

Feedback Block (FDBKB)

Business Requirements Document (BRD)

Version: 5.0

Date: Sep 18, 2026

Status: Final Review

Document Control

Version	Date	Author	Changes
5.0	Sep 20, 2026	Cross-Functional Team	Final Review

Table of Contents

1. Overview..... 1
1.1 Purpose and Scope..... 1
1.1.1 Purpose..... 1
1.1.2 In Scope (Feedback Block)..... 2
1.1.3 Out of Scope (for this stack only)..... 5
1.2 Business Value..... 6
1.2.1 For Citizens..... 6
1.2.2 For Business..... 7
1.2.3 For Government Agencies..... 7
1.2.4 For Government Leadership..... 7
1.3 Success Metrics..... 8
2. Feedback Block Overview..... 11
2.1 FDBKB Overview..... 11
2.1.1 Digital Government Ecosystem..... 11
2.2 FDBKB Roles and Responsibilities..... 13
2.3 Integration with Stacks..... 16
3. Usage Scenarios..... 19
4. Feedback Block Capabilities & Requirements..... 19
4.1 FDBKB Stack Capabilities..... 19
4.1.1 Feedback Form Management Capabilities..... 20
4.1.2 Feedback Collection Capabilities..... 20
4.1.3 Rating and Scoring System Capabilities..... 20
4.1.4 Feedback Processing and Analysis Capabilities..... 21
4.1.5 Dashboard and Reporting Capabilities..... 22
4.1.6 Handoff to UCS Capabilities..... 22
4.1.7 Integration and Interoperability Capabilities..... 22
4.1.8 Citizen Engagement and Transparency Capabilities..... 23
4.1.9 Data Privacy and Security Capabilities..... 24
4.1.10 Administrative and System Management Capabilities..... 24
4.2 FDBKB Functional Requirements..... 25
4.2.1 FR-1: Feedback Form Management Requirements..... 26
4.2.2 FR-2: Feedback Collection Requirements..... 27
4.2.3 FR-3: Rating and Scoring Requirements..... 28
4.2.4 FR-4: Feedback Processing and Analysis Requirements..... 28
4.2.5 FR-5: Dashboard and Reporting Requirements..... 29
4.2.6 FR-6: Handoff to UCS Requirements..... 30
4.2.7 FR-7: Integration and Interoperability Requirements..... 30

4.2.8 FR-8: Citizen Engagement and Transparency Requirements..... 31
4.2.9 FR-9: Data Privacy and Security Requirements..... 32
4.2.10 FR-10: Administration and System Management Requirements..... 33
4.3 Non-Functional Requirements..... 34
4.3.1 Performance..... 34
4.3.2 Availability..... 35
4.3.3 Scalability..... 36
4.3.4 Security..... 37
4.3.5 Usability..... 39
4.3.6 Compliance..... 39
4.3.7 Operational..... 41
5. Solution & Integration Architecture..... 42
5.1 System Architecture Overview..... 42
5.1.1 Architecture Principles..... 42
5.1.2 Architecture Layers..... 43
5.2 Technology Stack..... 45
5.3 Integration Architecture..... 48
5.3.1 DPI Integration Patterns..... 48
5.3.2 API Architecture..... 49
5.4 Deployment Architecture..... 51
6. Data, Security & Compliance..... 53
6.1 Data Management..... 53
6.1.1 Data Model (Indicative)..... 53
6.1.2 Data Lifecycle Management..... 54
6.1.3 Data Privacy Controls..... 55
6.2 Security Architecture..... 56
6.2.1 Defense in Depth..... 56
6.2.2 Security Operations..... 57
6.3 Compliance Framework..... 59
6.3.1 Regulatory Compliance..... 59
6.3.2 Audit and Reporting..... 60
6.3.3 Privacy Impact Assessment..... 61
7. Appendices..... 63
7.1 Glossary..... 63
7.2 API Reference Summary..... 66
7.3 Error Codes..... 68
7.4 Key Definitions..... 70
7.5 External References..... 71

8. Annexure: Example Scenario for Feedback Block..... 73
Scenario 1: Citizen Submitting Service Feedback via MyBhutan (Thimphu Thromde's own Feedback Node)..... 73

1. Overview

1.1 Purpose and Scope

1.1.1 Purpose

This Business Requirements Document (BRD) exists to define a single, standardized way for government agencies across Bhutan to capture, understand, and act on feedback from the citizens and businesses they serve — replacing today's fragmented mix of agency websites, the 1199 helpline, email, chatbots, and the legacy e-Kasel system, none of which share a common data model, a common way of scoring what citizens say, or a reliable path from a complaint to an agency's action. Without a shared standard, every agency has had to build its own collection tool, scoring approach, and escalation path — duplicated effort that also leaves leadership with no comparable, government-wide picture of service quality.

The Feedback Block (FDBKB) is being developed as a Digital Public Infrastructure (DPI) building block to solve this once, as a shared standard rather than a shared system: every agency, ministry, or dzongkhag operates its own Feedback Node, running the same codebase against its own data, connected through Bhutan's National Data Exchange (NDE). This gives citizens and leadership one coherent national feedback block while keeping each agency in control of, and accountable for, its own feedback, without a single national database becoming a single point of failure.

FDBKB's purpose is deliberately bounded: it is authoritative for capturing, validating, scoring, and categorizing what citizens and businesses are telling the government; not for resolving it. Where feedback needs a response, FDBKB hands the item to Unified Client Support (UCS).

Ultimately, this BRD is being developed to give agencies a faster, evidence-based way to improve services; to give citizens a consistent, trustworthy feedback channel through MyBhutan; and to give leadership comparable, national visibility into service quality, supporting Bhutan's Gross National Happiness (GNH) commitment.

The service addresses critical challenges in current feedback collection approaches:

- **Fragmented Collection:** feedback is scattered across agency websites, the 1199 helpline, email, chatbots, and the legacy e-Kasel system, with no shared analysis. Giving

every agency the same Node design, on a shared standard, fixes the fragmentation without needing one shared database.

- **No clear path from feedback to action:** FDBKB hands off to UCS via API (POST /tickets/from-fdbkb) once follow-up is needed.
- **Limited Demographic Analysis:** no reporting by age, gender, dzongkhag, or segment for targeted improvement.
- **Single point of failure:** one central platform means one outage or breach affects every agency at once. Each agency's own Node limits that risk to itself.

1.1.2 In Scope (Feedback Block)

The Feedback Block covers six capability areas, built once as a shared standard and run by every agency's own Node:

1. User Configuration Capabilities

- Feedback Form Management: one standard national form (multi-dimensional ratings, NPS, CSAT, CES, free-text comment) and a shared national taxonomy with agency sub-categories; only service name and category vary per Node.
- Service-Specific Configuration: map forms to services in the Public Service Registry; define collection points within multi-step workflows; handoff thresholds within national bounds (Section 4.1.6).
- Dashboard & Analytics Configuration: role-based dashboards at each Node, plus a federated national dashboard that queries across Nodes; alerting; scheduled reports; KPI definitions standardized nationally.
- User Access & Role Management: role-based access, departmental segregation, approval workflows; each Node manages its own users; the federation layer manages cross-Node roles (national oversight, ACC/RAA/PAC).
- Acknowledgment Configuration: acknowledgment messages and templates only. Everything after acknowledgment belongs to UCS, not FDBKB (see Section 1.1.3).

2. Application Functionalities

- Feedback Collection Engine: citizens submit feedback through MyBhutan, the national government services platform; agency systems (including the 1199 helpline) may submit via the authenticated API or batch upload, with each record carrying an NDI-verified CID

and recorded consent. MyBhutan lists government services and requires NDI/SSO sign-in to submit feedback; once signed in, the citizen sees the full service list. Selecting a service or GA system generates that service's feedback form, with context-aware prompts, progressive collection, and attachments. Submission routes to the Admin of the agency operating that service.

- Rating & Scoring: multi-dimensional ratings, NPS, CSAT, CES, NLP sentiment analysis, computed locally at each Node using a shared model/standard, with federated benchmarking across Nodes.
- Processing & Analysis: categorization/tagging, trend and pattern analysis, comparative analysis, root-cause support, local analysis at each Node; cross-Node comparative analysis performed through federated queries via the NDE rather than a shared database.
- Reporting & Visualization: real-time dashboards per Node, scorecards, demographic reports (age, gender, dzongkhag, gewog, segment), benchmarking, and a federated executive summary assembled from all Nodes.
- Handoff to UCS: where feedback needs action, FDBKB transfers it to Unified Client Support (UCS) via a single API call, so it can be worked and resolved there.
- Integration: MyBhutan as the citizen-facing channel; bidirectional Public Service Registry link, Notification Block, NDI/SSO authentication, UCS case handoff, analytics platform. The DPI services are shared as common federated services used by every Node.
- Citizen Engagement: feedback history, status tracking (mirrored from UCS once a case exists), public aggregated statistics, verification/update — a citizen's personal history may span multiple Nodes and is assembled via federated lookup keyed to their NDI identity.

3. Data Lifecycle Management

- Collection & capture with validation, metadata, transaction linkage, consent capture — performed at the originating Node.
- Storage: federated per-agency repositories (no single national repository); each Node stores its own agency's feedback with time-series storage, pseudonymization, versioning, archival; the NDE catalog indexes which Node holds which service's feedback for discovery purposes.

- Processing & enrichment: validation, dedup, text normalization, sentiment scoring — local to each Node, using shared, standardized models.
- Analysis: aggregation by service/department/time/demographics, statistical significance testing, trend forecasting, text mining — local analysis at each Node, with cross-Node aggregation performed by the federation layer through consented, standardized federated queries rather than a live central database join.
- Retention & archival: configurable retention by type/sensitivity, automated archival, secure deletion — governed by a shared national policy, enforced independently by each Node.
- Privacy & anonymization: minimization, automated anonymization for public reporting, pseudonymization, right to erasure per the Information, Communications and Media Act of Bhutan 2018, the National Digital Identity Act of Bhutan 2023, and the National Data Governance Framework 2025 — each Node is the data controller for the feedback it holds.

4. Communication & Data Exchange

- API Architecture: REST endpoints for submit/retrieve/status/analytics/handoff, implemented identically by every Node; WebSocket/SSE for live per-Node dashboards; async messaging (e.g., Kafka/RabbitMQ) local to each Node plus a federation-wide event topic; JSON/CSV/Excel/PDF/XML exchange formats.
- Integration Protocols: National Data Exchange (NDE, WSO2-based) as the primary federation fabric for secure cross-agency exchange, service discovery, and federated aggregation across Nodes; CloudEvents compliance; publish-subscribe patterns; webhook support; Public Service Registry bidirectional integration.
- External Systems: MyBhutan, which transparently routes citizens to the correct Node behind the scenes; BI tools; third-party survey platforms; Notification Block.

5. Security Controls

- Authentication: OAuth 2.0/OIDC via NDI (a shared, common service across all Nodes), MFA for admins, RBAC, JWT/API-key auth, service-to-service auth.
- Data Protection: TLS 1.3 in transit, AES-256 at rest, field-level encryption for sensitive content, privacy-by-design, consent management — enforced independently, but to a common standard, at every Node.

- Access Control & Audit: fine-grained access, segregation of duties, immutable audit trail (via the shared Journaling Block), periodic access review at each Node.
- API Security: rate limiting, input validation, injection protection, gateway controls, OWASP API Top 10 — each Node runs its own API gateway conforming to a shared security baseline.
- Threat Protection: DDoS protection, bot/spam detection, at each Node
- Fraud Prevention: duplicate detection, coordinated-pattern detection, CAPTCHA, reputation scoring — locally at each Node, with cross-Node duplicate signals shared through the federation layer.

6. Other Key Considerations

- Demographic Analytics: reporting by age group, gender, dzongkhag/gewog, rural/urban, service type — produced locally and federated nationally.

1.1.3 Out of Scope (for this stack only)

- Complex Grievance Case Management: formal, legally-timed adjudication remains with the (successor-to-e-Kasel) Grievance Management System. FDBKB is for service-improvement feedback, not formal complaints requiring legal process.
- Petition Management: signature collection, endorsement campaigns, and parliamentary submission stay with the dedicated Petition Service (PETB).
- Survey Creation and Management: complex branching, quota management, and survey design belong to the Survey Block (SRVB). FDBKB provides simple feedback forms only.
- Social Media Monitoring: direct platform integration for social listening is handled by separate tooling, not FDBKB.
- 1199 Call Center CRM: call routing, IVR, agent desktop and telephony stay in the CRM system, though 1199-sourced feedback reaches the relevant agency Node via API, submitted by the agent with the citizen's CID and recorded consent.
- Service-Specific Business Logic: agency-specific approval hierarchies and operational processes remain with individual agencies.

- Advanced AI/ML Model Development: FDBKB includes sentiment analysis and basic pattern recognition only; sophisticated predictive-governance modeling is out of scope for this stack.
- Financial Transaction Processing: FDBKB does not handle refunds, compensation, or any payment transactions related to service failures.
- Operating a Single National Feedback Database (explicitly out of scope under this concept)**:** GovTech Bhutan does not operate one central FDBKB datastore for all agencies; it operates the federation standard, the shared common services (NDI, NDE, JLB, NTFB integration contracts), and a national aggregation view built from consented, federated queries.
- Case Management and Agent Tooling: everything from ticket lifecycle through case closure belongs to UCS. FDBKB hands off a validated, scored feedback item via API and stops there; it keeps no case-handling queue, escalation rules, or agent-facing tools of its own; its only queue is the moderation queue for flagged submissions (FR-4.8, FR-9.6, FR-10.3).

