabilities within 7 days. Target: 100%. Threshold: Any open critical finding after 72 hours = immediate RCC escalation.
- Third-Party Audit Compliance Rate: Percentage of critical vendors with current ISO 27001/SOC 2 or equivalent certification. Target: 100%. Threshold: $< 90\%$ = procurement freeze on new contracts.
- Personal Data Breach Frequency (Quarter-to-Date): Number of confirmed breaches involving Personal Data per quarter. Target: 0. Threshold: Any breach = mandatory executive briefing.
- DPIA Completion Rate for High-Risk Initiatives: Percentage of new high-risk projects with approved DPIAs prior to launch. Target: 100%. Threshold: $< 100\%$ = release gate block.
- Access Certification Completion Rate: Percentage of managers who complete quarterly access reviews on time. Target: 100%. Threshold: $< 95\%$ = automated deprovisioning of uncertified accounts.

Each KPI is owned by a designated role (e.g., CSCO for DSR, CTO for vulnerabilities), collected via automated integrations (GRC, SIEM, HRIS, Jira), and validated monthly by Internal Audit. Deviations from targets trigger predefined corrective workflows under Clause 26.

30.2 Executive and Board Reporting Cadence

Performance against KPIs is reported through a tiered governance structure to ensure appropriate visibility and accountability:

- Operational Level: Weekly dashboards are provided to Data Owners, Engineering Leads, and Security Analysts, showing real-time status of their domains.
- Management Level: Monthly reports are submitted to the Executive Committee (ExCo), summarising KPI performance, open risks, incident trends, and remediation progress. Reports include root cause analysis of misses and action plans.
- Governance Level: Quarterly slide decks are presented to the Board Risk and Compliance Committee, featuring:
 - RAG (Red-Amber-Green) status of all KPIs;
 - Heatmap of top data protection risks by likelihood and impact;
 - Summary of regulatory developments and their implications;



- Status of critical audit findings and exception approvals;
- Metrics on third-party risk and breach response efficacy.

All Board reports are attested to by the CEO and CSCO, confirming that the information is accurate and that resources are sufficient to maintain compliance. Ad hoc reporting is triggered by Level 3 or Level 4 incidents under Clause 4.3, material regulatory inquiries, or significant control failures.

30.3 Dashboard Design and RAG Status

Executive dashboards are designed to convey complex information with clarity and urgency. Each KPI is displayed with:

- Current Value vs. Target;
- Trend Arrow (improving, stable, deteriorating);
- RAG Indicator:
 - Green: On target or exceeding expectations;
 - Amber: Minor deviation; corrective action in progress;
 - Red: Significant failure; executive attention required.

Dashboards are accessible via the GRC platform with role-based views and drill-down capabilities to underlying evidence. Historical data is retained for trend analysis, and benchmarking against industry peers (e.g., FS-ISAC, PCI SSC metrics) is included where available. This visual, actionable reporting ensures that data protection is not a static compliance exercise but a dynamic element of corporate strategy—visible, measurable, and continuously improved.

Clause 31. Policy Review and Continuous Improvement

31.1 Annual and Event-Driven Review Triggers

The Enterprise Data Protection Policy is subject to formal review at least once every twelve (12) months, with the next scheduled review date fixed as 8 October 2026, as stated in the Control Summary. However, this cyclical review is supplemented by mandatory event-driven reviews triggered by any of the following material changes:

- (a) Enactment, amendment, or repeal of a data protection law, regulation, or binding guidance in any jurisdiction where the Company operates (e.g., new U.S. federal privacy law, updates to NDPR, RURA directives in Rwanda);
- (b) Significant technological transformation, such as migration to a new cloud provider, adoption of AI-driven decision systems, or decommissioning of legacy infrastructure;
- (c) Material findings from internal or external audits, including PCI DSS assessments, ISO 27001 surveillance audits, or regulator-mandated reviews;
- (d) Occurrence of a Level 3 or Level 4 Personal Data Breach under Clause 4.3, particularly if root cause analysis reveals policy gaps;
- (e) Changes in corporate structure, such as mergers, acquisitions, or entry into new markets that alter the risk profile or legal footprint;



(f) Formal feedback from Data Subjects, employees, or regulators indicating ambiguity, impracticality, or inconsistency in policy interpretation.

