Digital Nation Pakistan

PAKISTAN DIGITAL AUTHORITY
Standards & Regulatory Publications

D-SERIES: DATA GOVERNANCE, MANAGEMENT & EXCHANGE

DNP-D.002 RA

Standards & Regulatory Publications —
WASL — Pakistan National Data Exchange Layer
Reference Architecture

Summary
This document defines the reference architecture of WASL, Pakistan's National Data Exchange Layer. It establishes the architectural philosophy, design principles, component model, trust boundary model, four functional planes, and interaction patterns that together constitute the authoritative architectural reference for WASL. It is the baseline against which downstream technical specifications, sectoral profiles, and conforming implementations shall be developed.

Document reference	DNP-D.002 RA
Title	WASL — Pakistan National Data Exchange Layer: Reference Architecture
Document type	Reference Architecture (RA)
Status	Published (PUB)
Version	1.0
Date	13 May 2026
Classification	PUBLIC
Published by	Pakistan Digital Authority
Persistent URL	https://standards.pda.gov.pk/DNP-D.002
DOI	https://doi.org/10.83282/dnp-d-002-ra
Gazette reference	N/A
Maturity Level	PILOT

FOREWORD

The Pakistan Digital Authority (PDA) is the statutory body established under the Digital Nation Pakistan Act, 2025, mandated to issue standards, frameworks, and technical publications in support of the National Digital Masterplan.

WASL — Pakistan's National Data Exchange Layer, is a foundational component of the country's Digital Public Infrastructure. It enables secure, consent-governed, standardized data exchange between federal and provincial entities, regulated private-sector institutions, and authorized platforms, while preserving federated data custody and cryptographic privacy by design.

This document defines the reference architecture of WASL. It is the authoritative architectural baseline against which all downstream WASL technical specifications (TS), sectoral profiles (PRF), and implementation guidelines (GDL) shall be developed. It is being issued at PILOT maturity.

This document was developed and approved by the PDA Standards Board under the procedures defined in DNP-X.001 FWK.

DOCUMENT CONTROL

Version	Date	Description	Approved By
1.0	13 May 2026	Initial publication — WASL Reference Architecture	PDA Standards Board (approved 13 May 2026)

Contact:
Pakistan Digital Authority
4th Floor, 5-A Constitution Avenue,
Sector F-5/1, Islamabad 44000, Pakistan
E-Mail: standards@pda.gov.pk
https://standards.pda.gov.pk

Table of Contents

Reference Architecture Diagram

The diagram below illustrates the WASL national data exchange model as a centrally routed, cryptographically protected, federated interoperability fabric. It should be interpreted as a logical architecture view, not a physical deployment topology.

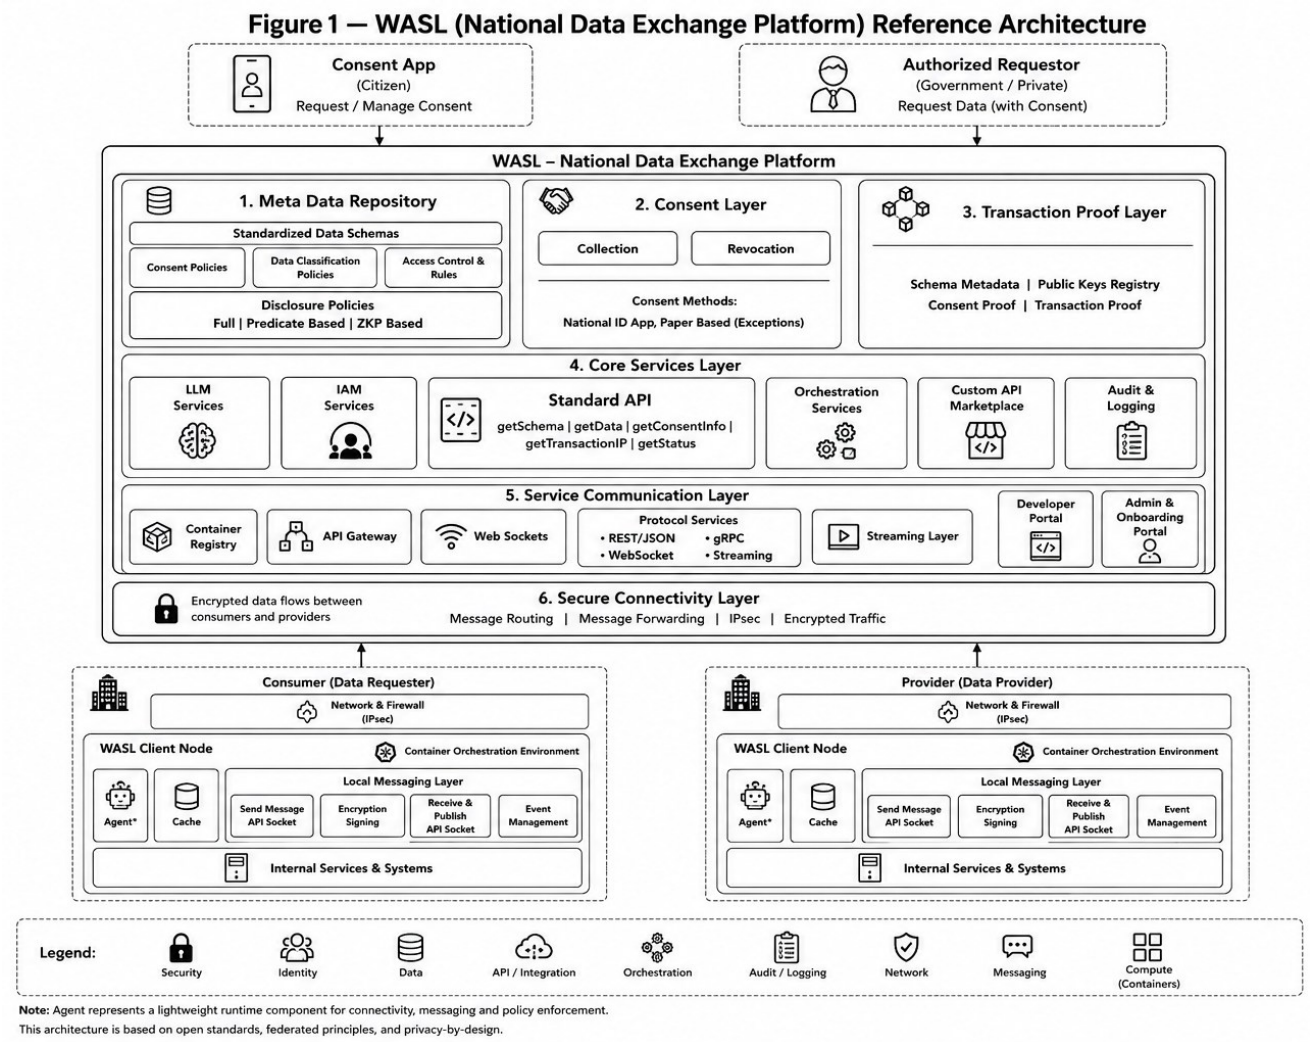

Figure 1: WASL Reference Architecture, illustrating the Consumer Client Node, Central WASL Platform (including Metadata Repository, Standard API, Consent Layer, Transaction Proof Layer, Identity and Access Management (IAM), Orchestration Services, Custom API Marketplace, API Gateway, Developer Portal, and Admin Portal), Provider Client Node, and Secure Connectivity Layer connecting them. The Consent App interaction at the top represents the Citizen Consent Plane.

1 Scope

This document defines the reference architecture of WASL, Pakistan's National Data Exchange Layer. It establishes the architectural philosophy, design principles, logical component model, trust architecture, consent framework, transaction proof model, AI-native capabilities, semantic interoperability, cryptographic agility, and cross-border interoperability design that together constitute the authoritative architectural reference for WASL.

It applies to:

- all technical specifications, profiles, implementation guides, and sectoral extensions published under the DNP designation that derive from or cite the WASL architecture;
- all government, regulated private-sector, and authorized entities participating in the WASL ecosystem as Data Providers, Data Consumers, or platform operators;
- all implementations of WASL Client Nodes, central platform components, and supporting subsystems.

This document does not govern:

- implementation-level technical specifications for individual WASL APIs, message formats, protocol bindings, or conformance test procedures, which are issued separately as Technical Specifications (TS) under the DNP-D series;
- the internal information systems, databases, or business logic of participating organizations;
- the legal, regulatory, and mandatory participation rules governing entities that must transact over WASL, which are issued separately as a Regulation (REG) under the DNP-D series;
- sector-specific data sharing rules, which are issued as sectoral Profiles (PRF) under the appropriate series;
- the lifecycle, issuance, or management of Verifiable Credentials, which are separately governed under the DNP-U (Digital Trust Services) series;
- AI model governance, evaluation, and lifecycle management, which are separately governed under the DNP-A (Artificial Intelligence) series. This document defines only the integration surface through which AI systems interact with WASL.

1.1 Non-Goals

WASL is deliberately bounded. The following are explicit non-goals of the WASL architecture, stated here to prevent scope drift in downstream specifications:

- WASL is not a data lake, data warehouse, or analytics store. It does not hold persistent copies of citizen data.
- WASL is not a national identity provider. It relies on PAK-ID for citizen authentication and consent acquisition.
- WASL is not a Verifiable Credential issuer, holder wallet, or credential lifecycle manager.
- WASL is not a payment rail. Settlement between Providers and Consumers, where applicable, occurs through systems governed under DNP-P.
- WASL provides the technical fabric for data exchange; enforcement of data protection, sectoral regulation, and consumer rights remains with the relevant statutory bodies.
- WASL is not a direct citizen-facing application. Citizens interact with WASL only indirectly through the Consent Plane.

2 References

The following references are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document applies.

2.1 Normative references — National

- Digital Nation Pakistan Act, 2025 (Act No. I of 2025).
- Electronic Transactions Ordinance, 2002 (Ordinance No. LI of 2002).
- Prevention of Electronic Crimes Act, 2016 (Act No. XL of 2016), as amended.
- DNP-X.001 FWK — Standards & Regulatory Publications: Nomenclature, Series Structure & Issuance Framework (03/2026), Pakistan Digital Authority.
- DNP-D.001 FWK — National Data Governance Framework, Pakistan Digital Authority. (To be issued)
- DNP-A.001 FWK — National AI Governance Framework, Pakistan Digital Authority. (To be issued)
- Pakistan Security Standards for Cryptographic and Information Technology Security Devices.
- The National Database and Registration Authority (NADRA) Ordinance, 2000

2.2 Normative references — International

- IETF RFC 2119 / RFC 8174 — Key words for use in RFCs to Indicate Requirement Levels.
- IETF RFC 6749 — The OAuth 2.0 Authorization Framework.
- IETF RFC 7519 — JSON Web Token (JWT).
- IETF RFC 8446 — The Transport Layer Security (TLS) Protocol Version 1.3.
- IETF RFC 9162 — Certificate Transparency Version 2.0.
- OpenID Foundation — OpenID Connect Core 1.0.
- OAuth 2.1 Authorization Framework (latest draft).
- OpenAPI Specification v3.1.
- JSON Schema Specification (Draft 2020-12).
- W3C Verifiable Credentials Data Model 2.0.
- W3C Decentralized Identifiers (DID) 1.0.
- W3C Data Catalog Vocabulary (DCAT) v3.
- W3C Open Digital Rights Language (ODRL) Information Model 2.2.
- ISO/IEC 27001 — Information security management systems — Requirements.
- ISO/IEC 27018 — Code of practice for protection of PII in public clouds.
- ISO/IEC 29100 — Privacy framework.
- NIST SP 800-207 — Zero Trust Architecture.
- NIST FIPS 203 — Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM).
- NIST FIPS 204 — Module-Lattice-Based Digital Signature Standard (ML-DSA).
- NIST FIPS 205 — Stateless Hash-Based Digital Signature Standard (SLH-DSA).

2.3 Informative references

- Dataspace Protocol (DSP) 2025-1 v1.0.0 — Eclipse Dataspace Working Group / International Data Spaces Association.
- DSSC Data Spaces Blueprint v3.0, Data Spaces Support Centre, 2025.