1.2 Business Value

1.2.1 For Citizens

- Voice and Empowerment through MyBhutan, regardless of which agency's Node ultimately handles the feedback.
- Transparency and Accountability: instant acknowledgment from FDBKB, followed by SLA-governed status updates, response, and closure; citizens see one continuous experience even though two different blocks are involved behind the scenes.
- Improved Service Quality from systematic issue identification at the Node and disciplined, SLA-tracked resolution through UCS.
- Personal Feedback History for authenticated users, aggregated across every agency Node they have interacted with, via federated lookup keyed to their NDI identity.
- Local accountability: citizens can be confident their feedback data is held by the agency they actually dealt with, not by an unrelated central authority.

1.2.2 For Business

- Targeted feedback on licensing, permits, taxation, and regulatory processes reduces compliance friction.
- Lower transaction costs as processes improve.
- Predictable service standards from published quality metrics, comparable across agencies through federated benchmarking.
- Sector-specific, anonymized trend visibility.
- Contributes to overall digital-ecosystem improvement.

1.2.3 For Government Agencies

- Evidence-based, proactive service improvement using real-time sentiment and trend data from an agency's own Node.
- Objective performance measurement (NPS, CSAT, CES) over time and across services, with the option to benchmark against federated national data.
- Resource optimization via demographic and geographic analysis.
- Staff performance insight from service-delivery feedback.
- Cross-agency learning through federated benchmarking (rather than one shared report).
- Compliance/audit support via Journaling Block (JLB) integration, common to every Node.
- Reduced support burden as recurring issues are fixed.
- Direct data ownership: each agency controls its own feedback data, its retention, and who may access it, subject to shared national standards.
- No duplicated case-management build: escalation, response, and resolution are handled by UCS rather than reimplemented inside FDBKB.

1.2.4 For Government Leadership

- Federated executive dashboards for real-time, government-wide satisfaction visibility, assembled on demand from every agency Node rather than read from one database.
- Policy impact assessment via before/after feedback analysis across the federation.
- Transparency and public trust through published aggregated statistics.
- Contribution to Gross National Happiness (GNH) wellbeing measurement in service delivery.

- Budget justification tied to satisfaction improvement.
- International benchmarking using standardized NPS/CSAT metrics.
- Cross-government coordination on cross-cutting issues, without requiring any single ministry to hold all agencies' data.

1.3 Success Metrics

The Feedback Block will be evaluated against the following key performance indicators. Metrics marked (federated) are computed by aggregating results from individual Nodes, and may carry a defined maximum staleness rather than being instantaneous.

Category	Metric	Target	Measurement Method
Adoption & Usage	Number of feedback entries received per month, summed across all Nodes (federated)	10,000+ by Month 6; 50,000+ by Year 1	Per-Node analytics rolled up via federated query
	Percentage of government services with an active Feedback Node collection point	80% of onboarded agencies' services by Year 1; 80% of all government services by Year 3	Service catalog integration tracking (per Node, verified via NDE)
	Number of government agencies operating their own Feedback Node	8+ agencies by Year 1; 50+ agencies by Year 3	NDE service registry / node registration records
Quality & Responsiveness	Average time to acknowledge feedback (FDBKB-owned)	≤ 60 seconds, automated (FR-2.5)	Per-Node automated acknowledgment timestamp tracking
	Percentage of feedback requiring follow-up that is successfully handed off to UCS as a valid ticket	≥ 99% within 3 seconds of the handoff decision	FDBKB handoff-API success metrics, cross-checked against UCS ticket-creation confirmations

	Time to respond, escalation compliance, and closure rate within SLA (UCS-owned — retained here only for traceability)	See UCS BRD Section 1.6 (First Contact Resolution Rate, P1/P2/P3 Response Time, SLA Breach Rate)	Tracked entirely in UCS; FDBKB surfaces the resulting status to citizens but does not independently measure it
Service Improvement	Average satisfaction score per service (5-point scale)	≥ 3.5 baseline; +10% annually	Weighted average computed at the owning Node from FDBKB rating data
	Net Promoter Score (NPS) for government services, federated	≥ 40 by Year 1	NPS calculated per Node, aggregated via federated query
	Number of service improvements implemented based on feedback (UCS-tracked action/resolution outcome, sourced from linked UCS tickets)	75+ improvements by Year 1	UCS action/resolution records, cross-referenced to originating FDBKB feedback_id
	Percentage of recurring issues resolved	≥ 70%	UCS resolution tracking; FDBKB contributes the pattern-detection signal that flags the recurrence (FR-4.3)
Technical Performance	System availability (uptime), per Node	≥ 99.5% Year 1, ≥ 99.9% steady state	Per-Node infrastructure monitoring (an outage at one Node does not count against others)

	Feedback submission API response time	≤ 500ms (95th pct)	Per-Node application performance monitoring
	Sentiment analysis accuracy	≥ 80% (Dzongkha), ≥ 85% (English)	Manual validation sampling across a cross-section of Nodes
	Federated query staleness (national dashboards)	≤ 15 minutes	Federation layer sync/query timestamp monitoring
Data Quality	Percentage of duplicate feedback successfully detected	≥ 95%	Per-Node duplicate detection algorithm validation
	Spam/bot feedback filtering accuracy	≥ 95%	Per-Node CAPTCHA and pattern detection effectiveness
	Feedback categorization accuracy	≥ 88%	Manual verification sampling
Integration	Number of Feedback Nodes successfully federated via NDE	8+ agencies by Year 1, scaling toward the full 50+ agency target	NDE registration and conformance testing records
	API usage by external systems (requests/month), summed across Nodes	20,000+ requests/month by Year 1	Per-Node API gateway analytics
	Federated query success rate	≥ 99.9%	NDE gateway and federated query broker logs
Governance & Compliance	Audit trail completeness (% of events logged to JLB)	100%	JLB integration monitoring, per Node

	Privacy compliance incidents	0 major incidents	Security audit and incident tracking, per Node
	Data retention policy compliance	100%	Automated retention rule execution
	Successful completion of security assessments	Pass all quarterly assessments	Third-party security audit reports, per Node
User Satisfaction	Citizen satisfaction with feedback process itself	≥ 80% satisfied	Meta-feedback survey
	Agency user satisfaction with feedback insights and tools	≥ 85% satisfied	Agency user survey
	Percentage of citizens who see their feedback leading to change	≥ 60%	Closed-loop tracking survey, drawing on UCS closure/resolution data linked back to the originating feedback_id

Measurement Frequency: Real-time metrics (system availability, API response times, submission volumes) are captured per Node; weekly, monthly, quarterly, and annual metrics are aggregated nationally through the federation layer on the cadence shown above.

2. Feedback Block Overview

2.1 FDBKB Overview

2.1.1 Digital Government Ecosystem

FDBKB is a core Communication building block within Bhutan's Digital Public Infrastructure (DPI), implemented as a decentralized, federated network of agency-operated Feedback Nodes rather than a single national system, providing the primary mechanism for systematic collection, validation, scoring, and analysis of citizen and business feedback on government services, institutions, personnel, and governance practices. Case resolution — routing to a responsible officer, escalation, response, action, and closure — is out of FDBKB's boundary and belongs to

the Unified Client Support (UCS) block, which FDBKB integrates with as a first-class, Critical dependency (see below and Section 2.3).

Position in DPI Ecosystem: FDBKB sits upstream of citizen support: it is where sentiment and structured/unstructured feedback enter the DPI, and UCS is where anything requiring follow-up is worked to resolution. Citizens reach FDBKB through MyBhutan, and each agency's Feedback Node integrates with:

Citizen Channel: MyBhutan: the national government services platform and the single citizen-facing channel for FDBKB. It is not a DPI building block that FDBKB integrates with as a peer; it is the front door through which citizens reach every Node. Citizens browse the government service list on MyBhutan and sign in via NDI/SSO to submit feedback. Selecting a service or GA system generates that service's feedback form, submitted to the Admin of the agency operating that service's Node. MyBhutan also provides personal history and action tracking via federated lookup.

Foundation Layer Integration:

- National Digital Identity (NDI): authenticates every MyBhutan submission and provides the common identity standard that lets a citizen's feedback be linked across Nodes without those Nodes sharing a database.
- National Data Exchange (NDE): Bhutan's operational interoperability platform (WSO2-based, with an embedded API gateway) — this is the federation fabric: it connects every agency's Node for secure feedback exchange, service discovery, and federated query/aggregation.
- Journaling Block (JLB): immutable audit trails for all feedback transactions, responses, and actions at every Node, for compliance and non-repudiation, without centralizing the underlying feedback content itself.

Service Layer Integration:

- Public Service Registry (PSR): links feedback to registered services and syncs quality scores back to the catalog from each Node.
- Notification Block (NTFB): sends the acknowledgment to citizens (FR-2.5). Critical-feedback alerts to agency staff appear on the Node's alert dashboard (Section 4.1.5).

Case Resolution Layer Integration:

- Case Resolution: where an item needs action, FDBKB transfers it to Unified Client Support (UCS) for resolution.

Key Architectural Characteristics:

- Node-Based Deployment: every participating agency, ministry, or dzongkhag administration deploys its own Feedback Node — a standardized, common codebase running the same capabilities described in Section 4 — rather than depending on one shared national instance
- Single Citizen Channel: MyBhutan is the only citizen-facing channel; citizens browse services and submit feedback after NDI/SSO sign-in. The only other entry routes are the authenticated API and batch upload used by agencies (including 1199 agents), and every record must carry an NDI-verified CID and recorded consent. No anonymous mode; no offline collection.
- Real-Time Intelligence: NLP sentiment analysis, pattern recognition, and trend forecasting computed locally at each Node, using a shared national model/standard, with federated correlation across Nodes.
- Split-Ownership Closed-Loop Feedback: submission, acknowledgment, scoring, and categorization are owned end-to-end by FDBKB at the responsible agency's Node; response, action, and closure are owned end-to-end by the linked case once one is created (Section 4.1.8).
- Localized Experience: full Dzongkha language support with culturally appropriate collection methods, standardized across every Node.

2.2 FDBKB Roles and Responsibilities

The Feedback Block supports multiple user roles with distinct responsibilities and access levels. Most roles are exercised locally, within a single agency's Feedback Node; a small number operate at the federation level, spanning Nodes.

Role	Description	Key Responsibilities

Citizen/Business User (MyBhutan)	Any citizen or business known to NDI after signing in on MyBhutan, individuals and businesses alike using the same NDI-backed session.	Browse services on MyBhutan; sign in via NDI/SSO; select a service or GA system to generate and submit its feedback form; track status; access personal history. Business users additionally see sector-specific aggregated trends for their industry. No anonymous path exists for either.
Agency Node Administrator (Focal Point)	The agency staff member — or, at larger agencies, small team — responsible end-to-end for that agency's own Feedback Node: the single destination for feedback submitted about that agency's services. Smaller agencies typically have one focal person covering everything below; larger agencies may split these duties across specialists.	Receive and review incoming feedback; categorize and moderate; review the system's handoff recommendation against the Section 4.1.6 criteria and confirm or override it; configure agency sub-categories and handoff thresholds within national bounds (Section 4.1.6); manage local users and permissions; track NPS/CSAT/CES and linked resolution outcomes; monitor Node performance and health; maintain DPI integrations; manage local retention. Responding to citizens happens through UCS, not FDBKB.
Node System Administrator	IT staff responsible for the technical operation, maintenance, and security of a single agency's Feedback Node infrastructure — distinct from the Agency Node Administrator's functional/configuration role above.	Monitor availability/performance; manage infrastructure/deployment; configure security controls; backup and disaster recovery; apply patches; monitor API usage/rate limits; manage integration endpoints.

Federation Administrator (GovTech, national level)	GovTech staff responsible for the federation standard, not for any single agency's data: the common codebase, data schema, security baseline, NDE conformance testing, and onboarding of new agency Nodes.	Define and publish the national data schema, category taxonomy, and API contract every Node must conform to; certify new Nodes before they join the federation; maintain the shared common services (NDI/NDE/JLB/NTFB integration contracts); monitor federation-wide health and conformance (not individual agencies' feedback content); coordinate national policy on retention, security, and handoff criteria; do not hold administrative access to any individual agency's feedback data by default
Analytics Analyst	Data analysts who analyze feedback trends, generate insights, and create reports, either within a single Node or, for cross-agency analysis, through federated queries.	Access feedback data for analysis (within their authorized Node, or via federated query where authorized); generate trend analysis and pattern recognition reports; create demographic analysis reports (age, gender, location); perform sentiment analysis validation; produce trend forecasts using the shared national models; create executive dashboards and scorecards; conduct comparative benchmarking analysis across Nodes; generate ad-hoc custom reports

Government Leadership	Senior government officials and decision-makers who use federated feedback insights for strategic planning and policy decisions.	Access federated executive dashboards assembled across Nodes; view cross-agency satisfaction metrics; monitor policy impact through feedback analysis; review anti-corruption and transparency feedback within the standard federated data; track GNH-related satisfaction indicators; make evidence-based policy decisions. Agency responsiveness is monitored through a separate leadership dashboard, which FDBKB's own dashboards may link out to rather than duplicate.
Monitoring & Oversight Authority	Independent oversight bodies (e.g., ACC, RAA, PAC) monitoring government responsiveness and service quality, with authorized federated access across Nodes for oversight purposes.	Access aggregated feedback data across Nodes for oversight, under a defined, audit-logged federated-access protocol; identify systemic governance issues from feedback content and sentiment; review corruption perception feedback within the standard federated data available to oversight roles.

