Zorinthia

These examples illustrate the structure, depth, and type of output produced during a Phase 1 diagnostic. They are anonymised and provided to support decision-making — not as case studies or endorsements.

Digital Referral Platform: Analytics & Responsible AI Strategy Assessment

Independent Assessment for Clinical Collaboration and Referral Platforms

---

Background Context

The organisation is a health technology startup operating a digital platform enabling clinical case collaboration, secure patient referrals between facilities, specialist consultations, and sharing of diagnostic images and reports. **The industry** — digital health and tele-collaboration — sits between clinical care, health IT, and regulated medical device software. **The operational size** includes hospitals, clinics, and private practitioners using the platform to seek advice, escalate cases, and coordinate patient transfers across regions.

Adoption had grown steadily. The volume of referrals increased. Clinical engagement was high.

Yet leadership faced a strategic limitation:

> "We know we are facilitating care. We do not know precisely how well, where the bottlenecks are, or where the biggest clinical impact lies."

The internal engineering team could produce operational metrics (logins, message counts, uptime). However, deeper clinical, operational, and strategic insights were not being surfaced.

An independent data strategy advisor was engaged to:

- Design a structured analytics roadmap
- Establish data governance foundations
- Identify responsible AI use cases
- Ensure regulatory and ethical alignment

Executive Summary

This document summarises an independent assessment of analytics and AI readiness for a digital referral and clinical collaboration platform. The platform is widely used by hospitals, clinics, and private practitioners, but leadership lacked clear visibility into referral patterns, response performance, and clinical impact.

Key analytics questions that could not be answered:

- How many referrals originate from hospitals versus clinics versus private practices?
- Which regions generate the highest referral volume?
- What is the average response time by specialty and by time of day?
- Which doctors respond fastest — and does speed correlate with outcome?
- Where are referral loops occurring (back-and-forth without resolution)?
- What proportion of cases result in patient transfer versus remote resolution?

The data existed inside the platform — but without:

- Consistent facility classification
- Standardised doctor identifiers
- Clear referral lifecycle definitions
- Timestamp and data quality governance

analytics outputs would be misleading.

The advisor recommended establishing master data governance (doctor master index, facility master list, referral taxonomy), standardising the referral lifecycle, and implementing data quality controls as prerequisites for any advanced analytics or AI. Only then could the platform safely pursue governed AI use cases such as image triage, documentation assistance, and triage prioritisation.

Executive-Level Assessment Dimensions

The assessment is framed around five key dimensions:

- **Executive level pattern** — Clarity on which analytics and AI initiatives are materially aligned to clinical impact and business strategy
- **The organisation** — A multi-tenant digital health platform serving public and private providers
- **The industry** — Regulated healthcare and digital health, with strict privacy, safety, and ethical requirements
- **The operational size** — Growing referral volumes across regions, specialities, and facility types
- **The governance condition** — Data exists but master data, lifecycle definitions, and AI governance are immature

---

Analytics Questions Leadership Could Not Answer

Leadership wanted to move beyond platform usage metrics.

Strategic questions included:

- How many referrals originate from hospitals versus clinics versus private practices?
- Which regions and facilities generate the highest referral volume?
- What is the average time to first response by specialty, region, and time of day?
- Which doctors respond fastest — and does speed correlate with case resolution or outcome markers?
- Where are referral loops occurring (back-and-forth without resolution)?
- What proportion of cases result in patient transfer versus remote resolution?

These questions required:

- Consistent facility classification (hospital, clinic, private practice)
- Standardised doctor identifiers (single practitioner ID across accounts)
- Timestamp integrity (created, assigned, responded, closed)
- Clear, shared referral lifecycle definitions

The platform contained rich interaction data. Governance did not.

Data Landscape and Integration Challenges

The platform interacted with multiple data sources.