- GovStack Information Mediator Building Block Specification — ITU/DIAL/GIZ/Estonia.
- Universal DPI Safeguards Framework — UNDP / Office of the UN Secretary-General's Envoy on Technology, 2023.
- G20 Framework for Digital Public Infrastructure, 2023.
- X-Road 7 / X-Road 8 Technical Specification — Nordic Institute for Interoperability Solutions (NIIS).
- European Interoperability Framework (EIF) — European Commission
- Model Context Protocol (MCP) Specification — Anthropic, 2024 onwards
- BBS Signature Scheme for Selective Disclosure — IETF CFRG draft.

3 Definitions

For the purposes of this document, the following definitions apply. Definitions are listed alphabetically.

3.1 AI Agent: A software system that uses an artificial intelligence model to autonomously plan, decide, or execute actions, including actions that invoke WASL APIs. An AI Agent may operate on behalf of an institution or, in future capability, on behalf of a citizen under a delegated consent artefact.

3.2 Central WASL Platform: The shared national control-plane providing identity management, metadata governance, consent orchestration, API governance, routing, audit, and transaction proof recording.

3.3 Client Node: A software component deployed within a Participating Entity's infrastructure that performs all cryptographic operations (signing, verification, encryption, decryption) on behalf of that Entity and mediates its participation in WASL data exchange.

3.4 Consent: An explicit, verifiable, revocable authorization granted by a natural person through an approved consent channel for the exchange of personal data concerning that person, for a specified purpose and duration.

3.5 Consent Artefact: A cryptographically signed, machine-verifiable digital record expressing the content, scope, purpose, duration, and signatory of a specific consent grant. Consent artefacts are independently verifiable by Provider, Consumer, and the Transaction Proof Layer.

3.6 Consumer: A Participating Entity that requests data through WASL for a defined and authorized purpose.

3.7 Control Plane: The functional layer comprising identity management, metadata governance, consent orchestration, access control, and developer services. Control-plane components coordinate trust and enforce policy but do not process business data.

3.8 Cryptographic Agility: The architectural property whereby cryptographic algorithms, key sizes, and protocol parameters can be updated or replaced without redesigning the system. In WASL, this enables migration to post-quantum algorithms without disrupting participating entities.

3.9 Data Blindness: The architectural property whereby the Central WASL Platform handles only encrypted payloads and is cryptographically prevented from reading the content of exchanged data, while retaining cryptographic proof that an exchange occurred.

3.10 Data Exchange Plane: The functional layer of WASL responsible for secure routing, protocol mediation, orchestration, and transaction lifecycle management between Consumer and Provider Client Nodes.

3.11 Data Product: A curated, governed, discoverable data offering published by a domain-owning agency through WASL, with a defined schema, quality guarantees, access controls, documentation, and an accountable product owner. The term originates in the Data Mesh paradigm.

3.12 Dataspace Protocol (DSP): An open specification developed by the Eclipse Dataspace Working Group and the International Data Spaces Association defining control-plane interactions (catalog access, contract negotiation, data transfer management) for interoperable data exchange. WASL's v2 interoperability roadmap targets DSP compatibility.

3.13 Delegated Consent: A consent artefact in which the citizen authorizes a designated agent, human or AI, to exercise consent decisions on their behalf, within defined scope, purpose, and duration boundaries. Reserved as an anticipated capability.

3.14 Federated Data Ownership: The principle whereby each Participating Entity retains custody, control, and serving responsibility for its own data and private keys, with WASL performing routing and control-plane functions only.

3.15 Metadata Repository: The authoritative control-plane registry maintaining machine-readable dataset definitions, schemas, provider bindings, consent policies, and classification rules for the WASL ecosystem.

3.16 Model Context Protocol (MCP): An open protocol, originating from Anthropic, that standardizes how AI applications connect to external data sources and tools. Referenced in this document as the anticipated integration surface for AI consumers of WASL.

3.17 Participating Entity: A federal, provincial, or regulated private-sector organization formally onboarded to WASL as a Consumer, Provider, or both, under a duly executed participation agreement.

3.18 Post-Quantum Cryptography (PQC): Cryptographic algorithms designed to remain secure against adversaries with access to cryptographically-relevant quantum computers. WASL architecture mandates cryptographic agility to enable migration to NIST-standardized (or other applicable) PQC algorithms.

3.19 PAK-ID: Pakistan's national digital identity platform operated by or in coordination with NADRA, used by WASL for citizen authentication and consent acquisition.

3.20 Provider: A Participating Entity that holds authoritative source data in its systems and serves that data through WASL in response to authorized and consent-verified requests.

3.21 Secure Connectivity Layer: The network-level component of the Central WASL Platform responsible for encrypted routing of payloads between Consumer and Provider Client Nodes.

3.22 Selective Disclosure: A cryptographic capability by which a Provider discloses only the specific fields required for a given purpose, rather than the full data record, while the response remains cryptographically verifiable.

3.23 Transaction Proof Layer (TPL): The component of the Central WASL Platform that records cryptographic proofs of each transaction (hashes, signatures, timestamps) without recording transaction content.

3.24 Verifiable Credential (VC): A tamper-evident, cryptographically signed digital credential conforming to the applicable Data Model. Issuance and lifecycle of VCs are governed under the DNP-U series and are outside the scope of this document.

3.25 WASL: Pakistan's National Data Exchange Layer, the subject of this document, meaning connection or linkage.

4 Abbreviations

Abbreviation	Definition
API	Application Programming Interface
ARB	Architecture Review Board
CA	Certificate Authority
CNIC	Computerised National Identity Card
CT	Certificate Transparency
DCAT	Data Catalog Vocabulary (W3C)
DCS	Department of Communication Security
DID	Decentralized Identifier (W3C)
DNP	Digital Nation Pakistan
DPI	Digital Public Infrastructure
DSP	Dataspace Protocol
DSSC	Data Spaces Support Centre
EDC	Eclipse Dataspace Components
FIPS	Federal Information Processing Standards
HSM	Hardware Security Module
IAM	Identity and Access Management
JWT	JSON Web Token
MCP	Model Context Protocol
ML-DSA	Module-Lattice-Based Digital Signature Algorithm (FIPS 204)
ML-KEM	Module-Lattice-Based Key Encapsulation Mechanism (FIPS 203)
mTLS	Mutual Transport Layer Security
NADRA	National Database and Registration Authority
NIIS	Nordic Institute for Interoperability Solutions
NIST	National Institute of Standards and Technology (US)
ODRL	Open Digital Rights Language (W3C)
OIDC	OpenID Connect
PDA	Pakistan Digital Authority
PKI	Public Key Infrastructure
PQC	Post-Quantum Cryptography
REST	Representational State Transfer
SKOS	Simple Knowledge Organization System (W3C)
SLO	Service Level Objective
TLS	Transport Layer Security
TPL	Transaction Proof Layer
VC	Verifiable Credential
WASL	Pakistan's National Data Exchange Layer
ZKP	Zero-Knowledge Proof

5 Architectural Philosophy

WASL is a centrally-routed, data-blind exchange architecture with federated data ownership and governance. Its design resolves a structural tension common to national data exchange systems: the need for shared routing and governance infrastructure without creating a centralized surveillance capability or requiring agencies to surrender custody of their data.

Two structural characteristics define the architecture.

5.1 What is centralized

Encrypted data payloads route through the WASL Secure Connectivity Layer. The central platform is on the data path but it cannot read the content of what it routes. Identity, access control, schema governance, audit logging, and consent signaling are managed centrally. This creates a shared control plane that all Participating Entities rely upon for trust establishment, policy enforcement, and auditability.

5.2 What is federated

Data ownership and custody remain entirely with the owning organization. Each agency, federal or provincial, owns, controls, and serves its own data from its own systems. WASL holds no copy of citizen data. Private key custody is also federated: all cryptographic operations occur within each Entity's own Client Node, never at the central platform.

Trust itself is federated. Neither Consumer nor Provider needs to extend blind trust to the central routing layer. Each party independently verifies cryptographic signatures at its own endpoint, using locally cached public keys. The central platform can prove a transaction occurred and holds cryptographic evidence of it but is mathematically prevented from reading the data that was exchanged.

5.3 The privacy guarantee

The privacy guarantee in WASL is cryptographic, not topological. End-to-end encryption between Consumer and Provider Client Nodes ensures that the routing layer handles only ciphertext it cannot decrypt. This is a stronger guarantee than a conventional API management platform or data broker, because it does not require trust in the operator of the central infrastructure.

The model is best described as a hub-and-spoke architecture with cryptographic data isolation and federated ownership. It is distinct from a fully peer-to-peer federated system (such as an X-Road 7 deployment) because encrypted traffic does transit the centre. However, it achieves data blindness equivalent to a peer-to-peer model through end-to-end encryption applied at the Client Node boundary.

5.4 Pakistan context

WASL's architecture is designed to respect Pakistan's constitutional distribution of authority. Provincial governments hold autonomous jurisdiction over several of the most valuable data domains, including land records, health systems, and local registries. A conventional centralized data platform would require provinces to cede control of their data to a federal body, constitutionally sensitive and practically unlikely to achieve broad participation.

WASL addresses this by adopting the principles of a Data Mesh architecture: each federal ministry and provincial agency is the designated data owner and controller for its own domain. WASL provides the exchange protocol and schema standards through which these sovereign data domains interoperate; it does not supersede or aggregate them. No province's data is held, controlled, or aggregated by the federal centre. Each province's agency serves its own data from its own systems.

Interoperability is achieved through common standards and a shared exchange protocol, not through data centralization.

5.5 Standards posture and protocol baseline

WASL's control-plane semantics align with the Dataspace Protocol (DSP) 2025-1 — an Eclipse Specification on an ISO transposition path, adopted by European data spaces, X-Road 8, and Gaia-X-aligned platforms. Alignment is at the level of catalogue vocabulary (DCAT), policy expression (ODRL), and contract-negotiation semantics. Pakistan-specific extensions, including the federated Transaction Proof Layer, the Citizen Consent Plane, PAK-ID identity binding, and cryptographic data blindness, operate alongside DSP and are specified in this RA. Protocol bindings are specified in the Technical Specification.

6 Core Design Principles

The following principles govern all architectural decisions in WASL. They are ordered by architectural primacy, not alphabetically, and collectively define what it means to be a WASL-conformant system.

Principle	Meaning in Practice
Data Blindness	The central platform handles only encrypted ciphertext. It can prove a transaction occurred but cannot read what was exchanged. The privacy guarantee is cryptographic, not administrative.
Cryptographic Three-Party Trust	Consumer, provider, and WASL each independently verify cryptographic signatures. Trust is established at the endpoints, not delegated to the center. No party is implicitly trusted.
Federated Data Ownership	Each federal or provincial agency owns and serves its own data. WASL does not aggregate or replicate data centrally. Private key custody remains with each participating entity.
Standardization First	All providers within a domain conform to PDA-published schemas. A consumer integrates once against the standard interface and can retrieve standardized data from any conforming provider, across provinces, departments, or sectors, with no additional integration effort.
Consent by Default	Access to personal data always requires explicit, verifiable citizen consent, including time-limited and purpose-limited grants and revocation. Consent is a mandatory, blocking prerequisite, not an optional feature.
Non-Repudiation	Every transaction is signed by both parties and recorded in the immutable Transaction Proof Layer. Neither party can deny a transaction occurred. The TPL provides legally defensible evidence of every data exchange.
Client Node Autonomy	All cryptographic operations, signing, encryption, decryption, and verification occur inside the entity's own infrastructure. The central platform never handles plaintext or holds entity private keys.
Cryptographic Agility and Post-Quantum Readiness	All cryptographic primitives are negotiated at the protocol level. WASL v1 deploys in hybrid (classical + post-quantum) mode for signatures and key encapsulation, with full migration to NIST-standardized PQC algorithms planned on the timeline defined in Section 14.
Semantic Interoperability	Beyond syntactic schema alignment, WASL adopts W3C controlled vocabularies (DCAT for data catalogue, ODRL for policy expression, SKOS for concept schemes) to ensure that the meaning of exchanged data is preserved across domains and across borders where applicable.
AI-Native by Design	AI agents are architecturally indistinguishable from any other authorized Consumer. They authenticate, consent, and audit through the same paths, with

	additional provenance metadata (model identifier, purpose, invoking principal). The architecture anticipates citizen-delegated AI agents as a future capability to be made a part of WASL.
Event-Driven Architecture	Providers push change notifications to subscribed consumers, eliminating polling and enabling near-real-time data synchronization without unnecessary load on the central platform.
Locally-Cached Discovery	Schema metadata and provider public keys are cached at each Client Node. Service discovery and request assembly can proceed without constant central round-trips, improving resilience and reducing platform load. The transformation and standardization for the consumer can be streamlined through Agentic AI. Guidelines can be issued including structured prompts to further standardize the process
Open Standards, No Lock-In	All APIs, schemas, and consent policies conform to open, internationally recognized standards owned by PDA. No participating entity or implementation is tied to any vendor's proprietary tooling. National data exchange infrastructure has a decades-long operational horizon.