2.3 Integration with Stacks

Component	Integration Purpose	Integration Type	Protocol	Criticality
MyBhutan	Citizen-facing channel. Lists government services and their published ratings and reviews; provides NDI-authenticated feedback submission routed to the correct Node and displays personal	Bidirectional	REST API, SSO (OIDC), WebSocket, Widget Embed /	Critical (3): the only citizen-facing channel

	feedback history. Case status, response, and resolution shown there are mirrored live from UCS via the linked ticket reference, not tracked independently by FDBKB.		JavaScript SDK	
Unified Client Support (UCS)	Feedback requiring action is transferred to UCS for resolution.	Node → UCS (ticket creation, one-way per feedback item); UCS → Node (status callback/webhook, one-way)	REST API (OpenAPI 3.0), Webhook callback, CloudEvents	Critical (3) — without this, feedback that needs action has nowhere to go
Journaling Block (JLB)	Creates immutable, tamper-evident audit trails for all feedback submissions, status changes, and handoff events at every Node. Ensures non-repudiation and compliance without centralizing feedback content itself. UCS logs its own resolution-related events to JLB separately, cross-referenced by ticket ID.	Node → JLB (one-way, per Node)	Event Stream Kafka CloudEvents	Critical (3) Mandatory for compliance
National Data Exchange (NDE)	The core federation mechanism. Enables secure, standardized data exchange across agency Nodes; provides service and Node discovery; supports federated query and cross-agency aggregation; ensures data sovereignty by keeping data at its Node of origin.	Bidirectional	REST/SOAP over NDE (WSO2)	Critical (3) — without this, independent Nodes cannot function as one coherent national service

National Digital Identity (NDI)	Provides authentication for every submission, common to every Node. Enables citizen identification and links a citizen's history across Nodes without those Nodes sharing a database. Supports consent management.	Node → NDI (consumed by each Node)	OIDC, OAuth 2.0, JWT	Critical (3) Every submission requires NDI authentication (FR-7.2)
Public Service Registry (PSR)	Links each Node's feedback to registered services. Retrieves service metadata for context. Publishes service quality scores back to the catalog from each Node. Enables service discovery for feedback forms.	Bidirectional	REST API (via NDE/API Ecosystem)	Important (2) Service linking and discovery
Notification Block (NTFB)	Sends acknowledgment notifications to citizens from the owning Node only (FR-2.5, FR-7.3, FR-8.2). All subsequent notifications are dispatched by UCS, not FDBKB, once a ticket exists.	Node → NTFB (one-way, per Node)	REST API Event-driven	Important (2) Enhances responsiveness

Integration Patterns:

- **Synchronous Integration:** REST APIs used for real-time feedback submission, retrieval, ticket handoff, and status updates — within a Node, across Nodes via the NDE, or from a Node to UCS.
- **Asynchronous Integration:** Event-driven patterns (e.g., Kafka, CloudEvents) for audit logging, analytics data streaming, and workflow triggers, per Node and federation-wide; UCS ticket-lifecycle events consumed the same way.
- **Federated Integration (the defining pattern of the decentralization concept):** NDE and PSR protocols for secure cross-agency data exchange, Node discovery, and federated

query/aggregation — used in place of any direct database-to-database link between agencies.

- **Widget Embedding:** JavaScript SDK and iframe embedding for MyBhutan integration, resolving to the correct Node with minimal coupling.
- **Real-Time Streaming:** WebSocket and SSE for live per-Node dashboard updates and notification delivery.

3. Usage Scenarios

Refer to Section 8 for detailed example scenarios

4. Feedback Block Capabilities & Requirements

This section outlines the key capabilities of the Feedback Block, organized by functional area. Unless stated otherwise, each capability is implemented identically at every agency's Feedback Node, using a shared national codebase, schema, and standard, so citizen experience stays consistent across agencies.

4.1 FDBKB Stack Capabilities

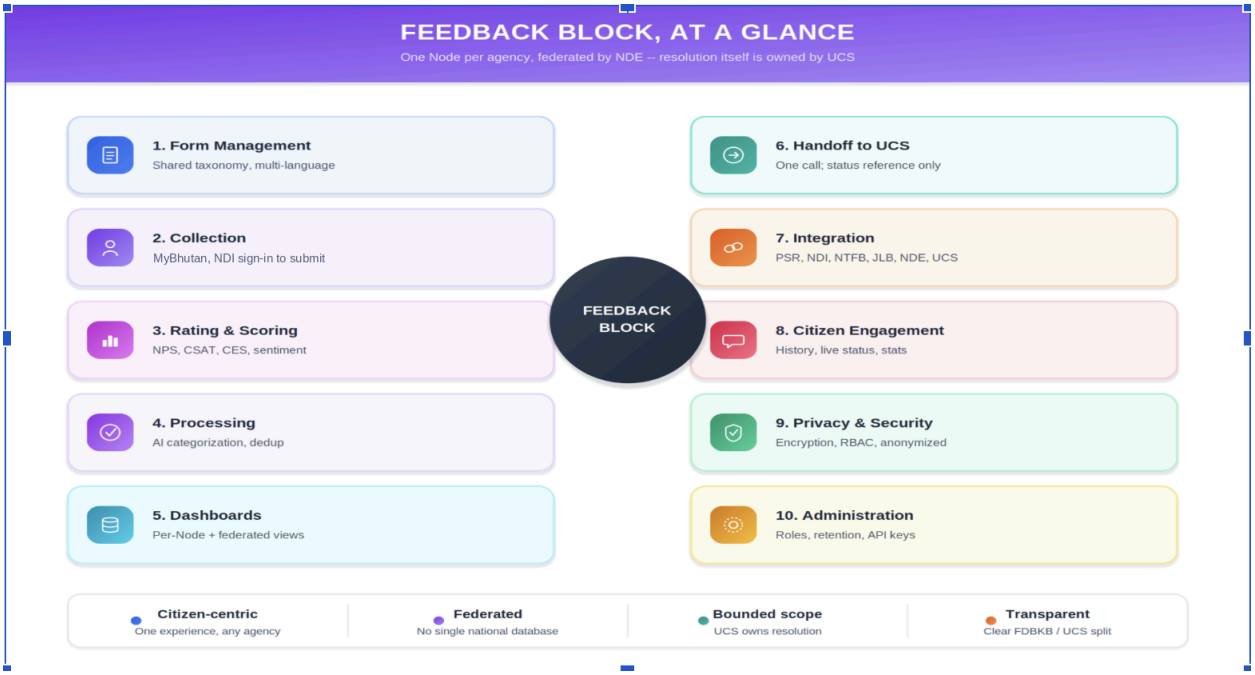


Figure 1: FDBKB Capabilities Diagram

4.1.1 Feedback Form Management Capabilities

- Standard National Form: every Node uses one fixed national feedback form, defined and versioned by the Federation Administrator: multi-dimensional ratings, NPS (0–10), CSAT (1–5), CES (1–7), and a free-text comment, in Dzongkha and English. Only the service name and category vary per Node; Nodes do not design or edit forms (FR-1.1).
- Feedback Types: service satisfaction ratings, suggestions, and complaints (handed off per Section 4.1.6); structured (ratings) and unstructured (free text). Formal grievances and surveys are out of scope (Section 1.1.3).
- Multi-Dimensional Rating System: standardized national dimensions (overall satisfaction, timeliness, ease of use, staff behavior, information clarity, accessibility, value for money) for comparability across Nodes.
- Trigger Configuration: the form is offered on MyBhutan after a transaction and as general feedback once signed in. A shared national frequency limit prevents feedback fatigue.
- Category and Taxonomy Management: a shared national taxonomy (service quality, process efficiency, user experience, accessibility, staff behavior, technical issues, policy concerns, transparency issues) maintained by the Federation Administrator; Nodes may add agency-specific sub-categories (FR-1.2).
- Form Versioning: the national form is versioned centrally and rolled out to all Nodes on the coordinated release train (NFR-7.1).

4.1.2 Feedback Collection Capabilities

- MyBhutan Integration: MyBhutan lists government services; submitting feedback requires NDI/SSO sign-in. Selecting a service generates that service's feedback form with context-aware pre-population and attachment support, and submission routes to the Admin of the agency operating that service. A general (non-transactional) feedback option is also available once signed in. API-based submission and CSV/Excel batch upload; every record must carry the citizen's NDI-verified CID and consent, or it is rejected (FDBKB-1003). Standard responsive design covers mobile browsers — no separate offline app.

4.1.3 Rating and Scoring System Capabilities

- Multi-Dimensional Ratings: collected and computed at each Node across standardized dimensions (overall satisfaction, timeliness, ease of use, staff behavior, information

clarity, accessibility, value for money); weighted scoring per dimension; aggregate scores using a nationally standardized calculation method.

- Net Promoter Score (NPS): standard 0–10 question at each Node; Promoters/Passives/Detractors classification; NPS trend tracking locally, with federated benchmarking against other Nodes.
- Customer Satisfaction Score (CSAT): 1–5 scale at each Node; percentage satisfied vs. total; trending and federated comparison.
- Customer Effort Score (CES): 1–7 ease-of-completion score at each Node; correlation with completion rates and abandonment.
- Sentiment Analysis: NLP classification (Positive/Neutral/Negative/Mixed/Highly Critical/Escalation-prone), computed locally at each Node using a shared national model, with intensity scoring (−1 to +1), emotion tags, and Dzongkha/English support.
- Rating Validation and Quality Control: each Node detects suspicious rating patterns and bot activity locally; reputation scoring; human verification for flagged submissions.

4.1.4 Feedback Processing and Analysis Capabilities

- Automated Categorization and Tagging: ML-based multi-label categorization performed locally at each Node against the shared national taxonomy.
- Automated Feedback Validation: each Node validates incoming feedback for mandatory fields, duplicate submissions, channel/source authenticity, data integrity, and spam/fraud signals before onward processing.
- Duplicate and Related-Case Detection: text-similarity and same-user/service/time-window matching within a Node; cross-Node duplicate signals (the same citizen contacting multiple agencies about a linked issue) are surfaced via the federation layer rather than a shared duplicate-detection database (FR-4.4).
- Trend Analysis and Pattern Recognition: time-series analysis of feedback volume and sentiment at each Node; seasonal pattern detection; cross-Node correlation performed via federated query rather than a shared database join.
- Comparative Analysis: cross-service comparison within a Node; inter-agency benchmarking, geographic comparison, and demographic segment comparison performed via federated query across Nodes, not by querying one shared database.

- Root Cause Analysis: correlation of low ratings with specific service attributes at each Node; recommendation engine for improvement priorities based on impact and frequency.
- Text Mining and Insights Extraction: keyword frequency analysis, phrase clustering, and theme identification, performed locally at each Node, in both supported languages.

4.1.5 Dashboard and Reporting Capabilities

- Per-Node: real-time operational dashboard, service performance monitoring, trend lines, CSAT/NPS reporting, scheduled reports, service quality scorecards, volume/trend reports, demographic reports, custom report builder, and an alert dashboard for critical feedback and anomalies.
- Federated (assembled across Nodes by the federation layer, per the staleness target in Section 1.3): agency performance comparison, geographic visualization, comparative benchmarking, executive summary, and a national alert view for GovTech and oversight roles.

4.1.6 Handoff to UCS Capabilities

The Node's Analysis Engine recommends handoff when any of the following apply: (a) any rating dimension is below 4/5 and sentiment is Negative or Mixed; (b) sentiment is Highly Critical or Escalation-prone, regardless of rating; (c) the category is staff behavior, transparency issues, or policy concerns; (d) the item matches a recurring-issue pattern (FR-4.3). The Agency Node Administrator confirms or overrides the recommendation. On confirmation, the Node transfers the item to UCS via a single API call, carrying the feedback content and its rating/sentiment/category, and UCS takes it from there. These thresholds are national defaults set by the Federation Administrator; Nodes may tighten but not loosen them.

4.1.7 Integration and Interoperability Capabilities

- Service Catalog Integration/Public Service Registry (PSR): each Node syncs bidirectionally with PSR for its own services.
- Identity Management Integration (NDI/SSO): a shared, common service used by every Node for authentication. (See FR-7.2.)

- Notification Block Integration (NTFB): used by each Node only for the initial acknowledgment (FR-2.5, Section 4.1.8); every notification from that point forward is dispatched downstream, once a case exists.
- Journaling Block Integration (JLB): a shared, common audit-logging service; each Node logs its own submission/handoff events, cross-referenced by case ID against the downstream system's own audit stream, not merged into one.
- National Data Exchange (NDE) Integration — the core federation mechanism: every Node registers itself and its services in the NDE catalog; NDE brokers secure cross-agency exchange, federated aggregation, and service/Node discovery; standardized API contracts ensure any authorized consumer can query any Node the same way.
- Unified Client Support (UCS): receives feedback transferred under Section 4.1.6 for resolution.

4.1.8 Citizen Engagement and Transparency Capabilities

- Personal Feedback History: assembled for the citizen via federated lookup across every Node they have submitted feedback to, keyed to their NDI identity, rather than read from one shared database.
- Feedback Status Tracking: tracking-ID-based status checking works against the specific Node that issued the tracking ID; where a case exists, the status shown is read live from the case system, and FDBKB does not maintain a second, independent status value once a case exists.
- Response Notifications: the initial acknowledgment is delivered by the owning Node via the shared Notification Block; all subsequent notifications (status change, response received, resolution, closure, CSAT request) are delivered downstream, not re-sent by FDBKB.
- Public Feedback Statistics: national-level public statistics are assembled by the federation layer from every participating Node's anonymized, aggregated output; no individual Node's raw feedback is exposed publicly.
- Improvement Impact Visibility: each Node publishes its own before/after stories, sourced from the downstream resolution/action records linked back to the originating feedback; a national showcase aggregates these across Nodes.

- Feedback Verification and Updates: performed against the owning Node for the feedback record itself; adding further information to an active case is performed in the case's own thread (Section 4.1.6) once one exists.
- Community Feedback Insights: aggregated across relevant Nodes via federated query, with no individual feedback exposed.