Internal Platform Data

- Referral requests (origin, destination, timestamps)

- Chat messages and case discussions
- Image uploads (e.g., radiology, dermatology, ophthalmology)
- Case attachments (PDF reports, lab results)
- User and role metadata (doctor, nurse, specialist, administrator)
- Audit logs (logins, access events)

External Data Sources

- Hospital information systems (HIS)
- Clinic EMRs
- Laboratory information systems (LIS)
- Radiology systems (PACS/RIS)
- National practitioner registries

Current challenges:

- Facility names entered inconsistently ("City Hospital", "City Hosp", "CH")
- Specialty labels varied ("Cardiology", "Cardiologist", "Heart Specialist")
- Duplicate doctor profiles (same practitioner, different email/phone)
- Patient identifiers sometimes incomplete, local, or anonymised for privacy

Without structured master data governance, any analytics (e.g., "top referring facilities") could produce distorted results.

Governance as the Foundation

The advisor recommended formalising three key foundations: master data, lifecycle, and data quality.

1. Master Data Domains

Domain	Description	Governance Needs
Doctor master index	Unique practitioner ID, specialty, region, facility affiliations	Controlled onboarding; deduplication; periodic review
Facility master list	Hospital, clinic, practice classification; location; ownership	Standard naming; type codes; region mapping
Referral type taxonomy	Reason for referral; urgency; specialty; care setting	Controlled list; change control; mapping from free text
Clinical specialty hierarchy	Hierarchy (e.g., Internal Medicine → Cardiology)	Single reference list; used across analytics and AI

This required:

- Controlled onboarding validation (e.g., registry verification)
- Periodic deduplication processes (merge duplicate doctor and facility records)
- Defined ownership of each master domain (data stewards)

---

2. Referral Lifecycle Definition

A standard lifecycle was defined to support consistent metrics:

1. **Referral created** — Request submitted by originator
2. **Specialist assigned** — Target specialist or team identified
3. **First response** — First clinically meaningful reply or acknowledgement
4. **Follow-up exchange** — Ongoing dialogue, additional images, questions
5. **Case closed** — Referral marked resolved or closed
6. **Outcome recorded** — Remote resolution vs. transfer vs. other outcome

Without this structure, response-time analytics and resolution metrics were inconsistent between teams and regions.

---

3. Data Quality Controls

Recommended controls:

- Mandatory timestamps at key lifecycle stages
- Automated anomaly detection for incomplete or stalled cases (e.g., created but never assigned; assigned but no response)
- Validation of facility types and specialties at data entry
- Role-based access logging to support privacy and audit requirements

In healthcare environments, analytics without quality assurance can create clinical and reputational risk.

Advanced Analytics Opportunities Identified

Once governance foundations were clarified, the advisor outlined structured insight opportunities.

Referral Flow Intelligence

- Referral origin counts and heatmaps by region and facility type
- Time-of-day and day-of-week response performance
- Specialty capacity analysis (volume vs. response time per specialty)
- Escalation frequency tracking (cases needing multiple handovers)

This would support:

- Workforce planning (where additional specialist capacity is needed)
- Capacity allocation (which regions or specialties to prioritise)
- Public-private partnership strategy (how public facilities and private specialists interact)

Clinical Efficiency and Collaboration Metrics

- Average time to first response by specialty, region, and facility
- Case resolution duration (created to closed)
- Repeat referral patterns (cases reopened or redirected)
- Specialist workload balancing (cases per specialist; outliers in response speed)

These insights could:

- Improve patient access to specialist input
- Reduce unnecessary patient transfers
- Highlight best-practice behaviour (fast responders with good outcomes)

AI Use Case Exploration (Governed and Ethical)

Leadership wanted to explore AI opportunities. The advisor emphasised that **AI must follow governance maturity**, not precede it.