7 Logical Component Model

The WASL architecture is composed of three structural domains: the Central WASL Platform, the WASL Client Node (deployed by each participating organization within its own infrastructure), and the Secure Connectivity Layer connecting them. This section describes each domain at a component-model level, that is, in terms of functional responsibilities and architectural boundaries, not in terms of implementation technology or specific protocol choices.

7.1 Central WASL Platform

The Central WASL Platform is the shared national control-plane and routing-plane. It provides the services that all participating entities rely on for trust establishment, governance, and routing. Critically, it does not store or process citizen data in plaintext. Its components are:

7.1.1 Identity and Access Management (IAM)

IAM is the trust anchor of the WASL ecosystem. It is responsible for establishing and verifying the identity of all participating organizations and their authorized applications. It issues verifiable access tokens, enforces access policies, and integrates with PAK-ID as the national identity provider. It supports both organizational identity (for agencies, banks, and other institutions) and citizen identity (for consent interactions). IAM uses open, standardized token-based authentication and supports role-based and attribute-based access control. No entity is admitted to WASL without explicit IAM-issued credentials.

For AI agents (Section 12), IAM issues tokens bound additionally to a model identifier and the invoking principal, so that the Audit and TPL records capture not only which organization made a request but which AI system, on whose behalf.

7.1.2 Metadata Repository

The Metadata Repository is the authoritative control-plane registry of the ecosystem. It maintains the machine-readable definitions, policies, and service metadata required for interoperable, policy-governed data exchange. It serves as the single source of truth for what datasets are available, how they are structured, which provider organizations can serve them, what access and consent policies apply, and what trust parameters are required for secure invocation.

The Metadata Repository is composed of six sub-registries:

- Data Schema Registry — standardized schemas per domain, versioned and published by PDA;
- Data Catalogue — dataset descriptions published in W3C DCAT v3 vocabulary, enabling semantic discovery and future DSP interoperability (see Sections 13 and 17);
- Consent Policy Registry — rules specifying which data fields require citizen consent, under what conditions, and with what validity constraints; policies expressed in W3C ODRL;
- Data Classification Policies — sensitivity tiers and corresponding minimum technical and governance controls;
- Provider Directory — binding of datasets to the authorized Provider organizations that serve them, including endpoint references and public key materials;
- Cryptographic Parameter Registry — the currently permitted algorithm suites, minimum key sizes, and deprecation schedules for all cryptographic primitives used across the WASL ecosystem (see Section 14).

Client Nodes cache relevant metadata locally, enabling service discovery and request construction without constant central round-trips. The Repository supports selective schemas including annotations for Zero-Knowledge Proof (ZKP) capable responses where technical and legal conditions are met.

7.1.3 Standard API Layer

The Standard API Layer defines the uniform interface contract through which all data exchange operations are performed within WASL. It provides a consistent, protocol-governed mechanism for requesting, retrieving, and verifying data across all Participating Entities. Rather than allowing each Provider to expose arbitrary APIs, WASL enforces a standardized interaction model, ensuring Consumers integrate once and can interact with multiple Providers without bespoke per-Provider work.

The standard API set encompasses the following core operations: retrieval of schema definitions (getSchema); retrieval of data for a specified subject from a designated Provider in the standardised schema format (getData); retrieval of the status and details of a consent record (getConsentInfo); query of the Transaction Proof Layer for the status and proof of a specific transaction (getTransactionTPLStatus); event subscription and notification (subscribe, publish, notify); and selective-disclosure requests (getDataDisclosed). These operations cover the entire data exchange lifecycle. All requests shall be authenticated, consent-verified, schema-compliant, and auditable.

7.1.4 API Gateway

The API Gateway is the central runtime enforcement and routing layer for synchronous request/response traffic. It is the single entry and exit point for all synchronous API calls, ensuring that every transaction is validated, governed, routed, and recorded according to WASL policies. It is a policy enforcement point, not a simple proxy, that integrates with IAM, the Metadata Repository, the Consent Layer, and the Transaction Proof Layer. No synchronous request may reach a Provider without passing through Gateway validation. The Gateway handles token authentication, consent validation, schema compliance, cryptographic parameter compliance, rate limiting, and routing.

Event-driven, asynchronous, and streaming traffic is handled by a separate Central-Platform component, the Event and Streaming Bus (Section 7.1.12), with its own distinct policy enforcement. The two are kept architecturally separate because synchronous and asynchronous patterns have different back-pressure, delivery-guarantee, and consent-re-evaluation semantics. Conflating them into one component reliably under-serves both.

7.1.5 Consent Layer

The Consent Layer manages the full lifecycle of citizen consent for access to personal data. It initiates and manages consent requests, captures and records citizen decisions, generates cryptographically verifiable consent artefacts, enforces consent validity at runtime, and supports revocation and expiry. Consent is not treated as a user interface event but as a cryptographically verifiable, policy-bound authorization artefact that must accompany every data request involving personal data.

Consent artefacts are expressed in a format derived from W3C ODRL and are forward-compatible with the W3C Verifiable Credentials Data Model 2.0, enabling consent grants themselves to be held and presented as credentials where the downstream use case requires it.

The Consent Layer supports multiple collection channels, including mobile application (PAK-ID), and at a later stage, SMS and USSD for feature-phone users, and assisted consent at service delivery points, to ensure inclusion of citizens without smartphone access, a requirement that is particularly acute in Pakistan's rural context. Consent notifications are delivered in Urdu and regional languages. Consent can be revoked at any time; revocation takes effect immediately for future transactions while historical records remain auditable.

The Consent Layer also hosts the delegated-consent scaffolding described. In v1, this scaffolding is structurally present but not operationally active; full delegated-consent for AI agents acting on behalf of citizens is a planned v2 capability.

7.1.6 Transaction Proof Layer (TPL)

The TPL provides a verifiable, tamper-evident record of all data exchange transactions. The TPL is described in full in Section 11, including its design rationale, threat model, and anchoring strategy. At the component-model level, it is composed of: a transaction proof store; a Merkle aggregation service; a public Certificate Transparency-style log serving as the external anchor; and a proof-verification API exposed to external auditors, regulators, and Participating Entities.

7.1.7 Orchestration Services

Orchestration Services coordinate multi-step, multi-system workflows that cannot be completed through a single API request. They coordinate interactions between IAM, the Metadata Repository, the Consent Layer, the API Gateway, and multiple Provider Nodes to execute complex transactions such as eligibility determinations, multi-agency verifications, and asynchronous processing flows. Orchestration is implemented as a state-driven workflow layer, not as ad hoc API chaining, so that complex flows remain reviewable, auditable, and restartable.

7.1.8 Audit and Logging

The Audit and Logging component provides comprehensive, structured, and tamper-aware recording of all system activities. It captures operational, security, and transaction-level events across all components, enabling regulatory compliance, system monitoring, anomaly detection, and forensic investigation. Unlike the TPL, which provides cryptographic evidence of specific transactions, the Audit and Logging layer provides system-wide visibility and behavioural traceability. Logs are retained according to national data governance requirements and do not contain raw personal data in plaintext. Log schemas align with OpenTelemetry semantic conventions (see Section 16) to enable standard tooling for observability and anomaly detection.

7.1.9 Custom API Marketplace and Developer Portal

The Custom API Marketplace allows Participating Entities to publish domain-specific APIs beyond the standard set, subject to PDA review and governance approval. This provides a controlled innovation layer within the WASL ecosystem, enabling sector-specific workflows, composite services, and analytics APIs while ensuring all interactions remain within the WASL trust,

authentication, consent, and audit framework. The Developer Portal provides schema discovery, sandbox access, credential management, API subscription, and usage monitoring, reducing onboarding friction and standardizing integration workflows for new participants.

7.1.10 Billing and Monetization

The Billing and Monetization component supports usage-based and other applicable billing models for data exchange over WASL. It provides automated purchase order and invoice generation, metering of API usage, and integration with government e-payment infrastructure.

Whether a given data exchange is charged, and at what rate, is determined by the applicable Billing Policy. The Policy distinguishes government-to-government exchange, government-to-regulated-private-sector exchange, and commercial-Consumer access; sets rate schedules where applicable; and governs approvals for Data Provider-specific charging arrangements.

WASL's architecture does not mandate a specific charging posture. The Billing Policy is where that posture is set. As a matter of Pakistan's Digital Public Infrastructure principles, the Policy is expected to favour free or nominal-cost arrangements for government-to-government exchange; but this is a policy choice under the Billing Policy, not an architectural commitment of the RA. The billing policy will be allowing the WASL owner, operator, and data providers to charge consumers for API access where applicable, with automated purchase order and invoice generation, and digital payment integration.

7.1.11 AI Integration Services

The AI Integration Services component provides the platform-side integration surface for AI-native capabilities described in Section 12. This includes model-identifier registry; AI-agent token binding (issued jointly with IAM); anomaly detection services consuming structured audit streams; and the MCP server surface through which AI agents discover WASL schemas and consent requirements. AI Integration Services are control-plane components: they do not process citizen data. They are described in Section 12.

AI Integration Services do not impose additional cryptographic key requirements beyond those applicable to the Entity's Client Node. An Institution deploying an AI agent consumer follows the same key-management flexibility defined in Section 14.2: HSM for high-risk institutions, software keystores acceptable for others. AI agents inherit the key management posture of their hosting organization.

7.1.12 Event and Streaming Bus

The Event and Streaming Bus is the Central-Platform-side counterpart to Client Node Event Management (Section 7.2.6). It is the policy enforcement and routing point for asynchronous, event-driven, and streaming traffic, deliberately kept separate from the API Gateway (Section 7.1.4) because the two interaction patterns have different operational characteristics: synchronous traffic is transactional and individually policy-evaluated; asynchronous traffic is long-lived, subscription-based, and requires ongoing consent re-evaluation over the subscription lifetime.

The Event and Streaming Bus is composed of: a subscription registry (recording which Consumers subscribe to which event streams, under which consent artefacts, for what duration); a fan-out routing service (delivering events from Provider Client Nodes to all authorized Consumer Client Nodes); delivery-guarantee machinery (at-least-once delivery with deduplication identifiers; dead-letter handling; replay from a configurable retention window); subscription-level policy enforcement (periodic re-verification that the underlying consent remains valid; automatic subscription suspension on consent revocation); and telemetry integration with Observability (Section 16.2). Like the Gateway, it is data-blind: event payloads are end-to-end encrypted between Provider and Consumer Client Nodes; the Bus routes ciphertext, not content.

Standard streaming patterns supported include change-data-capture feeds for subscribed data domains, notification events for status changes, and high-throughput telemetry streams for operational monitoring use cases. The full protocol binding, including transport choice, event-schema format, and delivery semantics, is specified in the associated Technical Specification.