4.1.9 Data Privacy and Security Capabilities

- Privacy-by-Design Features: enforced independently, to a shared national standard, at every Node.
- User Data Rights: exercised against whichever Node(s) hold a citizen's data; a citizen exercising a right to erasure, for example, may need this action at more than one Node if they have submitted feedback to multiple agencies, coordinated via the federation layer.
- Access Control and Authorization: RBAC/MFA/ABAC enforced independently at each Node; no Federation Administrator has blanket access to any individual Node's raw feedback by default — cross-Node access is itself an access-controlled, audited capability.
- Encryption and Data Protection: TLS 1.3, AES-256, field-level encryption, and key rotation, enforced identically at every Node.
- Audit and Compliance: audit logging to JLB from every Node; access logs, change history, and incident response procedures maintained per Node, conforming to a shared national compliance standard.
- Fraud Detection and Prevention: duplicate/spam/bot detection at each Node, with coordinated-campaign detection benefiting from cross-Node pattern sharing via the federation layer.

4.1.10 Administrative and System Management Capabilities

- User and Role Management: create/manage local user accounts, role assignment, bulk import — performed locally, per Node, by that agency's Agency Node Administrator.
- Configuration Management: Acknowledgment templates, handoff criteria (Section 4.1.6), and integration endpoints, configured per Node within bounds set by the Federation Administrator.

- Content Moderation: profanity filtering, flagging queue, moderator actions and audit trail, per Node.
- Data Retention and Archival: configurable retention by feedback type, automated archival and secure deletion, per Node within shared national policy.
- API Management: API key management, rate limiting, versioning, and deprecation notices, per Node, to a shared national API contract.

4.2 FDBKB Functional Requirements

This section defines the detailed functional requirements for the Feedback Block, organized by capability area. Requirements use MoSCoW prioritization (MUST/SHOULD/COULD). Unless a requirement explicitly says “federated” or “Federation Administrator,” it applies identically at every agency’s Feedback Node.

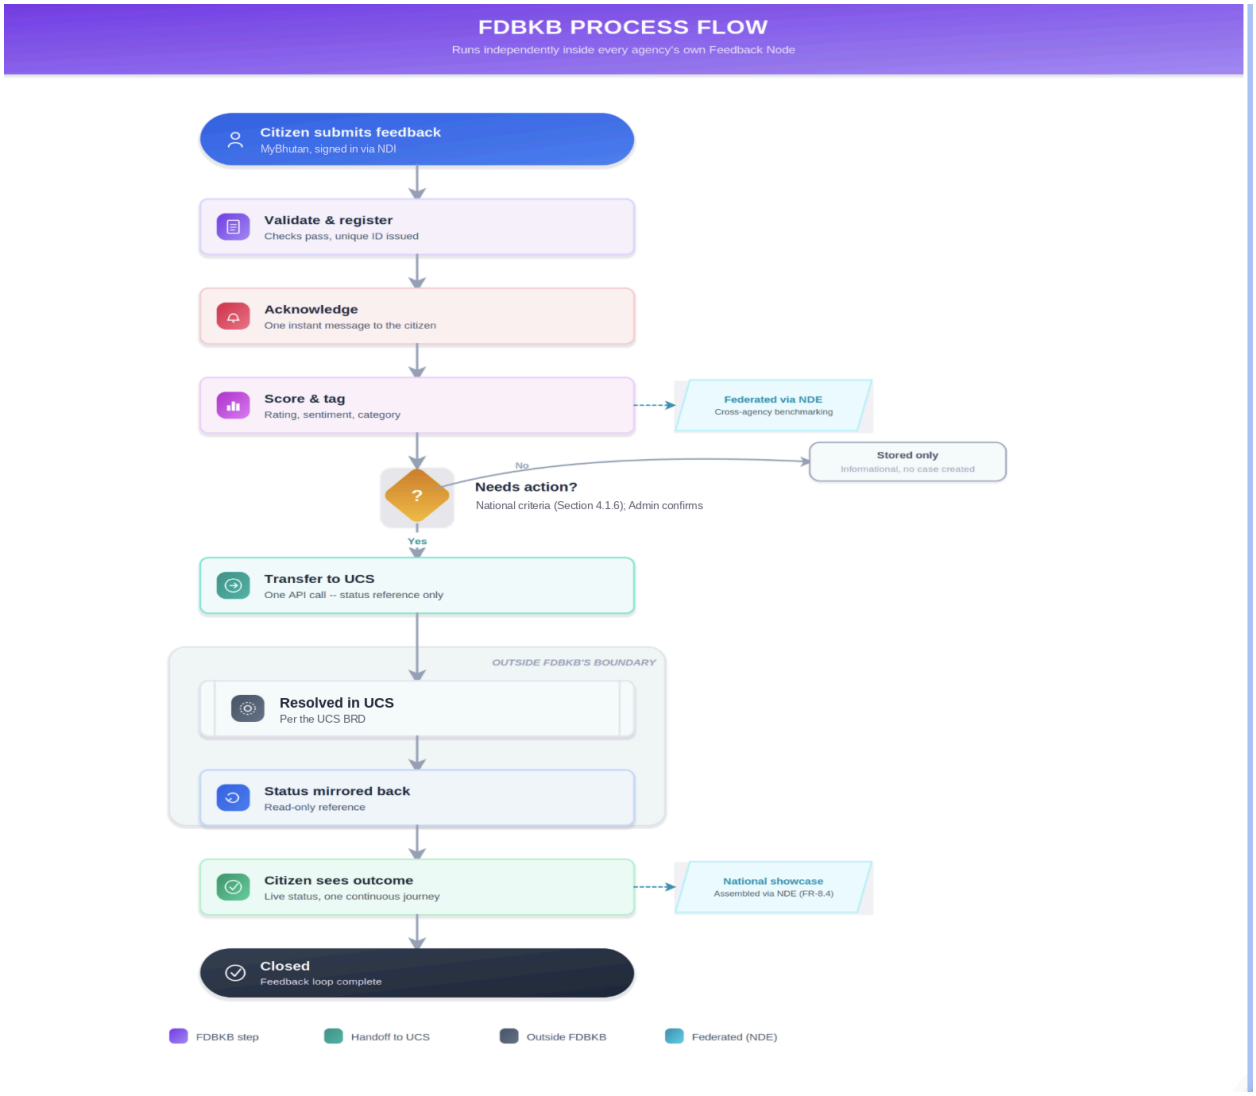


Figure 2: End-to-End Workflow of Feedback Block

4.2.1 FR-1: Feedback Form Management Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-1.1	Standard feedback form	MUST	One fixed, shared national form structure (multi-dimensional ratings, NPS, CSAT, CES, free-text comment) used identically at every Node for every service; only service name and category vary. No per-agency form design required.

FR-1.2	Category taxonomy	MUST	Shared hierarchical categories with agency sub-categories.

4.2.2 FR-2: Feedback Collection Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-2.1	MyBhutan service listing and submission	MUST	MyBhutan lists government services; submitting feedback requires NDI/SSO sign-in. Selecting a service or GA system generates that service's feedback form, responsive across desktop and mobile, with context-aware pre-population. Submission routes to the Admin of the agency operating that service.
FR-2.2	API-based submission	MUST	POST /api/v1/feedback/submit; OAuth 2.0; rate limiting; OpenAPI 3.0 documentation.
FR-2.3	Attachment support	SHOULD	Images (JPG/PNG) and PDFs, ≤2MB each, ≤5 files per submission; virus scan before acceptance; thumbnails; encrypted storage at the owning Node; download restricted to authorized Agency Node Administrators of that Node.
FR-2.4	General feedback	SHOULD	Not tied to a specific transaction, available once signed in on MyBhutan.
FR-2.5	Automated acknowledgment	MUST	Acknowledgment within seconds of submission, including a unique feedback ID and next steps (resolution timelines, where applicable, are communicated by UCS after handoff); logged for audit. This is FDBKB's only automated citizen-facing message.

4.2.3 FR-3: Rating and Scoring Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-3.1	Multi-dimensional ratings	MUST	Standardized national dimensions (Section 4.1.3); weighted aggregate score using nationally set weights.
FR-3.2	Net Promoter Score	MUST	0–10 question; Promoter/Passive/Detractor classification; trending and cross-agency benchmarking.
FR-3.3	Customer Satisfaction (CSAT)	MUST	1–5 satisfaction question; percentage satisfied vs. total; trending
FR-3.4	Sentiment analysis	MUST	NLP classification with intensity scoring, Dzongkha and English.
FR-3.5	Rating aggregation	MUST	Averages by service/agency/time period; outlier exclusion; minimum-sample threshold before display; national average assembled from agency averages.
FR-3.6	Customer Effort Score (CES)	MUST	1–7 ease-of-completion question; trending; correlation with completion and abandonment rates.

4.2.4 FR-4: Feedback Processing and Analysis Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-4.1	Automated categorization	MUST	ML-based multi-label categorization with confidence score and manual override.
FR-4.2	Trend Analysis	SHOULD	Time-series volume/sentiment trends; month-over-month and year-over-year comparison; alerting on significant change.

FR-4.3	Pattern recognition	SHOULD	Recurring-issue detection and clustering, with a signal raised when the same pattern appears across multiple agencies.
FR-4.4	Duplicate detection	MUST	24-hour auto-reject window, then flag as repeat.
FR-4.5	Comparative analysis	MUST	Cross-service/agency/geographic comparisons.
FR-4.6	Time-series storage	MUST	Feedback indexed by service, category, timestamp, sentiment, and segment; trend queries return within the performance SLA.
FR-4.7	Emotion detection	MUST	Adds a "Highly Critical" class to sentiment classification and alerts on negative-sentiment spikes.
FR-4.8	Feedback validation	MUST	Mandatory-field checks, authenticity checks, and spam flags, with human verification for flagged items.

4.2.5 FR-5: Dashboard and Reporting Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-5.1	Real-time dashboard	MUST	Live volume counter (≤30s refresh), sentiment distribution, drill-down, mobile-responsive.
FR-5.2	Service quality scorecards	MUST	NPS/CSAT/CES/average ratings with 30/90-day trends and traffic-light thresholds.
FR-5.3	Volume and Trend reports	MUST	Daily/weekly/monthly time series; split by submission route (MyBhutan form, API, batch upload); PDF/Excel export.

FR-5.4	Demographic analysis reports	MUST	Breakdown by age, gender, dzongkhag/gewog, urban/rural; minimum sample size (k-anonymity threshold set by the Federation Administrator) enforced before display.
FR-5.5	Comparative benchmarking	SHOULD	Agency-to-agency comparison matrices and rankings; exportable for presentations.
FR-5.6	Executive summary dashboards	MUST	Government-wide KPI summary (satisfaction, volume, handoff rate, linked UCS resolution status, top issues); scheduled distribution.
FR-5.7	Geographic feedback visualization	MUST	Interactive maps and heat maps, national → dzongkhag → gewog drill-down.

4.2.6 FR-6: Handoff to UCS Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-6.1	Handoff to UCS	MUST	Where an item meets the Section 4.1.6 criteria and the Agency Node Administrator confirms, FDBKB transfers it to UCS via a single API call and stores the returned reference, used only to display status. Decisions are logged to JLB.

4.2.7 FR-7: Integration and Interoperability Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-7.1	Service Registry integration (PSR)	MUST	Bidirectional sync with the Public Service Registry; feedback linked to services via service_id; quality scores published back.

FR-7.2	Identity integration (NDI)	MUST	NDI OIDC/OAuth 2.0 sign-in for every MyBhutan submission; API and batch submissions must carry an NDI-verified CID and recorded consent; no unauthenticated or anonymous submission path.
FR-7.3	Notification block integration (NTFB)	MUST	Acknowledgment notifications only; subsequent notifications are UCS's responsibility.
FR-7.4	Journaling block integration (JLB)	MUST	Immutable logging of submission, status change, and access events.
FR-7.5	National Data Exchange integration (NDE)	MUST	Every agency's deployment registers for discovery and supports federated queries.
FR-7.6	Node conformance certification	MUST	A new agency's Feedback Node completes automated conformance testing against the shared codebase, schema, taxonomy, and API contract (via the federation sandbox, NFR-7.7) before it is registered in the NDE Node Registry or permitted to serve federated queries; certification status is visible to the Federation Administrator and re-verified at each quarterly federation-wide conformance re-check (Section 6.3.2).

4.2.8 FR-8: Citizen Engagement and Transparency Requirements

ID	Requirement	Priority	Acceptance Criteria

FR-8.1	Personal feedback history	MUST	Chronological, filterable, exportable history across every agency the citizen has used; linked case status shown inline.
FR-8.2	Acknowledgment notification only	MUST	Delivery of the initial acknowledgment via email/SMS and the MyBhutan inbox.
FR-8.3	Public feedback statistics	MUST	aggregated statistics — service scores, top issues, volume trends; published on a regular schedule.
FR-8.4	Public improvement showcase	SHOULD	Each Node publishes anonymized before/after improvement stories sourced from UCS action/resolution records linked back to the originating feedback_id (Section 4.1.8); a national showcase aggregates these across Nodes on a published schedule, with no individual citizen or case identifiable.

4.2.9 FR-9: Data Privacy and Security Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-9.1	Data minimization	MUST	Only necessary fields collected; optional fields clearly marked; location captured at dzongkhag level, not precise coordinates.
FR-9.2	User data rights	MUST	Access, correction, erasure, and portability, consistent with the Information, Communications and Media Act of Bhutan 2018, the National Digital Identity Act of Bhutan 2023, and the National Data Governance Framework 2025.
FR-9.3	Access control	MUST	Role-based access with least privilege; staff see only their own agency's data; quarterly access review.

FR-9.4	Encryption in transit	MUST	TLS 1.3 with strong cipher suites.
FR-9.5	Encryption at rest	MUST	AES-256 database encryption; field-level encryption for sensitive content; keys held by the owning agency.
FR-9.6	Fraud detection	SHOULD	Duplicate and spam/bot detection, with a manual review queue.
FR-9.7	Security assessments	MUST	Quarterly vulnerability scanning and annual penetration testing.
FR-9.8	Erasure Handling	SHOULD	On erasure, individual data becomes unrecoverable while aggregated contributions may remain; erasure logged for audit.
FR-9.9	Federated query access control and logging	MUST	Every federated query is authenticated and authorized using a short-lived, scoped token issued for that specific query purpose (NFR-4.6); each federated query, together with the requesting role and purpose, is logged to JLB (Section 6.2.1, Layer 5) so cross-agency access is auditable and no query can silently reach another Node's data.