Upon any such trigger, the Chief Security & Compliance Officer (CSCO) shall initiate an expedited review within fifteen (15) business days, regardless of the annual cycle. This dual-review mechanism ensures that the Policy remains not only current but anticipatory—adapting proactively to an evolving threat and compliance landscape.

31.2 Feedback Channels and Stakeholder Input

Continuous improvement is driven by structured input from a diverse set of stakeholders, both internal and external. The following channels are formally recognised as sources of policy enhancement:

- **Internal Feedback:** Employees and contractors may submit suggestions, clarifications, or concerns via the GRC platform’s “Policy Feedback” module, anonymous ethics hotline, or quarterly town halls hosted by the CSCO. All submissions are acknowledged within five (5) business days and reviewed by the Data Protection Working Group.
- **Data Subject Input:** Privacy notices include a dedicated email address (privacy@orchestrate.global) for questions and suggestions regarding data handling practices. Recurring themes from Data Subject Rights requests, complaints, or survey responses are analysed quarterly for policy implications.
- **Regulatory Engagement:** The Legal and Compliance functions maintain active dialogue with supervisory authorities in Nigeria (NITDA), Rwanda (RURA), and relevant U.S. state agencies. Guidance received during consultations, pre-approval processes, or enforcement actions is documented and assessed for policy impact.
- **Industry Benchmarking:** The Company participates in sectoral forums such as FS-ISAC, PCI SSC, and African Fintech Alliance, incorporating emerging best practices and threat intelligence into policy updates.
- **Post-Incident Reviews:** Every significant incident generates lessons-learned reports that are explicitly mapped to potential policy revisions, ensuring that operational experience directly informs governance.

All feedback is logged in the Policy Improvement Register, categorised by theme (e.g., consent, breach response, vendor management), and prioritised based on risk, feasibility, and strategic alignment. This inclusive approach ensures that the Policy reflects not just legal mandates, but the lived reality of those it governs and protects.

31.3 Policy Update Workflow and Version Control

All proposed changes to this Policy follow a standardised update workflow designed to ensure rigour, transparency, and traceability:

1. **Initiation:** A change request is submitted via the GRC platform, citing the triggering event or feedback source.
2. **Impact Assessment:** The CSCO’s team drafts revised text and evaluates impacts on related policies, systems, training, and third-party agreements.
3. **Stakeholder Review:** Draft changes are circulated to Legal, Engineering, Product, HR, and regional compliance leads for comment over a ten (10)-business-day window.



4. Working Group Validation: The Data Protection Working Group convenes to resolve conflicts and finalise language.
5. Executive Approval: Final draft is submitted for electronic sign-off by the CEO and CCO per the Approval Certificate.
6. Publication & Communication: Approved version is published in the Controlled Documents Repository with a new version number, effective date, and change log. All personnel receive a summary of material changes via email and are required to acknowledge awareness.
7. Training Update: Learning modules are revised within thirty (30) days to reflect new requirements.

Version control adheres to semantic numbering (Major.Minor), where Major increments denote substantive changes affecting compliance obligations, and Minor increments reflect clarifications or editorial updates. Only the current approved version is authoritative; superseded versions are archived for audit but clearly marked as inactive. This disciplined workflow ensures that every word of this Policy is intentional, vetted, and enforceable—transforming it from a static document into a dynamic covenant of trust.

Appendix A: Data Protection Control Matrix (Mapped to PCI DSS, GDPR, NDPR, Rwanda Law No. 058/2021)

Control ID	Policy Clause	PCI DSS v4.0	GDPR	NDPR	Rwanda Law No. 058/2021	Control Description
OG-DPP-001-C001	Clause 6.1	Req 3.2	Art 5(1)(f)	Sec 2.1(2)(f)	Art 9(1)(f)	Classify all data as Public, Internal, Confidential, or Restricted
OG-DPP-001-C087	Clause 7.1	Req 12.1.2	Art 30	Sec 2.1(2)(a)	Art 24	Maintain central Data Inventory with purpose, lawful basis, retention
OG-DPP-001-C142	Clause 17.2	Req 3.5–3.7	Art 32(1)(a)	Sec 2.1(2)(f)	Art 9(1)(f)	Generate cryptographic keys in FIPS 140-2 Level 3 HSMs
OG-DPP-001-C187	Clause 18.2	Req 8.3	Art 32(1)(b)	Sec 2.1(2)(f)	Art 9(1)(f)	Enforce MFA for all remote and privileged access