7.2 WASL Client Node

The WASL Client Node is the primary execution and trust enforcement component deployed within each Participating Entity's own infrastructure. Every Participating Entity, whether Consumer, Provider, or both, deploys a Client Node. This is what distinguishes WASL from a conventional API management platform: the Client Node is the mechanism through which cryptographic data isolation is achieved, and through which each Entity retains full autonomy over its cryptographic keys and operations.

7.2.1 Secure Connectivity

The entry point of each Client Node is a secure network boundary enforced through firewall and encrypted communication channels with the WASL platform. All WASL traffic is routed through authenticated, encrypted site-to-site tunnels. Non-whitelisted traffic is blocked at the perimeter. Mutual authentication is enforced at both the network and application layers.

7.2.2 Local APIs

The Local APIs are the integration surface for the Entity's internal applications. A consuming application calls the Local API; the Client Node handles all WASL protocol complexity on its behalf, token acquisition, consent initiation, signature assembly, and request dispatch. This abstracts WASL protocol complexity from the application layer entirely, enabling organizations with limited technical capacity to participate.

7.2.3 Cache and Local Storage

The Client Node maintains locally cached copies of schema metadata, Provider public keys, cryptographic parameter sets, and event subscription records. This dramatically reduces load on the central platform and enables service discovery and request construction without constant central round-trips. Cached metadata is refreshed on demand, on a schedule, or upon expiry, using signed metadata bundles to prevent tampering.

7.2.4 Signature Verification and Encryption

This is the cryptographic core of the Client Node. On the Consumer side, it verifies the Provider's digital signature on every response, verifies the integrity of the consent record, and verifies the hash of the response data against the signed hash provided by the Provider. On the Provider side, it verifies the Consumer's digital signature on the inbound request, verifies the authenticity of the consent record, and verifies request integrity, before any internal data retrieval is initiated.

For end-to-end encryption, the Consumer encrypts the session key using the Provider's public key and encrypts the payload using that session key. The WASL platform routes the encrypted payload without the ability to decrypt it. The Provider's Client Node decrypts using its private key. This guarantees platform data blindness: the central node routes the message and records the proof, but is cryptographically prevented from reading the content.

The Client Node implements cryptographic agility at this boundary: the supported algorithm suites are negotiated per session using parameters published in the Metadata Repository, enabling migration from classical to hybrid to pure post-quantum cryptography without protocol-level disruption (see Section 14).

7.2.5 Fulfilment Layer

On the Provider's Client Node, the Fulfilment Layer receives inbound data requests, performs all verification steps (consent validity, Consumer signature, request integrity, payload decryption), calls the Entity's internal backend systems to retrieve the requested data, validates the response against the relevant published schema, signs the response with the Provider's private key, and simultaneously submits the Provider-side proof record to the central TPL. The Fulfilment Layer also implements selective-disclosure logic (Section 10.4) for schemas where this capability is enabled, returning only the specific fields the Consumer is entitled to receive under the applicable consent artefact.

7.2.6 Event Management

The Event Management component enables proactive, push-based data sharing between Entities. A Consumer subscribes to changes in specific data fields for specific identifiers from a Provider, subject to valid consent. When the Provider's internal systems record a change to a subscribed field, the Event Management component signs the change notification and pushes it to subscribed Consumers via the Secure Connectivity Layer, eliminating polling-based integration patterns and enabling near-real-time data synchronization.

7.2.7 AI Agent Integration Surface

The Consumer Client Node's Local APIs are callable by an AI agent operating within the consuming organization's infrastructure. From WASL's perspective, an AI agent is an authorized Consumer application, subject to the same authentication, consent, and audit requirements as any other application, with additional model-identifier binding. The design and governance for AI consumers is defined in Section 12. The Client Node surface also supports the Model Context Protocol (MCP), allowing AI agents to discover WASL schemas and consent requirements through a standardized protocol (Section 12.5).

7.3 Secure Connectivity Layer

The Secure Connectivity Layer is the trusted transport and routing fabric that connects WASL Client Nodes with the Central WASL Platform and, through it, with other Participating Entities. It provides the secure communications backbone over which all WASL traffic travels, including synchronous request-response traffic, asynchronous event notifications, consent-related signalling, and routing metadata.

Where end-to-end encryption is applied between Client Nodes, the Secure Connectivity Layer transports ciphertext and routing metadata without visibility into the underlying business payload. It supports reliable message delivery for both synchronous and asynchronous interactions, including retry, delivery acknowledgement, and dead-letter handling. The Secure Connectivity Layer is a neutral transport layer: it does not perform consent decisioning, schema validation, access control, or business payload interpretation.

8 Trust Boundary Model

WASL implements a three-party cryptographic trust model. No single entity, including the WASL platform itself, is implicitly trusted. Trust is established independently between Consumer and Provider through cryptographic verification, with WASL acting as a coordinating intermediary that can prove transactions occurred without reading their content.

8.1 The three parties

The three parties in every WASL transaction are: the Consumer (the Entity requesting data), the Provider (the Entity serving data from its systems), and the WASL platform. Each party holds its

own cryptographic key pair issued by the relevant entities. Each party independently verifies the signatures and credentials of the others.

8.2 Consumer trust establishment

Before a Consumer's request reaches a Provider, it shall carry: an IAM-issued access token authenticating the Consumer organization (and, for AI-originated requests, the bound model identifier); a signed consent artefact issued by the Consent Layer, proving the citizen approved the specific data access for the declared purpose; a Consumer signature over the request payload, proving the request has not been tampered with in transit; and a transaction identifier enabling end-to-end auditability. The Provider's Client Node independently verifies all artefacts before fulfilling any request.

8.3 Provider trust establishment

The Consumer's Client Node independently verifies the Provider's digital signature on every response before passing data to the consuming application. This verification uses the Provider's public key retrieved from the Metadata Repository, not from the WASL routing layer. The Consumer also verifies the hash of the response data against the signed hash provided by the Provider, ensuring the response has not been modified in transit through the central platform.

8.4 Platform trust establishment

The WASL platform establishes its role in the trust model through: IAM token issuance and validation; signed consent artefacts issued by the Consent Layer (verified independently by both Consumer and Provider); signed metadata and public key materials published through the Metadata Repository; and TPL proof records that cryptographically commit both parties to the transaction. The platform's trustworthiness does not depend on any party taking its word for it, all platform-issued artefacts carry cryptographic signatures that can be independently verified, and the TPL proofs are externally auditable via the Certificate Transparency log (Section 11).

8.5 Security model summary

WASL implements a Zero Trust security model throughout. Every API call is authenticated using token-based credentials, every connection is mutually verified, every network path is encrypted, and no Entity is implicitly trusted at any point. Access control decisions are evaluated at runtime against centrally defined but dynamically enforceable policies derived from the Metadata Repository. The following table summarizes the key security controls by layer.

Security Layer	Mechanism	What it Achieves
Organization identity	Token-based authentication with mutual node verification	Every request is attributable to a registered, authorized Entity
AI principal binding	Model identifier bound to access token via IAM	AI-originated requests are distinguishable and traceable to the specific model
Network transport	Encrypted site-to-site connectivity between Entity and platform (mTLS 1.3, hybrid PQC-ready)	All WASL traffic is confidential in transit at the network layer
Payload confidentiality	End-to-end encryption between Client Nodes (Consumer	Platform data blindness; routing layer cannot read payload content

	encrypts with Provider public key)	
Request integrity	Consumer signs request payload; Provider verifies signature	Tampering in transit is detectable
Response integrity	Provider signs response; Consumer verifies signature	Consumer is assured the response is genuine and unmodified
Consent verification	Signed consent artefact verified independently by Provider and TPL	No personal data is accessed without verified citizen authorisation
Non-repudiation	Both parties sign; TPL records hashes and signatures; external CT-log anchoring	Neither party can deny the exchange occurred; an external party can verify this without accessing WASL internals
Key management	Relevant PKI issues certificates; Entity private keys never leave Entity perimeter	Compromise of one Entity does not compromise others
Access control	Role-based and attribute-based policies, enforced at Gateway and Client Node	Only authorized Entities access authorised datasets for authorised purposes
Cryptographic agility	Algorithm suites negotiated per session from the Cryptographic Parameter Registry	PQC migration is possible without protocol-level disruption

9 Four Functional Planes

The WASL architecture is interpreted across four functional planes. These planes represent logical separations of concern, not physical deployment boundaries. Understanding the planes clarifies which components handle which types of functions and how they interact.

9.1 Control Plane

The Control Plane encompasses all components that coordinate trust and enforce policy but do not process business data. It includes IAM, the Metadata Repository, the Consent Layer, access control, audit and logging, developer services, and AI Integration Services. Control-plane components define and enforce the rules under which data exchange occurs: they govern who may participate, what may be exchanged, under what conditions, and how transactions are recorded, without seeing the content of the transactions themselves. It is the authoritative source for all trust and governance signals consumed by Client Nodes and the Data Exchange Plane.

9.2 Data Exchange Plane

The Data Exchange Plane handles the mechanics of moving requests and responses between Entities. It encompasses the API Gateway, the Secure Connectivity Layer, Orchestration Services, and the transport-level routing infrastructure. The Data Exchange Plane is responsible for enforcing gateway-level policy controls (authentication, consent validation, schema compliance), routing encrypted payloads to the correct destination, managing the transaction lifecycle, and queuing proof elements for the TPL. It does not interpret the business content of payloads.

9.3 Local Execution Plane

The Local Execution Plane is where all data retrieval, validation, encryption, decryption, and signing operations occur. It encompasses the Consumer and Provider Client Nodes and their integration with internal organizational systems. Every cryptographic operation takes place within the Local Execution Plane, within the infrastructure of the Participating Entity, not at the central platform. The Local Execution Plane is the mechanism through which WASL achieves both cryptographic data blindness and federated data ownership simultaneously.

9.4 Citizen Consent Plane

The Citizen Consent Plane encompasses all citizen-facing consent interactions. Consent is collected, managed, and revoked through citizen-facing channels that may include mobile applications (PAK-ID), SMS and USSD, assisted consent via field agents or service centres, and call-centre support. The Citizen Consent Plane reflects WASL's recognition that consent mechanisms must be inclusive and accessible regardless of digital literacy or device capability, a requirement that is particularly acute in Pakistan's context, where a significant portion of the population does not have access to smartphones or reliable internet connectivity.

The Citizen Consent Plane interacts with the Consent Layer in the Control Plane: the Consent Layer manages the policy and artefact lifecycle, while the Citizen Consent Plane manages the interaction channels through which citizens express their decisions.

10 Interaction Patterns and Transaction Lifecycle

This section describes how WASL components interact to execute a data exchange transaction from end to end. The description is at a logical level; specific protocol bindings and message formats are defined in the Technical Specifications published under the DNP-D series.

10.1 End-to-End Data Exchange Flow

A complete data exchange transaction proceeds through a sequence which is mentioned in detail in the Technical Specifications (TS) under the DNP-D series.

10.2 Consent flow

The consent flow is asynchronous. A getData request for personal data remains in a pending state until the citizen responds through an approved consent channel. The Consent Layer evaluates ODRL-expressed policy rules from the Metadata Repository to determine whether consent is required for the specific dataset and fields requested. Upon citizen approval, a signed consent artefact is generated and attached to the transaction request. Consent artefacts are purpose-bound, time-bound, and revocable. Revocation takes effect immediately for future transactions.

10.3 Event subscription and notification flow

A Consumer may subscribe to data change events for specific fields and subjects from a Provider, subject to valid consent. Subscriptions are recorded in the Provider's local storage. When the Provider's internal systems record a relevant change, the Event Management component on the Provider's Client Node signs and pushes a change notification to subscribed Consumers via the Secure Connectivity Layer. The Consumer's Client Node verifies the notification's authenticity before delivering the update to the consuming application. Event subscriptions remain active until explicitly revoked or until the underlying consent expires or is revoked.

10.4 Selective disclosure

Selected data schemas support selective disclosure. Instead of returning the full data record, the Provider returns only the specific fields the Consumer requires and is authorized to receive under

the applicable consent artefact. Selective disclosure is implemented cryptographically and its process is covered in the Technical Specifications (TS) under the DNP-D series. Selective disclosure is the default data-minimization mode for composite data products; full-record disclosure requires explicit justification recorded in the consent artefact.