4.2.10 FR-10: Administration and System Management Requirements

ID	Requirement	Priority	Acceptance Criteria
FR-10.1	User and role management	MUST	Create/update/deactivate accounts; role assignment; bulk import. Authentication delegated to NDI.
FR-10.2	System configuration	MUST	Acknowledgment templates, handoff criteria, and integration endpoints, with an approval workflow for changes. SLA timelines and escalation rules are configured in UCS.

FR-10.3	Content moderation	MUST	Profanity filtering, flagging queue, and an appeal path for rejected submissions.
FR-10.4	Data retention	MUST	Configurable retention by feedback type; automated archival and secure deletion.
FR-10.5	API management	MUST	Expiring keys; tiered rate limits.

4.3 Non-Functional Requirements

4.3.1 Performance

ID	Requirement	Priority	Target Metric
NFR-1.1	API response time (per Node)	MUST	Submission API: ≤500ms (95th pct), ≤1000ms (99th pct). Retrieval: ≤300ms (95th pct). Search: ≤1000ms (95th pct). Analytics: ≤2000ms (95th pct), under normal load, at each Node.
NFR-1.2	Page load time	MUST	FCP ≤1.5s, LCP ≤2.5s, TTI ≤3.5s; dashboard initial load ≤2s; measured on a 4G connection with a mid-range device, since this is representative of most citizen access outside the capital.
NFR-1.3	Database query performance (per Node)	MUST	Simple indexed queries ≤50ms (95th pct); complex joins/aggregations ≤500ms (95th pct); full-text search ≤200ms (95th pct); report generation ≤5s, at each Node's own database.
NFR-1.4	Concurrent user support (per Node, sized to agency)	MUST	Each Node sized to its own agency's expected load, with a baseline of ≥200 concurrent active users and burst capacity to 1,000 during high-visibility campaigns; ≥5 concurrent feedback submissions per second sustained per Node, bursting to 20/sec at peak;

			larger agencies (e.g., Ministry of Health, Ministry of Education) provisioned to the higher end of this range based on service volume.
NFR-1.5	Federated query performance	MUST	A federated query spanning up to 50 Nodes (e.g., a national executive dashboard) returns within 15 seconds for a cached/scheduled view, or within 60 seconds for an on-demand live federated query; individual Node timeouts do not block the overall response — a degraded Node is reported as "data unavailable" rather than failing the whole query.

4.3.2 Availability

ID	Requirement	Priority	Target Metric
NFR-2.1	System uptime (per Node)	MUST	99.5% in Year 1 rising to 99.9% at steady state (≤8h46m/year downtime) per Node, excluding planned maintenance; planned windows ≤44 min/month during 10 PM–2 AM. An outage at one agency's Node does not count against any other Node's availability target.
NFR-2.2	Service redundancy (per Node)	MUST	Each Node deployed active-active across 2+ availability zones where the agency's scale justifies it, or with an agreed shared-infrastructure fallback for smaller agencies; database replication with automatic failover (RPO ≤1h, RTO ≤4h); load-balancer health checks every 30s; stateless application tier.
NFR-2.3	Disaster recovery capability (per Node)	MUST	RTO 4h, RPO 1h, per Node; geographically separate backup for that Node's data; automated failover for critical services; DR drill twice yearly with

			documented findings, coordinated with the Federation Administrator's national DR calendar.
NFR-2.4	Graceful degradation	SHOULD	Core submission at a Node stays functional if that Node's own analytics fails; read-only mode on write failure; MyBhutan caches an in-progress submission client-side and retries on reconnect rather than offering a full offline mode, consistent with the single citizen channel in Section 2.1.1; circuit breakers on PSR/NDI/NTFB/NDE calls, so that a federation-layer or another agency's outage does not block local submission.
NFR-2.5	Monitoring and alerting	MUST	24/7 monitoring per Node; immediate alert on downtime; 5-minute alert at >5% API error rate; 15-minute alert at >20% response-time degradation; P1 response within 30 minutes, by that Node's own operations team, with a federation-level status view showing which Nodes are healthy.

4.3.3 Scalability

ID	Requirement	Priority	Target Metric
NFR-3.1	Data volume capacity (per Node, aggregate across the federation)	MUST	Aggregate across all Nodes: 500,000 feedback records in Year 1 (consistent with the ≥50,000/year adoption target in Section 1.3, with 10x headroom); 2 million cumulative records by Year 3, distributed across ~50 agency Nodes rather than one database. Each Node is individually provisioned to its own agency's share of this volume (e.g., 10GB–50GB initial database storage per Node, scaling to 200GB for the largest agencies); attachment storage similarly

			sized per Node. Query performance maintained as each Node's own data grows.
NFR-3.2	Horizontal scalability (per Node)	MUST	Stateless application tier at each Node; add/remove servers without downtime; auto-scale at 70% CPU up / 30% down; database read replicas; clustered cache layer (e.g., Redis), sized to that Node's own load.
NFR-3.3	Peak load handling (per Node)	MUST	Each Node handles 3x its own normal load during peak periods (e.g., a national campaign or post-budget-speech spike) without degradation; auto-scaling triggers within 2 minutes; queue-based async processing; quarterly load testing per Node to verify capacity against that Node's actual usage.
NFR-3.4	Geographic distribution	SHOULD	Static assets (no PII) served via CDN for every Node; since Bhutan has no domestic CDN point of presence, the nearest practical edge is in India (e.g., Kolkata/Delhi/Mumbai) — acceptable for static assets, but feedback content and PII are excluded from CDN caching and stay on each agency's own in-country infrastructure per the data-sovereignty principle in NFR-6.1.

4.3.4 Security

ID	Requirement	Priority	Target Metric
NFR-4.1	Authentication security	MUST	Authentication delegated to NDI for all users (FR-10.1); 30-minute admin / 2-hour citizen session timeout; MFA mandatory for admin access. Local break-glass admin accounts only: 12+ character

			complex passwords, 5-attempt lockout for 30 minutes, 90-day rotation.
NFR-4.2	API security (per Node)	MUST	OAuth 2.0 bearer tokens; tiered rate limiting (100/1,000/10,000 req/min) per Node; parameterized queries; output encoding; CSRF tokens on state-changing calls.
NFR-4.3	Data encryption standards (per Node)	MUST	TLS 1.3 only (no TLS 1.2 or below) externally; AES-256 at rest; field-level encryption for highly sensitive feedback; KMS-managed keys with 90-day rotation, with each agency controlling its own encryption keys rather than a shared central key authority.
NFR-4.4	Vulnerability management (per Node)	MUST	Quarterly scans; annual third-party penetration test; critical patches within 72h, high within 7 days; CI/CD dependency scanning; OWASP Top 10 verification — each Node's operating agency is accountable for its own patch cadence, against a shared national minimum standard.
NFR-4.5	Audit trail completeness	MUST	100% of security-relevant events (auth, authorization, data access, config change, user management) at every Node logged to JLB with timestamp/user/action/resource/IP/result; immutable; 7-year retention.
NFR-4.6	Federated access security	MUST	Any cross-Node query is authenticated and authorized separately from Node-local access, using short-lived, scoped tokens issued for a specific federated query purpose; a Node may decline or

			rate-limit federated queries it deems abnormal without breaching the federation SLA.

4.3.5 Usability

ID	Requirement	Priority	Target Metric
NFR-5.1	Ease of use	MUST	Median submission completed within 2 minutes; ≤3 clicks to any major function; inline help; progressive disclosure; clear, actionable error messages — consistent across every Node since they share the same codebase.
NFR-5.2	Responsive design	MUST	Mobile-first, 320px–3840px; ≥44x44px touch targets; no horizontal scroll; tested on iOS Safari, Android Chrome, desktop Chrome/Firefox/Edge.
NFR-5.3	User error prevention	SHOULD	Client-side validation; confirmation dialogs for destructive actions; 30-second autosave; unsaved-changes warning; clear required-vs-optional labelling.
NFR-5.4	Learnability	SHOULD	First-time citizen users can submit without training, regardless of which agency's Node they land on, since the interface is standardized within MyBhutan; contextual help; admin onboarding wizard; consistent navigation patterns.

4.3.6 Compliance

ID	Requirement	Priority	Target Metric

NFR-6.1	Data protection compliance (per Node as controller)	MUST	Compliance with the Information, Communications and Media Act of Bhutan 2018 and the National Digital Identity Act of Bhutan 2023, and alignment with the National Data Governance Framework 2025; privacy by design; data minimization; consent management; user rights (access/rectification/erasure/portability), exercised against whichever Node(s) hold the data; breach notification within 72 hours by the affected Node's operating agency; primary feedback data and PII stored in-country at each agency's own infrastructure (see NFR-3.4 on CDN scope).
NFR-6.2	GovStack standards compliance	MUST	GovStack Building Block alignment; CloudEvents for event-driven integration; OpenAPI 3.0 for all APIs; RFC 7807 error responses; documented versioning strategy — enforced identically at every Node through the shared codebase and NDE conformance testing.
NFR-6.3	Security standards compliance (per Node)	MUST	ISO/IEC 27001:2022 controls; OWASP Top 10 verification; Bhutan e-Government security standards adherence; quarterly security audits, each Node audited independently. FDBKB processes no payment data, so PCI DSS does not apply.
NFR-6.4	Audit and reporting	MUST	Annual compliance report to the GovTech Agency, assembled from every Node's own compliance status; quarterly security assessment per Node; monthly operational report per Node; a federation-level compliance dashboard; export support for external audits.

4.3.7 Operational

ID	Requirement	Priority	Target Metric
NFR-7.1	Deployment automation	MUST	CI/CD with automated unit/integration/E2E tests against the shared codebase; ≥80% code coverage; zero-downtime blue-green or rolling deploys, rolled out to Nodes on a coordinated release train; 5-minute rollback per Node; dev/staging/production parity.
NFR-7.2	Monitoring and observability (per Node, federated view)	MUST	APM with distributed tracing at each Node; infrastructure and centralized-log monitoring per Node; metrics dashboards; 24/7 on-call rotation per Node's operating agency, with a federation-level status view.
NFR-7.3	Backup and recovery (per Node)	MUST	Daily automated + 6-hourly incremental + weekly full backups at each Node; 7-daily/4-weekly/12-monthly retention; geo-redundant storage in 2+ locations within Bhutan or an approved in-country arrangement; monthly restore test; RTO 4h, RPO 1h.
NFR-7.4	Documentation	MUST	Auto-published OpenAPI/Swagger docs common to every Node; admin and developer guides; operations runbook per Node; architecture diagrams; incident-response playbooks, updated each release of the shared codebase.
NFR-7.5	Maintainability	SHOULD	Code maintainability index >70; technical-debt ratio <5%; documentation coverage >60%; linted coding standards; review on every pull request to the shared national codebase.

NFR-7.6	Support and training	MUST	Business-hours help desk (9 AM–5 PM Mon–Fri, Bhutan Standard Time) per Node's operating agency, plus a federation-level help desk for cross-agency/NDE issues; 24/7 emergency P1 support; two-day admin training workshop for each new Node's administrators; Dzongkha and English user documentation; SLA-tiered ticketing (P1 1h / P2 4h / P3 24h).
NFR-7.7	Environment management	MUST	Separate dev/test/staging/production per Node; Infrastructure as Code; secrets management held by each agency; environment-appropriate sandbox for partner testing, plus a shared federation sandbox for NDE conformance testing of new Nodes.

5. Solution & Integration Architecture

5.1 System Architecture Overview

5.1.1 Architecture Principles

- **Federation-First Design:** every capability that must work across agencies (national dashboards, cross-agency benchmarking, oversight access) is built on top of federated query and event patterns via the NDE, never on a shared database.
- **API-First Design:** all functionality via OpenAPI 3.0 REST APIs, versioned with backward compatibility, and identical across every Node so that any authorized client can integrate with any agency's Node the same way.
- **Event-Driven Architecture:** CloudEvents over Kafka for real-time flow from each Node to JLB and NTFB ; a separate federation-wide event topic carries cross-Node signals (duplicate/pattern alerts, Node health) without carrying feedback content itself.
- **Microservices per Node:** Feedback Submission, Rating Engine, Analysis Engine, Notification Manager (acknowledgment-only), and Admin Console run as an independent set of microservices within each Node, communicating via REST and message queues local to that Node.

- **Stateless Application Layer:** session state in a distributed cache (e.g., Redis) local to each Node for horizontal scalability.
- **Cloud-Native:** containerized (Docker), orchestrated (Kubernetes), Infrastructure as Code (Terraform) — the same deployment templates are reused by every Node, customized only for that agency's scale and any agency-specific branding.
- **Security by Design:** zero-trust with mutual TLS for service-to-service calls within a Node, and for Node-to-federation calls; defense in depth; secrets management (e.g., HashiCorp Vault) held independently by each agency.
- **No Central Point of Control:** the Federation Administrator (GovTech) sets and certifies the shared standard but does not operate, and by default cannot directly query any individual agency's Node data.

5.1.2 Architecture Layers

Each of the following layers exists within every Node, built from the same shared templates; the Federation Layer described afterward exists once, nationally, and connects the Nodes together.

Presentation Layer:

- Web Interface: the feedback form and personal history are rendered inside MyBhutan (React SPA components), responsive for desktop, tablet, and mobile — no separate kiosk mode or offline-capable PWA build.
- Admin Console: web-based administration interface for Node's government staff. Dashboard, reporting, configuration management. Role-based access control, local to the Node.