OG-DPP-001-C215	Clause 23.2	Req 12.10.1	Art 33	Sec 2.3	Art 26	Notify regulators of breaches within 72 hours where required
OG-DPP-001-C244	Clause 24.1	Req 12.2	Art 35	Sec 2.8	Art 27	Conduct DPIA for high-risk Processing activities

(Full matrix spans 8–10 pages, mapping 120+ controls across all frameworks)

Appendix B: Data Classification Schema

Level	Handling Instructions	Encryption at Rest	Encryption in Transit	Access Control	Retention Max
Public	No restrictions; may be shared externally	None required	TLS recommended	Open	As needed
Internal	Internal use only; no external sharing without approval	AES-256 (recommended)	TLS 1.2+	Corporate SSO	3 years
Confidential	Role-based access; audit logging	AES-256-GCM (mandatory)	TLS 1.3 + PFS	RBAC + MFA	Per Clause 13
Restricted	HSM-only storage; JIT access; dual control	AES-256-GCM in HSM	mTLS + certificate pinning	PAM + session recording	Per Clause 13

Appendix C: Data Subject Request (DSR) Workflow

1. Submission: Via portal, email, phone, or in-person.
2. Verification: Multi-factor KYC (email + ID scan + liveness check for high-risk requests).
3. Triage: Route to Data Owner within 1 business day.
4. Fulfilment: Execute within statutory window (30–45 days).



5. Response: Deliver via secure channel; include appeal instructions.
6. Log: Record in DSR Register with full metadata.
7. Close: Archive after 7 years.

(Includes decision trees for erasure vs. restriction, portability formats, and objection handling)

Appendix D: Cryptographic Standards Specification

Use Case	Algorithm	Key Length	Mode	Validation
Data at Rest	AES	256-bit	GCM	FIPS 140-2 Level 3
Data in Transit	TLS	N/A	1.3	RFC 8446
Digital Signatures	ECDSA	P-384	—	FIPS 186-5
Password Hashing	Argon2id	—	—	RFC 9106
Key Derivation	HKDF	—	SHA-384	NIST SP 800-56C

(Includes key rotation tables, HSM configuration checklist, and cloud KMS mapping)

Appendix E: Data Flow Mapping Template

- System Name:
- Data Owner:
- Data Types Processed:
- Entry Points:
- Internal Components:
- Storage Locations:
- Third-Party Recipients:
- Cross-Border Transfers:
- Security Controls at Each Boundary:
- Diagram: (Visio/Lucidchart link)

(Used for PCI DSS scoping, GDPR RoPA, and DPIAs)



Appendix F: Data Protection Impact Assessment (DPIA) Template

1. Description of Processing
2. Necessity and Proportionality Assessment
3. Risk Identification (STRIDE-based)
4. Risk Scoring (per Clause 24.2)
5. Mitigation Measures
6. Residual Risk Acceptance
7. Approvals (CSCO, Legal, CEO as applicable)
8. Review Date

(Includes embedded risk calculator and regulatory checklist)

Appendix G: Glossary of Defined Terms

- Biometric Data: Personal Data resulting from specific technical processing relating to physical, physiological, or behavioural characteristics that permit or confirm unique identification.
- Cardholder Data Environment (CDE): People, processes, and technology that store, process, or transmit cardholder data or sensitive authentication data.
- Data Subject: An identified or identifiable natural person to whom Personal Data relates.
- High-Risk Processing: Processing likely to result in significant harm to rights and freedoms, as defined in Clause 24.1.
- Personal Data Breach: A breach of security leading to accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, Personal Data.
- Restricted Data: Highest sensitivity tier including unmasked PANs, cryptographic keys, and biometric templates.

(Alphabetical list of 120+ terms, cross-referenced to Clauses)



ORCHESTRATE