10.5 Zero-Knowledge Proof (ZKP) responses

For a narrower category of use cases, schemas may support ZKP-based verification. Instead of retrieving a citizen's underlying personal data attribute (such as date of birth), a Consumer can request a cryptographic proof of a specific predicate (such as confirmation that the citizen is above a required age threshold). The Provider's Fulfilment Layer generates the ZKP against the relevant schema field and returns only the Boolean confirmation. The underlying data is never transmitted. ZKP capability is annotated in the Metadata Repository at the schema-field level and is an advanced data-minimisation capability, not a default response mode. The pertinent technical details on ZKP are covered in the Technical Specifications (TS) under the DNP-D series.

11 Transaction Proof Layer – Design Rationale and Threat Model

This section sets out the design rationale, threat model, and the reasons the adopted design is preferred over alternatives, for the Transaction Proof Layer. It also defines the authoritative external anchor: a Certificate Transparency-style log.

11.1 Purpose

The TPL provides a tamper-evident, independently verifiable record of every data exchange transaction conducted over WASL. Its purpose is to ensure that:

- neither Consumer nor Provider can credibly deny that a specific exchange occurred (non-repudiation);
- the platform cannot credibly fabricate, alter, or delete the historical record of exchanges (platform accountability);
- external auditors, regulators, courts, and Participating Entities can verify the integrity of the transaction record (external verifiability);
- the record contains cryptographic proofs, consent artefact references, and transaction metadata, but no underlying citizen data content, thereby preserving privacy while enabling independent auditability (privacy by design)

11.2 What the TPL records

For each transaction, the TPL records:

- a unique transaction identifier;
- the identifiers of the Consumer organization, Provider organization, and (where applicable) the bound AI model;
- a cryptographic hash of the request envelope;
- a cryptographic hash of the response envelope;
- the digital signatures of Consumer and Provider over their respective envelopes;
- the hash of the signed consent artefact, where applicable;
- precise timestamps from the WASL time service;
- the cryptographic parameter set used (for forensic traceability across algorithm migrations).

The TPL never records data content, field values, or payload plaintext. A reader of the TPL can establish that transaction "**T**" between Consumer "**C**" and Provider "**P**" occurred at time "*t*" with a consent artefact whose hash is "*h*" but cannot establish what data was exchanged.

11.3 Lifecycle

TPL records proceed through three states:

- **Pending:** the Consumer-side proof element has been recorded; the Provider-side proof element has not yet been received.
- **Endorsement:** both Consumer-side and Provider-side proof elements have been received and aggregated into a Merkle batch; the batch is circulated to the federated endorser set for signature.
- **Committed:** the quorum threshold of endorser signatures has been achieved; the batch is immutably appended to the federated ledger.

Pending records unresolved beyond a configurable timeout are marked "Expired" and are themselves recorded as evidence of a non-completion event. Endorsement failures, where a quorum cannot be achieved within the configured window, are recorded as Endorsement-Failed events, flagged for operational investigation, and do not silently disappear.

11.4 Federated endorsement model

The TPL is operated as a federated, multi-party, append-only ledger. PDA is one endorser among several. A transaction record is considered "Committed" only when a policy-defined quorum of endorsers has countersigned the batch. Merkle tree aggregation, cryptographic consistency proofs, and tamper-evident append-only semantics are invariants of the design (conforming either to the cryptographic primitives of IETF RFC 9162 or equivalent constructions); no single endorser, including PDA, can unilaterally modify history.

This model is a substantive strengthening of the trust posture relative to a single-operator append-only log. Under a single-operator model, trust rests on external monitoring to detect retroactive tampering. Under the federated model, tampering is not merely detectable after the fact but cryptographically infeasible without cross-institutional collusion among the endorser set. The federated model aligns with Pakistan's institutional tradition of inter-agency checks and balances and with the sensitivity of the citizen data whose exchange the TPL attests.

11.4.1 Endorser set — principles

The composition of the endorser set is deferred to subordinate regulation (see 11.7). The RA states the following principles that the regulation shall observe:

- institutional independence: endorsers shall include statutory institutions with independent legal mandates, such that no single branch of government exercises sole authority over the ledger;
- technical capability: endorsers shall have the operational capacity to run a witness node with high availability and to meet the signing-latency requirements of the WASL transaction flow;
- legal accountability: each endorser shall be subject to legal obligations regarding the integrity of its endorsement function, defined in the subordinate regulation;
- stability with reversibility: the endorser set shall be sufficiently stable to provide continuous operation but amendable through the process defined in subordinate regulation so that the set can evolve with the ecosystem.
- diversity of custody: private-key custody for TPL endorser signing keys shall be distributed across institutions with independent key-management controls. Higher-assurance hardware

protections, including certified HSM infrastructure, may be mandated by regulation for designated classes of endorsers.

[NOTE: This requirement applies only to TPL endorsers, not to participating data providers or consumers. See Section 14.2 for Entity key-management flexibility.]

11.4.2 Quorum

The endorsement quorum is a policy-defined threshold set by subordinate regulation, balancing two concerns: too low a threshold weakens the integrity guarantee; too high a threshold creates fragility against endorser downtime. The RA requires only that the quorum be (i) a strict majority of the endorser set, (ii) large enough that collusion is politically non-trivial, and (iii) achievable within the operational commitments of the transaction lifecycle. The specific numeric threshold is set in subordinate regulation.

11.4.3 Technology selection

The RA is technology-neutral on the specific ledger implementation. Any construction that satisfies the federated-endorsement, append-only, cryptographic-consistency, and data-blindness requirements of this section is conformant. Candidate implementation families, ranging from Trillian-style federated append-only logs with multi-party witness signatories, to permissioned distributed ledgers, to purpose-built multi-signature Merkle-tree constructions, will be evaluated in the downstream Technical Specification against delivery, operational, and governance criteria. Technology selection at the RA level is neither necessary nor desirable: it would foreclose options that the TS evaluation may surface as preferable, and it would couple the architectural commitment to a specific product whose lifecycle WASL does not control.

11.5 Threat model and design alternatives considered

The TPL design resists the following threat classes:

Threat	Description	Mitigation
Participant repudiation	A Consumer or Provider later denies that an exchange occurred or claims its terms were different.	Both parties sign their envelopes; signatures and hashes are recorded in the TPL; TPL records are anchored to an append-only CT log. Repudiation is mathematically infeasible.
Platform tampering (retroactive)	PDA or a compromised insider retroactively alters the TPL to delete or modify historical records.	Append-only CT log; Merkle consistency proofs; independent witness monitoring. Any tampering becomes detectable at the next monitor checkpoint.
Platform tampering (forward)	PDA selectively omits recording certain transactions.	Consumer and Provider Client Nodes retain local proof records; any Participating Entity can query the TPL and detect if its own transactions are missing. Selective omission is detectable by affected parties.
Privacy leakage through audit	The audit surface itself reveals citizen data to auditors or intruders.	TPL records only hashes and signatures, never content. Even full read access to the TPL reveals no personal data.
Forgery of consent	A Consumer presents a falsified consent artefact.	Consent artefacts are signed by the Consent Layer; the Provider's Client Node verifies the signature against the Consent Layer's public key before fulfilment; the

		consent hash is bound into the TPL record.
Cryptographic algorithm compromise	A signature algorithm used in historical records is broken (e.g. RSA-2048 against a future quantum adversary).	The cryptographic parameter set is recorded per transaction. Records remain verifiable under the original parameters; a regulator or court can interpret confidence accordingly. Cryptographic agility (Section 14) limits forward exposure.

11.6 Interaction with DNP-U (Digital Trust Services)

The TPL is not a Verifiable Credential issuer. However, a TPL record may constitute evidence consumed by an external credential issuer governed under the DNP-U series. For example, a TPL-proven exchange of an educational record between a university and a regulator might serve as evidence inputs to the issuance of a VC asserting that record's authenticity. The boundary is that WASL attests the exchange; the issuer attests the claim. This separation keeps the TPL architecturally simple and preserves the DNP-U series' jurisdiction over credential lifecycle.

11.7 Companion regulation — TPL Federation Regulation

Operationalization of the federated TPL requires a companion regulatory instrument, the TPL Federation Regulation, to be issued as a REG under the DNP-D series and gazette-notified prior to WASL production launch. This RA names the Regulation as a required companion instrument but does not preempt its content. The Regulation shall specify:

- the initial roster of endorser institutions and the criteria for admission;
- the numeric quorum threshold;
- procedures for endorser admission, removal, and rotation, including quorum rules for amending the roster itself;
- operational obligations of each endorser (availability, signing-latency, HSM custody requirements, incident response obligations);
- legal accountability regime for endorser conduct, including consequences for failures of integrity;
- the inter-institutional MoU framework under which endorsers operate.

The endorser roster itself is a governance question requiring inter-institutional agreement. The RA expresses the architectural commitment to a federated model; the Regulation gives it operational substance. Absent the Regulation, WASL cannot be admitted to production operation.

12 AI in WASL

This section defines how artificial intelligence systems interact with WASL, both as components of the WASL platform itself and as external Consumers of WASL services. It is written on the premise that AI is not a peripheral optimization but a primary mode of interaction with national data infrastructure for the coming decade. The architecture distinguishes two categories of AI integration that require different governance treatment: AI in WASL (operational capabilities under PDA's control) and AI as a Consumer of WASL (external AI systems invoking WASL APIs). Governance obligations specific to AI model lifecycle, evaluation, and safety are governed under the DNP-A series; this section defines only the integration architecture.

12.1 AI in the WASL platform

PDA operates the following AI-native capabilities within the WASL Control Plane. Each operates on structured metadata, audit streams, or schemas, never on citizen data content, preserving data blindness.

12.1.1 Consent comprehension assistance

Consent artefacts are dense machine-readable documents. Citizens with limited digital literacy cannot meaningfully consent by reading them. The Consent Layer offers AI-generated, bias-audited plain-language summaries of consent requests, rendered in Urdu and regional languages, including a non-technical explanation of the specific data fields, the requesting organization, the declared purpose, the duration, and the consequences of granting or withholding consent.

This capability is operated under DNP-A governance and is subject to explicit quality and fairness evaluation. The AI summary is never a substitute for the underlying consent artefact, which remains the legally binding record; the summary is an accessibility aid with legal weight defined by subordinate regulation.

12.1.2 Anomaly detection in the audit and TPL streams

AI models consume the structured audit and TPL streams to detect anomalous access patterns across agencies, for example, unusual query volumes, out-of-hours access, cross-agency query patterns inconsistent with legitimate workflows, or access patterns consistent with data exfiltration. Detections are surfaced to Security Operations Centre and to the affected Data Provider. The model operates on structural metadata only: timestamps, identifiers, frequencies, hashes. It does not read payload content.

Anomaly detection is a recommending system, not an enforcement system. Any action taken on a detection, suspension of access, investigation, referral, follows the human-in-the-loop governance pattern defined under DNP-A.

12.1.3 Schema mapping and semantic integration assistance

New Participating Entities onboarding as Data Providers frequently have internal data structures that do not match the PDA-published schema for their domain. AI assists this mapping by proposing field-level correspondences between an Entity's internal schema and the relevant DNP-published schema. All proposed mappings are reviewed and approved by the onboarding Entity and by PDA before being committed. The AI accelerates a task that requires integration effort; it does not bypass human review.

12.1.4 Quality and drift monitoring

AI services monitor schema usage, response latency, error rates, and data quality indicators (completeness, freshness, consistency) across the ecosystem, identifying domains where quality is drifting and where intervention may be needed. These insights feed into the Data Product ownership model described in Section 22.

12.1.5 API design and conformance assistance

AI supports the API lifecycle at three points. At design time, API authors use AI assistance to draft OpenAPI 3.1 specifications, propose schema structures consistent with the National Data Vocabulary (Section 13), draft ODRL-expressed consent policies, and produce developer-facing documentation. All AI-drafted artefacts are reviewed and approved by the publishing authority before committal to the Metadata Repository; AI is a productivity tool, not an authorization path.