API Gateway Layer:

- API Gateway: a single entry point for each Node's external requests (e.g., Kong or AWS API Gateway). Rate limiting, request validation, API key management. SSL termination and request routing. A lightweight routing/discovery layer, informed by the NDE catalog, directs incoming MyBhutan requests to the correct agency Node's gateway.

- Authentication & Authorization: OAuth 2.0 / OIDC integration with NDI/SSO, the same shared identity provider used by every Node. JWT token validation. Role-based access control enforcement, local to the Node.
- Load Balancer: application load balancer for that Node's own traffic distribution. Health checks and automatic failover. SSL/TLS termination.

Application Services Layer:

- Feedback Submission Service: handles feedback submission via the MyBhutan form, API, and batch upload for that Node's agency. Validation, deduplication, routing. Integration with NDI/SSO for authentication.
- Rating & Scoring Service: calculates NPS, CSAT, CES scores for Node's own services. Multi-dimensional rating aggregation. Sentiment analysis using shared NLP models.
- Analysis Engine: automated categorization using machine learning; pattern recognition; trend analysis, all scoped to that Node's own data; also computes the handoff recommendation against the Section 4.1.6 criteria for the Agency Node Administrator to confirm (FR-6.1).
- Handoff Adapter: a lightweight adapter that calls UCS once per handoff and stores the reference returned.
- Notification Manager: multi-channel notification delivery (email, SMS, push) for Node's own citizens, scoped to the single acknowledgment message only (FR-2.5/FR-8.2). Integration with NTFB. Template rendering and personalization.
- Reporting & Analytics Service: dashboard data aggregation for that Node; report generation (PDF, Excel, CSV); real-time metrics computation, scoped to the Node's own data.
- Integration Hub: manages Node's own integrations with PSR, JLB, and — critically — the NDE, through which it participates in the federation.

Data Layer:

- Primary Database: PostgreSQL for Node's own transactional data (feedback, users, configurations). Supports ACID transactions, JSONB for flexible schema. Read replicas for read scalability, all local to the Node.

- Document Store: MongoDB for Node's own unstructured feedback text and attachments metadata. Flexible schema for varying feedback formats. Horizontal sharding for scalability, local to the Node.
- Cache Layer: Redis for Node's own session management and API response caching. Distributed cache with clustering. Pub/sub for real-time updates within the Node.
- Object Storage: S3-compatible object storage for Node's own file attachments. Encrypted storage with versioning. Lifecycle policies for archival, local to the Node.
- Search Index: Elasticsearch for full-text search of that Node's own content. Indexed feedback content for fast retrieval. Faceted search and aggregations, local to the Node; federated search spans Nodes only through the Federation Layer described below.

Integration Layer:

- Message Broker: Apache Kafka, one instance per Node for that Node's own event streaming, plus the federation-wide topic described above. Guaranteed delivery and ordering.
- NDE Adapter: National Data Exchange (NDE, WSO2-based) integration for secure cross-agency data exchange, present in every Node as the door into the federation. Service discovery and registration via the NDE catalog.
- External API Clients: REST clients for PSR, NDI, NTFB, JLB, and UCS — implemented once in the shared codebase, configured per Node with that agency's own credentials. Retry logic and circuit breakers. Connection pooling and timeout management.
- **Federation Layer (national, operated by GovTech):** NDE Node Registry (service_id → owning Node, certification status); Federated Query Broker (issues scoped, short-lived tokens, fans queries out to Nodes, returns privacy-filtered aggregates and flags unavailable Nodes; FR-9.9, NFR-1.5); Cross-Node Signal Bus (duplicate, pattern, and Node-health signals, with no feedback content); and a federation status view. It holds no feedback data of its own.

5.2 Technology Stack

Layer	Technology (Example)	Purpose

Frontend	React 18+, TypeScript, Tailwind CSS	Component-based UI development, embedded in MyBhutan. Type safety. Utility-first styling. Identical across every Node.
Backend Framework	Node.js (Express) or Spring Boot (Java)	High-performance asynchronous I/O. Enterprise-grade framework, deployed per Node.
API Gateway	Kong Gateway or AWS API Gateway	API management, rate limiting, authentication, routing, one instance per Node.
Primary Database	PostgreSQL 15+	ACID transactions, JSONB support, advanced indexing, read replicas — one database per Node, never shared across agencies.
Document Database	MongoDB 6+	Flexible schema for unstructured feedback, per Node. Horizontal sharding within that Node.
Cache	Redis 7+	In-memory data store per Node. Session management. Pub/sub messaging.
Search Engine	Elasticsearch 8+	Full-text search per Node. Faceted search. Real-time indexing, local to that Node's data.
Object Storage	MinIO or AWS S3	File attachment storage per Node. Encrypted. S3-compatible API.

Message Broker	Apache Kafka 3+	Event streaming per Node, plus the shared federation-wide topic. Guaranteed delivery. High throughput.
Authentication	OAuth 2.0 / OIDC (Keycloak or Auth0), fronting NDI	NDI integration, common across the federation. SSO. Token-based authentication.
Federation / Integration Fabric	National Data Exchange (NDE), WSO2-based	Node discovery, federated query brokering, cross-agency secure exchange.
Containerization	Docker 24+	Consistent deployment environments, identical container images used by every Node.
Orchestration	Kubernetes 1.28+	Container orchestration per Node. Auto-scaling. Self-healing. Rolling updates.
CI/CD	GitLab CI/CD or GitHub Actions	Automated testing, building, deployment of the shared codebase, released to Nodes on a coordinated train.
Infrastructure as Code	Terraform 1.5+	Cloud infrastructure provisioning, one reusable module set applied per Node.
Monitoring	Prometheus + Grafana	Metrics collection, alerting, visualization, per Node, with a federation-level status rollup.
Logging	ELK Stack (Elasticsearch, Logstash, Kibana)	Centralized logging within each Node. Log aggregation and search.

APM	New Relic or Datadog	Application performance monitoring per Node. Distributed tracing, including federation-layer calls.
NLP/ML	Python (scikit-learn, spaCy, Hugging Face)	Sentiment analysis. Text categorization. Pattern recognition — a shared national model distributed to and run locally within every Node.
Secrets Management	HashiCorp Vault or AWS Secrets Manager	Secure storage of credentials, API keys, certificates — a separate instance per agency, so no agency's secrets are held by another.

5.3 Integration Architecture

5.3.1 DPI Integration Patterns

Synchronous RESTful Integration:

- Public Service Registry (PSR): each Node calls GET /services for service metadata retrieval and POST /services/{id}/quality-scores to publish its own satisfaction ratings; webhook callbacks for service updates.
- National Digital Identity (NDI): OAuth 2.0 / OIDC authorization flow, common to every Node. GET /userinfo for profile retrieval. Token validation endpoint for session management.
- Notification Block (NTFB): each Node calls POST /notifications/send with template_id, recipient, channel, for the acknowledgment message only. Synchronous acknowledgment with async delivery. Webhook for delivery status callbacks.
- Unified Client Support (UCS): on a confirmed handoff, each Node calls POST /tickets/from-fdbkb with the feedback content, ratings, sentiment, category, and NDI identity/consent, and stores the returned ticket reference; UCS sends status updates back via webhook (CloudEvents).

Asynchronous Event-Driven Integration (per Node):

- Journaling Block (JLB): each Node publishes its own audit events to a Kafka topic (e.g., fdbkb.{agency}.audit.events). CloudEvents format with metadata (user, action, timestamp, resource). Fire-and-forget pattern with Kafka guaranteed delivery. UCS separately publishes its own case-lifecycle audit events to JLB, cross-referenced by ticket ID.

Widget Embedding for MyBhutan:

- JavaScript SDK: lightweight SDK (<50KB) for easy integration, common across every Node's forms. Configuration: service_id, language, theme (the service_id is resolved by MyBhutan, via the NDE catalog, to the correct owning Node before the widget loads). API: FDBKB.init(), FDBKB.show(), FDBKB.hide().
- Iframe Embedding: sandboxed iframe for security isolation, pointed at the correct Node's endpoint. PostMessage API for cross-origin communication. Context passing: service_id, user_context, session_token.
- Server-Side Rendering: Server-side widget rendering for faster first load on low-bandwidth connections, rendered by the owning Node. Progressive enhancement for JavaScript-disabled browsers.

5.3.2 API Architecture

Endpoint	Method	Purpose
/api/v1/feedback/submit	POST	Submit new feedback with ratings and text, to the Node that owns the named service
/api/v1/feedback/{id}	GET	Retrieve feedback by ID with full details, from the Node that issued the ID
/api/v1/feedback/{id}/status	PUT	Update feedback status (admin only), at the owning Node

/api/v1/feedback/{id}/handoff	POST	Transfer the item to UCS once the Agency Node Administrator confirms the handoff (FR-6.1), and store the returned reference
/api/v1/feedback/user/{userId}	GET	Get personal feedback history for an authenticated user **at this Node**; the federation layer calls this endpoint across every Node the user has used to assemble a full history
/api/v1/forms	GET	Retrieve the current national standard form (read-only at each Node)
/api/v1/forms/{id}	POST	Publish/update the national standard form (Federation Administrator only)
/api/v1/reports/generate	POST	Generate a custom report with filters, scoped to this Node's own data
/api/v1/analytics/dashboard	GET	Retrieve real-time dashboard data (volume, sentiment, NPS/CSAT/CES) for this Node, filterable by date range and service, scoped to this Node's own data (FR-5.1)
/api/v1/federation/nodes	GET	List certified Feedback Nodes and the services each owns, from the NDE Node Registry (federation layer only; FR-7.5, FR-7.6)
/api/v1/federation/query	POST	Submit an authorized cross-Node query, using a scoped, short-lived token and logged to JLB, and receive aggregated,

		privacy-filtered results, with any unavailable Nodes reported as partial (federation layer only; FR-9.9, NFR-1.5)

API Standards: RESTful principles with resource-oriented URLs; JSON request/response format with UTF-8 encoding; standard HTTP status codes; standardized error responses following RFC 7807 Problem Details; API versioning in URL path (/api/v1/, /api/v2/); pagination for list endpoints; HATEOAS links; OpenAPI 3.0 specification documentation — all common across every Node so that switching between agencies' Nodes requires no client-side rework.

5.4 Deployment Architecture

Per-Node Cloud Infrastructure:

- Each agency's Feedback Node is deployed independently, sized to that agency (see NFR-1.4 and NFR-3.1); larger agencies deploy across 2+ availability zones, while smaller agencies may share a common, GovTech-provided hosting arrangement while still maintaining logically and cryptographically separate data.
- Kubernetes cluster per Node (or a shared multi-tenant cluster with strict namespace isolation for smaller agencies), with auto-scaling based on Node's own CPU/memory.
- Load balancer with SSL termination and health checks, per Node.
- CDN for static assets only (no PII), shared where convenient across Nodes.

Database Deployment (per Node):

- PostgreSQL: primary-replica configuration with synchronous replication, one deployment per Node. Automated failover using Patroni or equivalent. Connection pooling via PgBouncer.
- MongoDB: replica set with 3+ members, per Node. Sharding for horizontal scalability (future) within that Node. Automated backup to that agency's own S3-compatible storage.

- Redis: cluster mode with master-replica replication, per Node. Persistence via RDB snapshots and AOF logs.

Disaster Recovery (per Node):

- Each Node's backup lives in a geographically separate location from its primary, within Bhutan or an approved in-country arrangement.
- Database replication to that Node's own DR site with <1 hour lag.
- Automated failover procedures documented and tested per Node.
- RTO: 4 hours, RPO: 1 hour, per Node.
- DR drills twice yearly, coordinated on a shared national calendar by the Federation Administrator so drills don't collide with shared-service (NDI/NDE/JLB) maintenance windows.

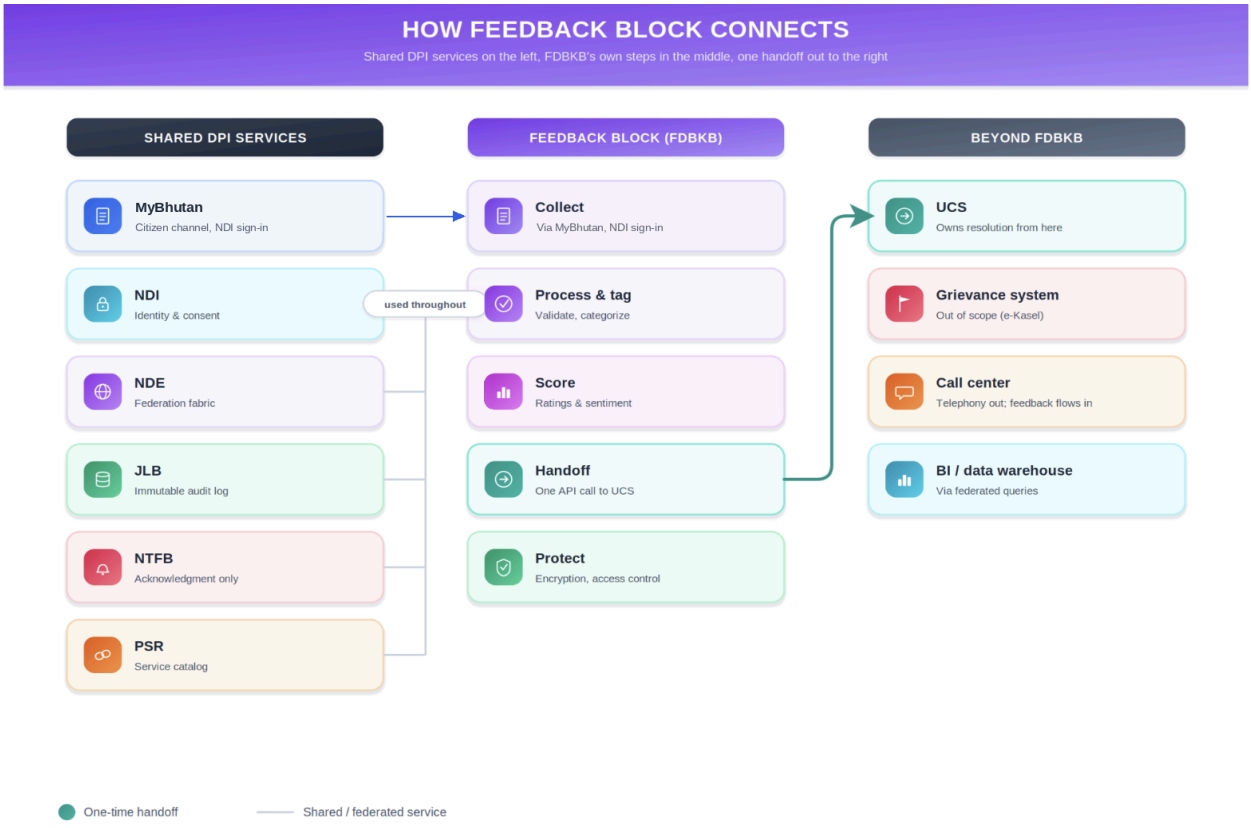


Figure 3: Indicative Integration of Feedback Block

6. Data, Security & Compliance

6.1 Data Management

6.1.1 Data Model (Indicative)

Every Node maintains its own instance of the following data model for the feedback it owns; there is no single national table for any of these entities. The NDE catalog additionally maintains a lightweight federation-level index (service_id → owning Node) so that clients and the federated query broker know where to send a request.