At onboarding, AI-driven conformance testing exercises a candidate Provider's implementation against the published schema, generating positive and adversarial test cases, identifying schema

drift, consent-policy mis-enforcement, and integration gaps before a Data Product goes live. Similar AI-driven re-certification runs are executed on a scheduled cadence for deployed Data Products.

At integration, AI assists in generating Fulfilment Layer adapter code that translates between a Provider's internal schema and the PDA-published schema. Generated adapter code is subject to security review, static analysis, and conformance testing before deployment; the AI accelerates an integration task that typically takes weeks of engineering effort but does not remove the Provider's accountability for the correctness of its production Fulfilment Layer.

12.2 AI as a Consumer of WASL

AI systems, whether classical machine-learning models, large language models, autonomous agents, or composite systems combining multiple AI components, may invoke WASL APIs as Consumers. This is increasingly the default integration pattern for new applications: AI agents driving loan underwriting, benefits eligibility determination, fraud detection, case triage, and composite service delivery all require access to authoritative citizen and organizational data.

From WASL's perspective, an AI agent is a Consumer application that satisfies all the usual requirements, with additional identity binding. The governance principle is architectural parity: AI-originated requests are subject to the same authentication, consent, and audit obligations as human-originated requests, with added provenance so that an AI-driven decision can always be traced to the specific model, version, and invoking context.

12.2.1 AI identity binding

An access token issued to an AI-driven application is additionally bound to:

- a registered model identifier (naming the specific model, version, and training provenance, from the AI model registry governed under DNP-A);
- the invoking principal (the human operator, service account, or orchestrating system that initiated the AI action);
- the declared purpose of the AI action (for consent policy evaluation).

These bindings are recorded in the TPL entry for every transaction, so that the audit record can answer not only 'which organization requested this data?' but 'which AI model requested it, on whose behalf, for what purpose?'.

12.2.2 Consent when AI is the Consumer

When an AI agent acts on behalf of an institution, for example, a bank's loan underwriting AI, the consent model is unchanged from a human-operated loan officer: the citizen consents to the institution receiving the data for the declared purpose, and the institution is accountable for ensuring that its AI uses the data only as consented. The consent artefact is not materially different; what changes is the TPL record, which captures the AI provenance so that the institution's conduct is later auditable.

12.2.3 AI agents acting on behalf of citizens

A growing pattern internationally involves AI agents acting on behalf of individual citizens — negotiating with service providers, comparing offers, completing applications. This capability is not operationally enabled in WASL v1. However, the v1 architecture includes the scaffolding necessary to support it in v2, specifically:

- delegated-consent artefact types in the Consent Layer schema, expressing that citizen C authorizes agent A, within scope S, for purpose P, for duration D;
- agent registration in the Metadata Repository, with clear identification of the principal (citizen) and the agent (AI system);

- TPL capture of the delegation chain, so that every agent action is traceable to the authorizing consent;
- revocation primitives that allow the citizen to terminate an agent's authority with immediate effect for future transactions.

Full operational enablement of citizen-delegated AI agents requires resolution of several policy questions, liability chain, model accreditation, required safeguards against agent manipulation of citizens, that are deliberately deferred to v2 and to parallel policy development under the DNP-A series.

12.3 Governance guardrails

All AI interactions with WASL, whether AI in the platform or AI as Consumer, are subject to the following architectural guardrails:

- **Parity of audit obligation.** AI-originated actions are recorded with provenance that is at least as rich as human-originated actions, not less.
- **No autonomous consent granting.** An AI agent cannot grant consent on a citizen's behalf except under a pre-existing, citizen-signed delegated-consent artefact. AI cannot bootstrap its own authority.
- **No privileged data access.** AI agents have no access privileges that equivalent human operators do not have. There are no 'AI-only' data endpoints.
- **Purpose binding.** Every AI-originated request declares a purpose; the consent artefact and access policy are evaluated against that declared purpose.
- **Model accountability.** TPL records carry the bound model identifier, so that adverse actions traceable to a specific model, version, or training corpus can be investigated under the governance procedures of DNP-A.
- **Revocability.** Any consent enabling AI access is revocable with immediate effect for future transactions.
- **Human accountability for AI-drafted artefacts.** Schemas, APIs, adapter code, consent policies, and documentation drafted with AI assistance remain subject to human review and sign-off before committal to the authoritative Metadata Repository or deployment to production. AI accelerates but does not replace the accountable author.

12.4 Model Context Protocol (MCP) integration surface

The Model Context Protocol (MCP) is an emerging open protocol standardizing how AI applications discover and invoke external data sources and tools. WASL exposes an MCP server surface through the Developer Portal, allowing MCP-compatible AI clients to discover WASL schemas, consent requirements, available datasets, and integration patterns through a standardized protocol rather than through bespoke integration work.

The MCP surface does not bypass any WASL security control. An MCP client still obtains IAM tokens, still obtains consent artefacts, and still transits the Gateway. MCP is a discovery and invocation convenience layer, not a privileged channel. WASL's MCP surface is one of the first national DPI components to offer such an interface, positioning Pakistan favourably as the primary interaction pattern for enterprise and government AI shifts towards agentic architectures.

12.5 Alignment with AI governance

This section defines the integration architecture between AI systems and WASL. Governance of the AI systems themselves, model registration, evaluation, safety testing, fairness auditing, bias assessment, lifecycle management, and incident response, is governed under the DNP-A (Artificial

Intelligence) series. An AI Consumer of WASL shall be registered, evaluated, and governed under DNP-A; WASL enforces the technical integration requirements defined here, and defers substantive AI governance to the relevant DNP-A instruments.

13 Semantic Interoperability

Syntactic interoperability, agreement on message formats and field names, is a necessary but insufficient condition for successful data exchange. The deeper requirement is semantic interoperability: agreement on what the fields mean, how terms are defined, and how concepts map across domains and across borders. Most large-scale integration efforts fail on semantics, not on syntax.

13.1 Vocabularies

WASL adopts the following W3C-standard vocabularies as the semantic backbone:

- **DCAT (Data Catalog Vocabulary) v3** — for describing datasets and data services. All dataset descriptions in the Metadata Repository are published in DCAT. This aligns WASL with European open data portals and with the DSP catalog model (Section 17). PDA governs mandatory metadata publication standards across all ministries and provinces.
- **ODRL (Open Digital Rights Language) v2.2** — for expressing consent and access policies. Policies in the Consent Policy Registry are expressed in ODRL. ODRL is machine-actionable: a Client Node can evaluate a policy algorithmically without interpreting natural-language legal text.
- **SKOS (Simple Knowledge Organization System)** — for representing controlled vocabularies, concept schemes, and thesauri. A national Data Vocabulary, governed by PDA in coordination with domain custodian agencies, provides the authoritative term definitions for all schemas.
- **JSON Schema (Draft 2020-12) and JSON-LD** — for schema definition and linked-data expression of data records, enabling semantic annotation of field values with stable URIs.

13.2 The National Data Vocabulary

PDA will maintain a National Data Vocabulary at standards.pda.gov.pk/vocabulary, providing stable, persistent URIs for fundamental concepts used across Pakistan's data ecosystem, administrative units (divisions, districts, tehsils), institutional identifiers, classification codes, and reference values. Schemas reference vocabulary terms by URI rather than by string, eliminating the ambiguity that plagues cross-agency data integration.

The National Data Vocabulary is itself expressed in SKOS and is openly published. It is maintained through the DNP-D issuance process: new terms and revisions follow the same working-draft, public-comment, final-draft lifecycle as other DNP publications.

13.3 Domain ontologies

Where domain complexity warrants it, PDA and the relevant custodian agency may jointly publish a domain ontology extending the National Data Vocabulary, for example, a Healthcare Ontology (in coordination with the Ministry of Health and provincial health departments), a Land Records Ontology (in coordination with provincial Boards of Revenue), or a Tax Ontology (in coordination with FBR). Domain ontologies are published as Profile (PRF) publications under the DNP-D series.

13.4 Cross-lingual and cross-script considerations

Pakistan's data ecosystem operates across Urdu, English, and several regional languages. The National Data Vocabulary provides multi-lingual labels for each concept. Schemas are defined in English (for global interoperability) but include SKOS labels in Urdu as a first-class field, so that citizen-facing presentations can render terminology in the language the citizen understands.

14 Post-Quantum Cryptography and Cryptographic Agility

National data exchange infrastructure has a decades-long operational horizon. Over that horizon, the advent of cryptographically-relevant quantum computers is not a speculative concern: a transition to post-quantum cryptography (PQC) is now a planned activity across major national infrastructures globally. NIST finalized the first three PQC standards in August 2024 (FIPS 203, 204, 205). WASL's architecture commits to PQC readiness from v1, with a staged migration.

14.1 Cryptographic agility

WASL's protocol is cryptographically agile: algorithm suites, key sizes, and parameter choices are negotiated per session using the Cryptographic Parameter Registry published in the Metadata Repository. This design isolates cryptographic migration from protocol migration, PQC algorithms can be introduced, preferred, and eventually mandated without redesigning the protocol. Each TPL record carries the parameter set used, so that forensic interpretation of historical records remains possible under the original parameters.

The Stage 1 (v1 launch) algorithm suite is specified as follows. These are the minimum requirements and can be upgraded based on the nature of client data. Where Pakistan Security Standards (PSS) specify additional requirements or restrictions on any of these primitives, PSS prevails (see 14.1.1).

- **Signatures:** ECDSA on NIST curves P-256 and P-384; Ed25519 (RFC 8032); RSA-PSS (RFC 8017) with minimum 3072-bit keys. RSA PKCS#1 v1.5 signatures are not permitted for new issuance.
- **Key encapsulation / key agreement:** ECDH on Curve25519 (X25519), P-256, and P-384; RSA-OAEP with minimum 3072-bit keys.
- **Symmetric encryption:** AES-256-GCM and ChaCha20-Poly1305.
- **Hash functions:** SHA-256 and SHA-3-256; SHA-384 permitted where curve strength warrants.
- **Transport:** mutual TLS 1.3 (RFC 8446) with mandatory forward secrecy.

Deprecated and forbidden for new issuance in Stage 1: RSA PKCS#1 v1.5 signatures; MD5; SHA-1; RSA keys below 3072 bits; ECDSA on curves below P-256; pre-TLS 1.3 transport. Historical records signed under deprecated primitives remain verifiable under their original parameters with confidence interpreted accordingly. All components within WASL will undergo a thorough security evaluation.

14.1.1 Alignment with Pakistan Security Standards

The Cryptographic Parameter Registry shall at all times conform to the currently-ratified Pakistan Security Standards (PSS). Where PSS specifies algorithm suites, key-length requirements, or deprecation schedules, those specifications are authoritative for WASL; the algorithm suite in 14.1 is the architectural baseline, applicable only to the extent not superseded by PSS. Conflicts between PSS and the international references cited in Section 2.2 are resolved in favour of PSS. PSS revisions shall be reflected in the Cryptographic Parameter Registry within a bounded window (specified in the operating regulation), with Participating Entities notified through the standard Entity-notification channel and transition windows aligned to PSS-mandated timelines.

14.2 Key management obligations

All private keys remain in each Entity's own infrastructure and must be protected through appropriate cryptographic key protection mechanisms:

For TPL Endorsers: Hardware Security Module (HSM) with FIPS 140-3 Level 3 certification is mandatory.

For Data Providers and Consumers: HSM recommended for high-risk data domains (civil registry, health records, tax data, land records) and such categories will be mentioned in the Technical Specifications document; software keystores with strong operational security (encryption at rest, access control, audit logging) are acceptable for lower-risk use cases. Specific mechanisms are defined in the Technical Specification based on data classification.

The PKI does not hold or escrow Entity private keys. Key rotation, revocation, and recovery procedures are specified in the associated Technical Specification.

15 Privacy-Preserving Analytics

A recurring use case for national data exchange is population-level analytics: policymakers, planners, and researchers need aggregate insight from distributed agency data without creating a new centralized data store and without compromising individual privacy. WASL supports privacy-preserving analytics as a first-class capability pattern, distinct from the record-level exchange described elsewhere in this document.

15.1 Aggregate query with differential privacy

A Consumer (typically a research institution, federal planning body, or provincial statistical authority) may issue an aggregate query to a Provider or to a domain of Providers. The Provider executes the query within its own systems and returns an aggregate result (count, mean, distribution) protected by differential privacy, mathematically guaranteed noise calibrated to the privacy budget associated with the query class. Individual records are never transmitted. Privacy budgets are accounted centrally through the Metadata Repository so that cumulative disclosure risk across queries is bounded.

15.2 Federated analytics

For analytics that require joint computation across multiple Providers, for example, a national health statistic combining data held separately by provincial health departments, WASL supports federated-analytics orchestration. Each Provider computes its local contribution; an aggregation component combines local contributions into a global result without any Provider disclosing its underlying records. Secure multi-party computation or secure aggregation techniques are used; the specific protocol family is defined in the relevant Technical Specification.

15.3 Confidential computing clean rooms

For analytics workloads that require richer joint processing than aggregate queries or federated analytics allow, WASL specifies a Confidential Computing pattern: multiple Providers deposit encrypted data into a hardware-isolated enclave (Trusted Execution Environment) operating under a jointly signed policy. The enclave executes a pre-registered analytic query and returns only the output; no Provider can read another Provider's input. Clean rooms are a heavyweight capability requiring explicit multi-party agreement; they are reserved for high-value joint analytics and are not a default WASL interaction pattern.

15.4 Analytical integrity

All privacy-preserving analytics transactions are recorded in the TPL with the query identifier, participating Providers, privacy parameters (where applicable), and the hash of the returned aggregate. The TPL record enables later verification that a published statistic derives from the claimed analytic run.

16 Resilience, Observability, and Service Levels

As a foundational DPI component, WASL shall meet resilience and observability standards commensurate with its criticality. This section sets out the architectural requirements; specific quantitative targets are defined in the operating-regulation and the Technical Specifications.

16.1 Availability and continuity

The Central WASL Platform operates in a multi-region active-active posture across geographically separated data centres, with deterministic failover. Client Nodes maintain locally cached metadata sufficient to complete in-progress transactions during brief central-platform outages. The exact targets for Recovery Time Objective and Recovery Point Objective values specified in the operating regulation, aligned with ISO 22301 business continuity standards, are mentioned in the pertinet DNP D series WASL Solution Architecture and Deployment Guide. The specific values are set at a level commensurate with WASL's criticality as national DPI.

16.2 Observability

All WASL components emit structured telemetry conforming to OpenTelemetry semantic conventions. Observability is a first-class concern of the architecture, not an operational afterthought. Metrics, traces, and logs are correlated by transaction identifier, enabling end-to-end tracing of a request across Consumer Client Node, Gateway, routing layer, Provider Client Node, and back.

The Analytics Dashboard component exposes aggregate telemetry to Participating Entities (for their own transactions), to the Operations Centre (for platform-wide monitoring), and to the AI anomaly-detection services (Section 12.1.2).

16.3 Degraded-mode operation

The architecture defines explicit degraded modes for partial outages. If the Central Metadata Repository is temporarily unavailable, Client Nodes continue operating on locally cached metadata (with cache staleness communicated to consuming applications). If the Consent Layer is temporarily unavailable, in-flight consent grants complete but new consent requests queue with a defined service-level commitment for resolution. If the TPL is temporarily unavailable, transactions may proceed with deferred TPL capture (with strict time-bounded commitment to retrospective anchoring). Each degraded mode is explicitly specified rather than emergent.

16.4 Disaster scenarios

The architecture's resilience posture is designed to withstand loss of any single data centre, regional connectivity disruptions, and major cryptographic compromise (through cryptographic agility, Section 14). Loss of scenario planning for broader national-level disruption is addressed in the operational-regulation and is not within scope of this architectural reference.

17 Digital Public Infrastructure Principles Alignment

WASL is conceived and designed as a core component of Pakistan's Digital Public Infrastructure (DPI). This section maps WASL's architecture against the DPI principles most relevant to national data exchange systems. Where alignment is partial or conditional, that is stated explicitly.

17.1 Interoperability

Schema standardization is the primary interoperability mechanism in WASL. All providers within a domain conform to centrally published schemas. A consumer integrating once against the standard getData API can receive data from any conforming provider, across provinces, departments, or sectors, with no additional integration effort. Semantic interoperability is addressed through the Metadata Repository, which maintains authoritative field definitions, identifier standards, and data classification policies agreed across custodian agencies. Legacy system integration is addressed through the Fulfilment Layer in the Client Node, which acts as an adapter between an entity's internal systems and the WASL standard protocol.

17.2 Minimalist and Reusable Building Blocks

The standard API set is deliberately minimal, four standardized endpoints cover the entire data exchange lifecycle without prescribing how consuming applications must use the data. WASL integrates with PAK-ID for citizen authentication and consent notification rather than building a parallel identity system. The Client Node has a lightweight deployment footprint. The TPL records only cryptographic proofs, not data content, the most minimal possible audit record that still provides full non-repudiation. Shared infrastructure (the Metadata Repository, Consent Policy Registry) is consumed by all participants rather than each maintaining their own copy.

17.3 Federated and Decentralised by Design

Data ownership and custody remain with the owning organization. WASL does not aggregate, copy, or store citizen data centrally. Every getData call retrieves data live from the provider's own systems. Client Nodes are deployed within each entity's own infrastructure, giving entities full control over their cryptographic keys, local storage, and fulfilment logic. Trust is decentralized even though routing is not: the three-party cryptographic trust model ensures that neither consumer nor provider needs to extend blind trust to the central platform. Schema metadata and provider public keys are cached locally. The event subscription model allows providers to push change notifications without the central platform acting as a content intermediary. This architecture is well-suited to Pakistan's constitutional structure, where provinces hold autonomous authority over many of the most valuable data domains.

17.4 Security and Privacy by Design

Data blindness is a structural property of WASL, not an administrative control. The central platform is cryptographically prevented from reading the data it routes. Consent is a mandatory blocking step for all personal data access, not an opt-in feature. The TPL provides non-repudiation and transparency as built-in protocol features. A Zero Trust security model is applied throughout: every API call is authenticated, every connection is mutually verified, and every network path is encrypted. Privacy minimization is enforced architecturally: the TPL records only hashes and signatures, never data content.

17.5 Diverse and Inclusive Ecosystem Innovation

The standardized schema and API model means any registered organization, including small startups, insurtech firms, or provincial health authorities, can integrate once and access the same data infrastructure as large financial institutions or federal ministries. The Custom API Marketplace allows providers to publish domain-specific APIs beyond the standard schema. Open standards ensure no participating organization is locked into proprietary tooling. The Developer Portal and sandbox environment lower the barrier to integration testing for new participants.

17.6 Once Only Principle

A citizen's civil registry data, tax filing status, land ownership records, or educational credentials etc., once recorded with the authoritative agency, can be retrieved by any authorized consuming agency through WASL with citizen consent, without the citizen submitting paper documents or filling forms again. Event subscriptions extend this proactively: when a citizen's address changes in the civil registry, subscribed agencies receive automatic notifications, keeping downstream records synchronized without requiring the citizen to update each agency separately.

18 Conformance Criteria

This section defines what it means for a system, implementation, or Participating Entity to be 'WASL-conformant'. Conformance is assessed at three levels: Platform Conformance (for the Central WASL Platform and the Secure Connectivity Layer); Client Node Conformance (for implementations deployed by Participating Entities); and Data Product Conformance (for data served by a Provider).

Implementation-level technical specifications for specific conformance test procedures, protocol bindings, and certification workflows are defined in downstream Technical Specifications published under the DNP-D series. The criteria below define the capability-level requirements that all conformant implementations shall satisfy.

18.1 Platform Conformance Requirements

A conformant Central WASL Platform implementation shall:

- maintain an authoritative Metadata Repository publishing standardized, versioned schemas per data domain, DCAT dataset descriptors, ODRL-expressed consent policies, data classification policies, Provider directory, and Cryptographic Parameter Registry, in machine-readable format accessible to Client Nodes;
- operate an IAM component that issues and validates access tokens using standardized token-based authentication, supports AI model identity binding, enforces role-based and attribute-based access control, and integrates with PAK-ID as the national identity provider;
- enforce a Consent Layer that requires citizen consent as a mandatory, blocking precondition for all access to personal data; generates ODRL-expressed, cryptographically signed consent artefacts; supports multiple consent collection channels including mobile and assisted channels; supports time-bound, purpose-bound, and revocable consent; and includes delegated-consent scaffolding (operationalised per v2 roadmap);
- operate a Transaction Proof Layer that records cryptographic proofs (hashes, digital signatures, timestamps, consent references, bound AI model identifiers) of every transaction without recording data content; supports Pending and Committed lifecycle; and anchors proof batches to a Certificate Transparency-style log conforming to RFC 9162;
- operate an API Gateway that enforces token authentication, consent validation, schema compliance, rate limiting, and cryptographic parameter compliance on all inbound requests, and routes requests to the correct Provider Client Node using metadata-driven routing;
- maintain append-only, tamper-evident audit logs of all platform events, including authentication decisions, consent decisions, API calls, and routing decisions;
- publish and conform all APIs to OpenAPI 3.1 specifications;
- emit structured telemetry conforming to OpenTelemetry semantic conventions;
- support cryptographic agility per Section 14, including the ability to deploy hybrid classical-PQC cryptographic suites. cryptographic agility is a platform-level capability and does not impose specific key-management technology (HSM vs. software) on Participating Entities. Key-management flexibility is defined in Section 14.2;
- never store or process citizen data in plaintext at any platform component.

18.2 Client Node Conformance Requirements

A conformant Client Node implementation shall:

- perform all cryptographic operations i.e. signing, verification, encryption, decryption, within the Entity's own infrastructure perimeter; no private key material shall be transmitted to or processed by the Central WASL Platform;
- protect private keys through mechanisms compliant with Section 14.2 (HSM, software keystore, or managed key service). The choice of mechanism is Entity's responsibility based on risk profile; HSM is mandatory only for TPL endorsers;
- establish encrypted, mutually authenticated site-to-site connectivity to the WASL Secure Connectivity Layer using a PKI-issued certificate, and enforce mutual authentication at the application layer above the network tunnel;
- retrieve and locally cache schema metadata, Provider public keys, cryptographic parameter sets, and consent-policy references from the Metadata Repository, and use locally cached materials for request construction and response verification;
- verify, on every inbound getData response, the Provider's digital signature and the hash of the response data before passing data to the consuming application;
- verify, on every inbound getData request, the Consumer's digital signature, the consent artefact signature, and request integrity before initiating any internal data retrieval;
- validate every outbound response against the relevant published schema before signing and returning it;
- submit a Provider-side proof record to the central TPL simultaneously with every response dispatch;
- implement end-to-end payload encryption (Consumer encrypts with Provider public key) for all data exchanges involving personal or confidential data, ensuring the routing layer cannot read payload content;
- support selective disclosure and ZKP responses for schemas where these capabilities are enabled;
- expose Local APIs that abstract all WASL protocol complexity from internal consuming applications;
- support event subscription and notification using the WASL standard event management protocol;
- for AI-Consumer Client Nodes, include the model identifier in token requests and carry the binding through to TPL records;
- implement cryptographic agility per Section 14.

18.3 Data Product Conformance Requirements

All Data Providers within a domain shall:

- conform their getData responses to the PDA-published JSON Schema for their domain, including all mandatory fields, field types, cardinality, and validation constraints;
- publish a DCAT dataset descriptor for each data product served, including semantic annotations referencing National Data Vocabulary terms;
- validate all responses against the published schema in the Fulfilment Layer before signing and returning them;
- follow the PDA schema versioning and backward compatibility rules, maintaining support for schema versions within the published deprecation timeline;

- accept and fulfil getData requests expressed using the standard request envelope defined in the applicable Technical Specification;
- nominate a Data Product Owner accountable for the accuracy, currency, and documentation of the data product.