Entity	Key Attributes	Storage (Example)
Feedback	feedback_id (globally unique, Node-prefixed), service_id, user_id (always populated, from NDI), submission_timestamp, submission_route (MyBhutan/API/batch), ratings, text_content, category, sentiment_score, fdbkb_status, linked_case_reference	PostgreSQL (structured data), MongoDB (text content), Elasticsearch (search index) — all local to the owning Node
User	user_id (CID), name, email, phone, preferred_language, preferred_contact_method, age_group, gender, dzongkhag, consent_timestamp	PostgreSQL, synchronized from NDI, at each Node the citizen has interacted with
Service	service_id, service_name (Dzongkha/English), owning_agency, category, description, avg_satisfaction, nps_score, total_feedback_count	PostgreSQL, synchronized from PSR, at the owning agency's Node; also indexed (service_id → owning Node) in the NDE federation catalog

Form	form_id, form_name, service_id, questions (JSONB), triggers, version, status (draft/published), created_by, created_at	PostgreSQL (JSONB); national form published by the Federation Administrator, replicated read-only to every Node
Attachment	attachment_id, feedback_id, filename, file_type, file_size, storage_path, uploaded_timestamp, virus_scan_status	PostgreSQL (metadata), S3 (file content) — both at the owning Node
AuditLog	event_id, event_type, user_id, resource_id, action, timestamp, ip_address, user_agent, result	JLB (via Kafka stream, common shared service), plus recent logs cached in PostgreSQL at each Node. UCS logs its own case-lifecycle events separately, cross-referenced by case reference.

6.1.2 Data Lifecycle Management

Collection Phase (at the originating Node):

- Data minimization: collect only essential information for feedback processing.
- Consent capture: explicit consent for data usage at submission time on MyBhutan.
- Validation: client-side and server-side validation for data quality.
- Encryption: TLS 1.3 for data in transit during submission to the Node.

Active Phase (0–2 years, at the owning Node):

- Hot storage in PostgreSQL/MongoDB local to that Node for fast access.
- Real-time indexing in Elasticsearch local to that Node for search.

- Backups per NFR-7.3 (daily, 6-hourly incremental, weekly full; 7-daily/4-weekly/12-monthly retention), at that Node.
- Available for reporting and analytics within the Node, and for display of UCS-mirrored status; available for federated queries only through the authorized federation broker.

Archival Phase (2–7 years, at the owning Node):

- Automated migration to that Node's own cold storage (e.g., S3 Glacier or equivalent).
- Compressed and encrypted archives, keys held by the owning agency.
- Metadata retained in Node's own RDBMS for discovery.
- Retrieval possible within 24 hours for audit/legal requirements, coordinated by the owning agency.

Deletion Phase (after 7 years, at the owning Node):

- Automated deletion after retention period expires, independently at each Node.
- Secure deletion: cryptographic erasure (delete encryption keys), performed by the owning agency.
- Audit trail of deletion events logged to JLB.
- Legal hold mechanism to prevent deletion if under investigation, applied at the Node holding the relevant record.

6.1.3 Data Privacy Controls

User Rights Implementation (across whichever Nodes hold the citizen's data):

- Right to Access: self-service view in MyBhutan of all personal feedback data, assembled via federated lookup across every Node the citizen has used.
- Right to Rectification: update contact information and correct factual errors, at each relevant Node.
- Right to Erasure: delete feedback with exceptions (legal/audit requirements); a citizen exercising this right may need it actioned at more than one Node, coordinated via a federation-level request tracker.

- Right to Data Portability: export personal data in machine-readable format (JSON, CSV), combined across Nodes on request.
- Right to Object: opt-out of non-mandatory processing (marketing, research), settable per Node or, where the citizen prefers, propagated to all Nodes via the federation layer.
- Right to Restrict Processing: limit processing to specific purposes only, at each relevant Node.

6.2 Security Architecture

6.2.1 Defense in Depth

Layer 1: Network Security (per Node)

- Virtual Private Cloud (VPC) with private subnets for database and application tiers, one per Node.
- Network Access Control Lists (NACLs) and Security Groups for traffic filtering, configured independently by each agency.
- Web Application Firewall (WAF) with OWASP Core Rule Set, at each Node.
- DDoS protection with rate limiting at the edge, at each Node.
- VPN for administrative access, no direct public access to infrastructure, for each Node independently.

Layer 2: Application Security

- Input validation on all user inputs (whitelisting, length limits, type checking).
- Output encoding to prevent XSS attacks.
- Parameterized queries to prevent SQL injection.
- CSRF token validation on all state-changing operations.
- Content Security Policy (CSP) headers to mitigate XSS.
- Security headers: HSTS, X-Frame-Options, X-Content-Type-Options.

Layer 3: Authentication & Authorization

- OAuth 2.0 / OIDC integration with NDI for citizen authentication — a shared, common service used by every Node.

- Multi-Factor Authentication (MFA) required for administrative users, enforced independently at each Node.
- Role-Based Access Control (RBAC) with least privilege principle, local to each Node.
- Service-level data isolation: Agency Node Administrators only see their own agency's Node's feedback — by construction, since there is no shared database to accidentally expose another agency's data.
- JWT token-based API authentication with short expiry (15 minutes).
- API key rotation every 90 days, per Node.

Layer 4: Data Security

- TLS 1.3 for all data in transit with perfect forward secrecy.
- AES-256 encryption for data at rest (database, backups, file storage) at each Node.
- Field-level encryption for highly sensitive feedback content, using that Node's own agency-held keys.
- Key management (e.g., HashiCorp Vault, AWS KMS) with key rotation, held by each agency independently — no agency's data can be decrypted using another agency's keys, and GovTech does not hold a national master key.
- Database Transparent Data Encryption (TDE), per Node.

Layer 5: Audit & Monitoring

- Comprehensive audit logging to JLB for all security events, from every Node.
- Every federated query, with the requesting role, purpose, and scoped token ID, is logged to JLB (FR-9.9).
- Real-time security monitoring with SIEM integration, per Node, with a federation-level rollup of alert status (not raw data).
- Intrusion Detection System (IDS) / Intrusion Prevention System (IPS), per Node.
- Security Information and Event Management (SIEM) for threat correlation, per Node.

6.2.2 Security Operations

Vulnerability Management (per Node, to a shared national standard):

- Quarterly vulnerability scanning using automated tools (Nessus, Qualys), by each Node's own operating agency.
- Annual penetration testing by certified ethical hackers (CEH, OSCP), of each Node.
- Patch management: critical vulnerabilities patched within 72 hours, high within 7 days, by each Node's own team.
- Dependency scanning in CI/CD pipeline (Snyk, OWASP Dependency Check), applied to the shared codebase before release to all Nodes.
- Container image scanning for vulnerabilities before deployment to any Node.

Incident Response:

- Security Incident Response Team (SIRT) with 24/7 on-call rotation, at each Node's operating agency, supported by a federation-level SIRT for cross-agency or NDE-layer incidents.
- Incident response playbooks for common scenarios (data breach, DDoS, ransomware), standardized nationally, executed locally.
- Incident severity classification: P1 (critical), P2 (high), P3 (medium), P4 (low).
- Response SLAs: P1 (30 minutes), P2 (2 hours), P3 (8 hours), P4 (24 hours).
- Post-incident review and root cause analysis within 48 hours, by the affected Node, with findings shared to the federation for lessons-learned purposes.
- Data breach notification: affected users within 72 hours by the affected Node's operating agency, regulators as required.

Security Testing:

- OWASP Top 10 compliance verification in every release of the shared codebase.
- Static Application Security Testing (SAST) in CI/CD pipeline.
- Dynamic Application Security Testing (DAST) on staging, before release to Nodes.
- Security code reviews for all high-risk changes to the shared codebase.
- Regular security training for the development team and for each Node's local administrators.

6.3 Compliance Framework

6.3.1 Regulatory Compliance

Regulation/Standard	Requirements	Implementation
Information, Communications and Media Act of Bhutan 2018	Lawful data processing, security measures, user rights protection, data localization	Consent management, data encryption, user rights portal, in-country data storage — each agency as data controller for its own Node's data
National Digital Identity Act of Bhutan 2023	Governs NDI authentication, consent, and citizen control of personal data	NDI is the identity basis for every submission at every Node (FR-7.2); consent is captured against the citizen's NDI identity at submission; only the attributes needed for feedback are requested from NDI (FR-9.1)
National Data Governance Framework 2025	Data minimization, purpose limitation, storage limitation, accountability, transparency	Privacy by design, data lifecycle management, privacy notices, DPO appointment at each agency, with a federation-level DPO coordinating cross-Node consistency
e-Government Security Standards	Access control, encryption, audit logging, vulnerability management, incident response	RBAC, TLS 1.3/AES-256, JLB integration, quarterly scans, SIRT team — implemented per Node against the shared national standard

GovStack Building Blocks	API standards, event-driven architecture, interoperability, security protocols	OpenAPI 3.0, CloudEvents, NDE compliance, OAuth 2.0/OIDC — the same specification followed by every Node
WCAG 2.1 Level AA	Perceivable, operable, understandable, robust web content for all users including disabilities	Screen reader support, keyboard navigation, color contrast, alt text, semantic HTML — built once into the shared codebase, inherited by every Node
ISO/IEC 27001:2022	Information Security Management System (ISMS), risk assessment, security controls	Security policies, risk register, control implementation, internal audits — each agency maintains its own ISMS scope for its Node, aligned to a shared national baseline

6.3.2 Audit and Reporting

Internal Audits

- Monthly compliance checks against the security controls checklist, at each Node.
- Quarterly access reviews: verify user accounts, roles, permissions still appropriate, at each Node.
- Quarterly security assessments: vulnerability scans, configuration reviews, at each Node.
- Quarterly federation-wide Node conformance re-check by the Federation Administrator (FR-7.6).
- Annual ISO/IEC 27001:2022 internal audit by an independent auditor, scoped per agency, with a federation-level summary compiled by GovTech.

External Audits:

- Annual penetration testing by a certified third party (CEH, OSCP), of each Node (or a representative sample of smaller Nodes, at the Federation Administrator's discretion, provided every Node has passed initial conformance certification).
- Annual WCAG 2.1 AA accessibility audit by a qualified accessibility auditor (e.g., IAAP-certified), of the shared codebase.
- Compliance audit by GovTech Bhutan on demand, of any Node.
- Support for oversight body audits (ACC, RAA, PAC) with federated audit trail access, coordinated through JLB and the federation layer.

Compliance Reporting:

- Monthly operational reports per Node: uptime, performance, incident summary.
- Quarterly security reports per Node: vulnerabilities, patches, incidents, risks.
- Annual compliance report: regulatory adherence, audit results, improvements, assembled nationally from every Node's own report.
- Data breach reports: within 72 hours to affected users and regulators, by the affected Node's operating agency.
- Annual transparency report: feedback statistics, handoff rates, improvements made, assembled by the federation layer from every Node.

6.3.3 Privacy Impact Assessment

A Privacy Impact Assessment (PIA) has been conducted, identifying the key privacy risks and mitigations below.

Privacy Risk	Impact	Mitigation
Unauthorized access to sensitive feedback at any single Node	Critical	RBAC with least privilege, service-level data isolation (structural, since there is no shared database), MFA for admin, field-level encryption for corruption reports at the owning Node

Data breach exposing PII at one agency's Node	Critical	AES-256 encryption at rest with agency-held keys, TLS 1.3 in transit, penetration testing, incident response plan, breach notification procedures — contained to the affected Node, since no other agency's data is stored on the same infrastructure or encrypted with the same keys
Excessive data retention at any given Node	Medium	Automated deletion after 7 years, archival to cold storage after 2 years, legal hold mechanism for ongoing investigations, enforced independently by each Node against the shared national policy
Lack of user control over personal data spread across multiple Nodes	Medium	Self-service access in MyBhutan for data access, rectification, erasure, portability, assembled via federated lookup; consent management; opt-out mechanisms, per Node or propagated federation-wide at the citizen's request
Third-party data sharing without consent	High	No third-party sharing except DPI components (with data sharing agreements); explicit consent for research/analytics use; federated queries are themselves access-controlled and logged (FR-9.9), so cross-agency sharing cannot happen silently
Inconsistent implementation across many independently	Medium-High	Mandatory Node conformance certification before joining the federation (FR-7.6); shared, versioned codebase and schema; quarterly

operated Nodes (new risk introduced by decentralization)		federation-wide conformance re-checks; smaller agencies may opt into a shared, GovTech-hosted Node deployment rather than operating their own infrastructure, without changing who is the legal data controller

7. Appendices

7.1 Glossary

Term	Definition
API (Application Programming Interface)	Set of protocols and tools for building software applications, enabling systems to communicate
CID (Citizen Identification Number)	Unique identifier assigned to Bhutanese citizens for government transactions
CSAT (Customer Satisfaction Score)	Metric measuring satisfaction with a service on a 1-5 scale, reported as percentage of satisfied users.
Handoff	The act of an FDBKB Node calling UCS's ticket-creation API to transfer a feedback item, with full context, into UCS's case-management system for resolution. The defining integration point between the two blocks (Section 4.1.6).
UCS (Unified Client Support)	The DPI building block responsible for resolving feedback once FDBKB transfers it there. See the UCS BRD for full specification.
Decentralized Concept	The FDBKB architecture whereby each agency operates its own Feedback Node with its own data, rather than all agencies sharing one central FDBKB database

DPI (Digital Public Infrastructure)	Foundational digital systems enabling delivery of government services at scale
Dzongkhag	District-level administrative division in Bhutan (20 dzongkhags total)
Federated Query / Federated Aggregation	A query or report that is assembled by calling multiple Feedback Nodes' own APIs and combining the (privacy-filtered) results, rather than by querying one shared database
Federation Administrator	The national role (GovTech Bhutan) responsible for the shared codebase, schema, security baseline, and Node conformance certification — distinct from an Agency Node Administrator, who manages a single agency's own Node
Feedback Node	A single agency's, ministry's, or dzongkhag's own independently operated deployment of the FDBKB codebase and data store
FDBKB (Feedback Block)	Communication building block for collecting, validating, scoring, and analyzing citizen feedback, implemented as a federated network of agency-operated Feedback Nodes. Response, escalation, and resolution are handed off to and owned by UCS — see this table's UCS entry.
GA System	A Government Agency's own internal system or application.
Gewog	Block-level administrative division in Bhutan, subdivision of dzongkhag
GovStack	Framework of reusable digital building blocks for implementing e-government services
JLB (Journaling Block)	DPI building block providing immutable audit logging for compliance and non-repudiation, fed by every Node independently

K-Anonymity	Privacy technique ensuring each record is indistinguishable from at least k-1 other records; applied both at each Node and again at federation-level aggregation
MoSCoW	Prioritization technique: MUST have, SHOULD have, COULD have, WON'T have
MyBhutan	Bhutan's national government services platform and the single citizen-facing channel for FDBKB.
NDI (National Digital Identity)	DPI building block providing secure digital identity and authentication services, used identically by every Node
NLP (Natural Language Processing)	AI technology enabling computers to understand, interpret, and generate human language
NPS (Net Promoter Score)	Metric measuring likelihood to recommend on 0-10 scale, calculated as % Promoters minus % Detractors
NTFB (Notification Block)	DPI building block for multi-channel notification delivery (email, SMS, push notifications), used identically by every Node
OAuth 2.0	Industry-standard protocol for authorization enabling secure delegated access
OIDC (OpenID Connect)	Identity layer on top of OAuth 2.0 enabling authentication and user profile retrieval
PSR (Public Service Registry)	DPI building block maintaining catalog of government services with metadata
RBAC (Role-Based Access Control)	Security approach restricting system access based on user roles and permissions

RPO (Recovery Point Objective)	Maximum acceptable amount of data loss measured in time (e.g., 1 hour), applied per Node
RTO (Recovery Time Objective)	Maximum acceptable time to restore service after disaster (e.g., 4 hours), applied per Node
SLA (Service Level Agreement)	Commitment specifying expected service performance and response times
SSO (Single Sign-On)	Authentication scheme allowing users to access multiple systems with one set of credentials
WCAG (Web Content Accessibility Guidelines)	International standards for making web content accessible to people with disabilities
NDE (National Data Exchange)	Bhutan's operational interoperability platform (WSO2-based) enabling secure, standardised information sharing between government systems, with an embedded API Gateway — the federation fabric that connects independent Feedback Nodes into one national capability

7.2 API Reference Summary

Complete API documentation for the shared codebase is available in OpenAPI 3.0 specification at each Node's own /docs endpoint (e.g., https://fdbkb.{agency}.gov.bt/docs), and a consolidated federation-level directory at https://api.fdbkb.gov.bt/docs lists every registered Node and the services it owns.

Key API Endpoints (identical contract at every Node):

Endpoint	Method	Description

/api/v1/feedback/submit	POST	Submit new feedback with ratings, text, category, attachments, to the owning Node
/api/v1/feedback/{id}	GET	Retrieve feedback by ID with full details and history, from the owning Node
/api/v1/feedback/{id}/status	PUT	Update feedback status (Agency Node Administrators only)
/api/v1/feedback/user/{userId}	GET	Personal feedback history at this Node; the federation layer assembles the full history across Nodes
/api/v1/feedback/{id}/handoff	POST	Transfer the item to UCS once the Agency Node Administrator confirms the handoff (FR-6.1), and store the returned reference
/api/v1/reports/generate	POST	Generate custom report with specified filters and format (PDF, Excel, CSV)
/api/v1/analytics/dashboard	GET	Retrieve real-time dashboard data (volume, sentiment, NPS/CSAT/CES) for this Node, filterable by date range

		and service, scoped to this Node's own data (FR-5.1)
/api/v1/federation/nodes	GET	List certified Feedback Nodes and the services each owns, from the NDE Node Registry (federation layer only; FR-7.5, FR-7.6)
/api/v1/federation/query	POST	Submit an authorized cross-Node query, using a scoped, short-lived token and logged to JLB, and receive aggregated, privacy-filtered results, with any unavailable Nodes reported as partial (federation layer only; FR-9.9, NFR-1.5)

- Authentication: all API requests require an OAuth 2.0 Bearer token in the Authorization header, issued via NDI and honoured identically by every Node.
- Rate Limiting: 100 requests/minute (basic tier), 1,000 requests/minute (standard tier), 10,000 requests/minute (premium tier), per Node.
- Response Format: JSON with UTF-8 encoding.
- Error Handling: RFC 7807 Problem Details format.

7.3 Error Codes

Code	Message	Resolution

FDBKB-1001	Invalid feedback submission	Check required fields: service_id, rating or text content. Verify data types and formats.
FDBKB-1002	Service not found	Verify service_id exists in the Public Service Registry / NDE Node Registry. Use /api/v1/federation/nodes to confirm which Node owns the service.
FDBKB-1003	User not authenticated	Submit via an NDI-authenticated MyBhutan session or, for API/batch submissions, include an NDI-verified CID and recorded consent (FR-7.2).
FDBKB-1004	Feedback not found	Verify feedback_id and that you are querying the correct Node (the ID prefix indicates the owning Node). Check if the user has permission to access this feedback.
FDBKB-1005	Duplicate feedback detected	Similar feedback recently submitted to this Node. Wait 24 hours or provide different details.
FDBKB-2001	Unauthorized access	The user does not have permission for this operation at this Node. Contact that Node's administrator for access.
FDBKB-2002	Invalid authorization token	Token expired or invalid. Refresh authentication token via NDI.
FDBKB-3001	File upload failed	Check file size (<2MB), supported formats (JPG, PNG, PDF), and network connection.
FDBKB-3002	Virus detected in attachment	The uploaded file contains malware. Scan file locally and upload a clean version.

FDBKB-4001	Rate limit exceeded	Too many requests to this Node. Wait before retrying. Current limit: 100 req/min (basic), 1,000 req/min (standard), 10,000 req/min (premium).
FDBKB-5001	Internal server error	Unexpected error occurred at this Node. Retry after 5 minutes. If persists, contact that Node's support with request_id.
FDBKB-5002	Integration service unavailable	External service (PSR, NDI, NTFB, JLB, NDE) temporarily unavailable. Retry after 5 minutes.
FDBKB-6001	Federated query partially failed	One or more Nodes did not respond within the federated-query timeout; the response includes a list of unavailable Nodes so results can be interpreted as partial. Retry, or narrow the query to available Nodes.
FDBKB-6002	Node not certified	The target Node has not completed federation conformance certification and cannot serve federated queries until certified (see FR-7.6).

7.4 Key Definitions

Feedback: Structured or unstructured input from citizens, businesses, or stakeholders about government services, institutions, personnel, or governance practices. Includes ratings, text comments, sentiment, and contextual metadata. Held at the Feedback Node of the agency that owns the relevant service.

Authenticated Feedback: the only submission mode — every feedback item is tied to an NDI-verified identity, through an NDI-authenticated MyBhutan session or, for API and batch submissions, an NDI-verified CID with recorded consent, enabling personalized follow-up and history tracking across every Node the citizen has used.

Sentiment Analysis: Natural Language Processing technique to determine emotional tone of text feedback, run locally at each Node using a shared national model. Classifies as Positive, Neutral, Negative, Mixed, Highly Critical, or Escalation-prone, with intensity and confidence scores (Section 4.1.3). Feeds FDBKB's own analytics.

Closed-Loop Feedback: The complete lifecycle from submission through acknowledgment, review, response, action, and impact measurement. The citizen experiences one continuous loop because FDBKB mirrors UCS's status rather than operating a second, independent loop.

Feedback Node: A single agency's, ministry's, or dzongkhag's own deployment of the FDBKB codebase, holding that agency's own feedback data independently of every other agency.

Federation Layer: The set of national, shared mechanisms (the NDE Node registry, the federated query broker, the cross-Node signal bus) that connect independent Feedback Nodes into one coherent national service.

Service Owner: Government agency or department responsible for a specific service, and for operating (directly or via a shared GovTech-hosted arrangement) its own Feedback Node. Primary consumers of feedback insights; improvement actions themselves are tracked in UCS once feedback is handed off there.

Agency Node Administrator (Feedback Reviewer): Agency staff responsible for reviewing and categorizing feedback within their own agency's Node, and for confirming or overriding the system's handoff recommendation (Section 4.1.6). Does not respond to citizens or manage case status once handed off — that is UCS's responsibility.

Web Collection: NDI-authenticated submission through MyBhutan, via widget or API.

Handoff: The transfer of a feedback item from FDBKB to UCS for resolution.

Pattern Recognition: Analytical capability to identify recurring issues, themes, or trends across multiple feedback submissions, performed locally per Node, with a federation-wide signal raised when the same pattern is detected across multiple Nodes.

7.5 External References

Reference	URL / Source
Bhutan DPI Blueprint	GovTech Bhutan, May 2025

GovStack Building Block Specifications	https://govstack.global/specifications
Information, Communications and Media Act of Bhutan 2018	Royal Government of Bhutan
National Digital Identity Act of Bhutan 2023	Parliament of Bhutan, 2023- https://parliament.bt/uploads/topics/16920888542499.pdf
WCAG 2.1 Guidelines	https://www.w3.org/WAI/WCAG21/quickref/
ISO/IEC 27001:2022 Information Security Standard	https://www.iso.org/standard/27001
CloudEvents Specification v1.0	https://cloudevents.io/
OpenAPI Specification 3.0	https://spec.openapis.org/oas/v3.0.0
OAuth 2.0 Framework (RFC 6749)	https://datatracker.ietf.org/doc/html/rfc6749
OpenID Connect Core 1.0	https://openid.net/specs/openid-connect-core-1_0.html
National Data Exchange (NDE) Architecture	GovTech Bhutan / WSO2 platform documentation
National Data Governance Framework 2025	National Statistics Bureau, July 2025
OWASP Top 10	https://owasp.org/www-project-top-ten/
RFC 7807 Problem Details for HTTP APIs	https://datatracker.ietf.org/doc/html/rfc7807

8. Annexure: Example Scenario for Feedback Block

This scenario illustrates a typical citizen journey: feedback is held at the owning agency's Node, federated via the NDE, and, where follow-up is needed, handed off to UCS for resolution.

Scenario 1: Citizen Submitting Service Feedback via MyBhutan (Thimphu Thromde's own Feedback Node)

Context: Tashi, a 35-year-old citizen from Thimphu, has just applied for a building permit through MyBhutan, signed in with NDI. He wants to give feedback on his experience.

On the confirmation page, Tashi clicks "Share Your Feedback." MyBhutan resolves, via the NDE Node Registry, that the Building Permit service is owned by Thimphu Thromde's Feedback Node. It loads the standard national form from that Node, pre-filled with the service name and his NDI-verified identity. If he had not been signed in, the same button would first prompt him to sign in with NDI. He rates the service across the standard dimensions, comments that processing-time expectations were unclear, and submits.

The Node validates the submission, runs sentiment analysis, and issues a Node-prefixed feedback ID (e.g., TT-FDB-2026-001234). It stores the record in its own database and logs the submission to the Journaling Block. Within seconds, it sends FDBKB's only citizen-facing message, the acknowledgment, via NTFB (FR-2.5). Because Tashi rated "Timeliness" below 4/5 and the comment was classified Negative, the Node recommends handoff under the Section 4.1.6 criteria. Thimphu Thromde's Agency Node Administrator reviews and confirms it.

The Node transfers the item to UCS with the feedback content, its sentiment and category, and Tashi's NDI identity and consent. UCS returns a reference, which the Node stores against the record. From here, FDBKB only reads status; it does not manage the case.

Later, Tashi checks his feedback in MyBhutan, which shows the current status inline, mirrored from UCS via the linked reference. When UCS records the resulting improvement, the public improvement showcase (FR-8.4) publishes it in anonymized form.

Outcome: FDBKB's role ended at acknowledgment and handoff; resolution happened in UCS.