18.4 What WASL-conformance does not require

WASL conformance does not require:

- any specific vendor, product, or technology for the Client Node implementation, provided all capability-level requirements above are met;
- any specific container orchestration technology, provided the Client Node operates within a secure, isolated environment within the Entity's infrastructure;
- any specific cryptographic algorithm beyond the parameter sets published in the Cryptographic Parameter Registry;
- any specific transport protocol at the network layer, provided the chosen mechanism satisfies the encrypted, mutually authenticated connectivity requirement;
- participation in the Custom API Marketplace; the standard API set alone is sufficient for core WASL participation;
- use of AI capabilities; AI-native features are optional extensions, not a conformance floor.

19 Entity Onboarding Overview

Any organization, government agency, provincial department, or regulated private-sector entity, may participate in WASL as a Data Consumer, a Data Provider, or both. The onboarding process spans three domains: governance onboarding (registration, approval, legal agreements); trust provisioning (PKI identity and credentials); and technical onboarding (connectivity, Client Node deployment, integration). Only upon completion of all stages is an Entity permitted to participate in production data exchange.

The flow and pertinent details pertaining to Entity Onboarding are covered in the Technical Specifications (TS) document under the DNP-D series.

20 Data Architecture Evolution

WASL's reference architecture is designed not only as a data exchange protocol but as the connective tissue for a broader national data architecture. As WASL matures and participation deepens, it will progressively embody the principles of a Data Mesh architecture and, in later versions, a Data Fabric layer.

20.1 Data Mesh — v1

Under the WASL Data Mesh model, each federal ministry and provincial agency is designated as the data owner and data controller for its respective domain. NADRA owns civil registry data; FBR owns tax data; provincial land record authorities own land record data within their jurisdiction. Each agency is responsible for the accuracy, currency, and governance of its data as a Data Product, an offering with a defined schema, quality guarantees, access controls, documentation, and a responsible product owner.

The WASL Client Node is the self-serve platform in the data mesh model. Official container images are published that any Participating Entity can deploy without requiring bespoke integration support for each deployment. The Developer Portal, sandbox environment, and schema catalogue enable organizations to discover available Data Products, test integrations, and go live independently. Data

Governance is Federated: global standards are defined centrally by PDA but enforced locally at each domain's Fulfilment Layer, balancing domain autonomy with the consistency that Consumers rely upon.

20.2 Data Fabric — v2+

As the ecosystem matures, a Data Fabric layer will be progressively integrated atop the Data Mesh, providing unified metadata intelligence, automated governance enforcement, cross-domain data discovery and cataloguing, and AI and analytics enablement across distributed data sources. The Data Fabric does not change the ownership or custody model, data remains with owning organizations, but provides a semantically rich, automatically governed view across the distributed landscape.

This evolution follows the 'mesh on fabric' pattern observed in mature enterprise data architectures: domain teams own and publish Data Products (mesh principles) while a unified metadata and integration layer (fabric technology) provides seamless connectivity, automated lineage, and cross-domain discovery. WASL's Metadata Repository, enriched with DCAT, ODRL, and SKOS, is the seed of this fabric layer, and its architecture is designed with this evolution trajectory in mind.

Dimension	Data Mesh (v1)	Data Fabric (v2+)
Primary question	Who owns what data and who is accountable?	How do we make distributed data navigable, intelligent, and AI-ready?
Data location	Remains with domain owner	Remains distributed; fabric provides virtual unified view
Data Governance model	Federated — global standards, local execution per domain	Automated enforcement via active metadata and policy layer
Key component	Domain ownership, Data Sharing Agreements, Fulfilment Layer	Active metadata catalogue, lineage tracking, AI-driven discovery
Timeline	v1 — Foundational	v2 / v3 — Mature ecosystem

21 Compliance with International Standards

The WASL architecture is designed to align with the following internationally recognized standards. Specific protocol-version requirements and technical bindings are defined in the Technical Specifications published under the DNP-D series.

Category	Standard	Application in WASL
API Standards	OpenAPI Specification v3.1	All WASL APIs and schema definitions conform to OAS v3.1
Data Modelling	JSON Schema (Draft 2020-12), JSON-LD	All data schemas published in JSON Schema; linked-data expression via JSON-LD
Semantic	W3C DCAT v3, W3C ODRL v2.2, W3C SKOS	Data catalogue, policy expression, controlled vocabularies

Category	Standard	Application in WASL
Identity	W3C Verifiable Credentials Data Model 2.0, W3C DID 1.0	Forward-compatibility of consent artefacts and VC-readiness
Security	OAuth 2.1 / OIDC, JWT (RFC 7519), mutual TLS 1.3 (RFC 8446)	API authentication, authorization, and mutual authentication
Security	PKI / X.509	Digital certificates for connectivity, signing, and encryption
Security	NIST Cybersecurity Framework; ISO/IEC 27001:2022; FIPS 140-3	Cybersecurity posture; ISMS; HSM certification
Cryptography	NIST FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA)	Post-quantum cryptographic algorithms per staged migration (Section 14)
Privacy	ISO/IEC 27018, ISO/IEC 29100	PII protection; privacy framework
Resilience	ISO 22301	Business continuity and disaster recovery
Architecture	NIST SP 800-207 (Zero Trust)	Identity-based access with continuous validation
Observability	OpenTelemetry Semantic Conventions	Structured telemetry across all components
Interoperability	Dataspace Protocol (DSP) 2025-1	v2 roadmap alignment for cross-border data spaces
DPI Reference	GovStack Information Mediator BB	Reference DPI capability model (Section 18)

22 Comparison with Global Data Exchange Layers

WASL's architecture draws on lessons from successful national and regional data-exchange implementations. The following table situates WASL in the global landscape; note that the table describes architectural models as at 2026 and that several of the peer systems are themselves evolving.

Platform	Country / Region	Year	Model	Key Distinguishing Feature
X-Road 7	Estonia / Finland	2001 (v1); 2021 (v7)	Decentralized federation with central identity	Fully federated; platform blind; each participant runs own Security Server; open-source (DPG, MIT)
X-Road 8 (beta 2025)	Estonia / Finland	Q4 2026 (planned)	Federated + Dataspace Protocol	Built on DSP; Gaia-X interoperable; light-context consumer mode; explicit convergence with European data spaces

Platform	Country / Region	Year	Model	Key Distinguishing Feature
APEX	Singapore	2017	Centralised	Government API gateway; platform sees data; focused on G2G service exposure
API Setu / DEPA-AA	India	2020 onwards	Centralized API marketplace; consent-driven data via DEPA / Account Aggregator	Broad private-sector participation; mature consent-artefact ecosystem; data blindness achieved at AA layer only for specific flows
GSB	UAE (federal)	2009	Centralized	Federal G2G service bus; aligned with UAE Digital Data Interoperability Framework (TDRA/DGOV); platform sees data
AD Connect	Abu Dhabi	circa 2020	Centralised, data-exchange via TAMM super-app	Unified Data Exchange platform integrated with TAMM; GovStack-aligned reference implementation; 900+ APIs across 50 AD government entities; platform sees data
Yasser	Saudi Arabia	circa 2018	Centralized	Government integration hub operated by SDGA; platform sees data
MyData	Korea	2020 onwards	Sector-federated, consent-driven	Citizen-controlled personal data; multi-sector deployment; strong consent ledger
Open Finance	Brazil	2021 onwards	Regulatory sector federation	Strong consent and standardised API regime enforced by central bank
Midata / OOTS	Switzerland; EU	2013; 2023	Decentralised; federated + routed	Citizen-controlled; evidence brokering across member states
GAIA-X / DSP	EU (multi-nation)	2020 onwards	Federated data spaces	Sovereign cloud federation; data stays with provider; control plane via DSP; no central data store
WASL	Pakistan	2026	DSP-aligned core + Federated + Data-Blind	DSP 2025-1 adopted as core control-plane protocol from v1; cryptographic data blindness; federated multi-party TPL with cross-institutional endorsement; citizen consent and PAK-ID

Platform	Country / Region	Year	Model	Key Distinguishing Feature
				binding as WASL extensions; AI-native; DPG Client Node

Annex I (Informative): AI-Consumer Governance Pattern (Delegated-Consent Scaffolding)

This annex describes the architectural scaffolding present in WASL v1 to enable a future capability: AI agents acting on behalf of individual citizens under a delegated-consent artefact. The scaffolding is present in v1 but the capability is not operationally enabled; operationalization requires policy resolution and a formal PDA Standards Board decision, planned for the v2 cycle.

I.1 Motivation

A growing class of digital interactions internationally involves AI agents acting on behalf of individual citizens, comparing service offers, completing applications, negotiating with providers, retrieving and consolidating personal records. Enabling this pattern responsibly requires an architectural mechanism by which a citizen can authorize an AI agent to act within a bounded scope, purpose, and duration, with robust revocation, traceability, and liability clarity.

The WASL v1 architecture includes the scaffolding to support this pattern in v2, so that v2 does not require redesigning core components.

I.2 Delegated-consent artefact

A delegated-consent artefact is a citizen-signed consent record of the form: 'Citizen C authorizes agent A (AI system with model identifier M, operated by organization O) to exercise consent decisions within scope S, for purpose P, for duration D, with revocation address R.' Delegated-consent artefacts are signed by the citizen through the Consent Plane, recorded in the Consent Layer, and verifiable by any Provider.

I.3 Agent registration

AI agents operating as citizen-delegated principals are registered in the Metadata Repository with: the model identifier and version; the operating organization; the scope of authority offered to citizens; safeguards certified under DNP-A; liability provisions; and the revocation endpoint.

I.4 TPL capture

Every delegated-agent action is recorded in the TPL with the full delegation chain: citizen identifier (or pseudonym), agent model identifier, operating organization, governing delegated-consent artefact hash, and action-specific parameters. This provides complete traceability in support of later accountability questions.

I.5 Revocation

A citizen can revoke a delegated-consent artefact at any time through the Consent Plane, with immediate effect for future transactions. The revocation propagates to all registered Providers within a bounded service-level commitment. Historical actions taken under the artefact remain auditable via the TPL.

I.6 Why this is deferred to v2

Operationalizing delegated AI agents raises substantive policy questions that deserve explicit resolution rather than architectural presumption:

- Which classes of decisions may an AI agent make on behalf of a citizen? Financial commitments, healthcare consent, legal filings?
- What is the liability chain when an AI agent takes an action adverse to the citizen's interest?
- What safeguards are required against AI agents manipulating citizens into granting broader delegation than intended?

- What accreditation regime governs the AI systems that may offer themselves as citizen-delegated agents?

These questions are not architectural; they are governance questions that require parallel development under DNP-A. The v1 architecture ensures that when these policy questions are resolved, the architectural substrate is ready

Machine-Readable Metadata


```
{
  "@context": "https://standards.pda.gov.pk/schema/dnp",
  "@type": "DNPPublication",
  "identifier": "DNP-D.002 RA (05/2026)",
  "publishingAuthority": {
    "code": "PDA",
    "name": "Pakistan Digital Authority"
  },
  "title": "WASL — Pakistan National Data Exchange Layer: Reference Architecture",
  "titleUrdu": "وصل — پاکستان قومی ڈیٹا ایکسچینج لیئر: حوالہ فنِ تعمیر",
  "series": "D",
  "type": "RA",
  "status": "PUB",
  "maturityLevel": "PILOT",
  "classification": "PUBLIC",
  "normativeReferences": [
    "DNP-X.001 FWK (03/2026)",
    "DNP-D.001 FWK",
    "DNP-A.001 FWK",
    "IETF RFC 2119",
    "IETF RFC 8446",
    "ISO/IEC 27001"
  ],
  "dateApproved": "2026-05-13",
  "dateEffective": "2026-05-13",
  "dateReview": "2029-05-13",
  "jurisdiction": "PK",
  "gazetteReference": null,
  "persistentURL": "https://standards.pda.gov.pk/DNP-D.002"
}
```