1. Medical Image Analysis

Potential applications:

- Ophthalmology image triage (retinal scans for diabetic retinopathy)

- Dermatology lesion analysis (risk scoring; triage assistance)
- Radiology pre-screening (flagging possible fractures or critical findings)

Benefits:

- Flag urgent cases for faster specialist review
- Prioritise specialist attention
- Reduce turnaround time for high-risk cases

Requirements:

- Regulatory clearance as clinical decision support tools
- Bias and performance evaluation across populations
- Clear clinical oversight (AI as assistive, not autonomous)

---

2. Oncology Video and Imaging Review

Potential applications:

- AI-assisted review of biopsy imaging
- Video endoscopy pattern detection
- Tumour boundary estimation and change over time

These would function strictly as **clinical decision support** — aiding, not replacing, specialist judgement.

---

3. LLM-Based Clinical Documentation Assistance

Large language models could support:

- Summarising referral threads into structured case summaries
- Extracting key symptoms and findings from narrative messages
- Generating structured discharge summaries or referral feedback letters

Safeguards required:

- Strict data privacy controls and de-identification
- Human review and sign-off before communication
- Clear logging of AI-generated content vs. clinician edits

---

4. Triage Prioritisation Models

Predictive models could:

- Identify high-risk referrals based on pattern recognition
- Prioritise urgent cases in specialist queues
- Suggest optimal specialist matching based on case content and history

Prerequisites:

- Structured training data with outcome labels
- Bias monitoring and periodic model retraining
- Governance on acceptable false-negative/false-positive trade-offs

---

5. Population Health Insights

By aggregating anonymised referral trends, the platform could:

- Detect outbreak signals (e.g., sudden spike in respiratory cases)

- Identify chronic disease clusters by geography
- Support public health planning and resource allocation

Governance was critical to ensure **privacy, consent, and ethical use**.

Ethical and Regulatory Considerations

In a medical context, analytics and AI carry heightened responsibility.

The advisor recommended:

- Establishing a **clinical AI oversight committee** (clinicians, ethicists, data scientists)
- Documenting model validation processes and performance metrics
- Implementing robust data anonymisation and pseudonymisation standards
- Aligning with national health data regulations and guidance on AI in healthcare

Innovation without governance would expose the platform to reputational, clinical, and regulatory risk.

Strategic Roadmap

Phase 1: Governance and Foundations (Months 1–3)

- Define and implement master data domains (doctors, facilities, referral types)
- Establish referral lifecycle and mandatory timestamp requirements
- Set up data quality monitoring and anomaly detection

Phase 2: Core Analytics (Months 4–6)

- Build referral flow and performance dashboards (origin, response time, resolution)
- Implement specialist workload and capacity analytics
- Provide leadership with baseline referral and response metrics

Phase 3: AI Readiness (Months 7–9)

- Curate training datasets with governance approval
- Evaluate potential AI vendors or research partners
- Establish AI oversight processes and documentation standards

Phase 4: Pilot AI Use Cases (Months 10–12)

- Pilot one image analysis use case in a governed environment
- Pilot LLM-based documentation assist with limited user group
- Monitor performance, safety, and clinician feedback

Advanced AI deployment would only proceed after successful pilot evaluation.

Outcome

The engagement delivered:

- A structured analytics roadmap aligned to governance maturity
- Defined master data governance domains and data stewardship roles
- Prioritised referral performance dashboards and clinical efficiency metrics
- A phased, governed AI exploration plan

- An ethical oversight framework appropriate for medical context

Leadership gained clarity that:

- Analytics must be built on clean, governed data
- AI use cases must align with clinical value and regulatory standards
- Deeper insight is possible — but only with disciplined data architecture and governance

---

Takeaway

In digital health platforms, data is abundant but sensitive.

True strategic value comes from:

- **Structured referral lifecycle governance**
- **Master data discipline** (doctors, facilities, specialties)
- **Clinically aligned analytics** that answer real operational and care questions
- **Responsible AI deployment** under ethical and regulatory oversight

The advisor's role was not to introduce algorithms immediately.

It was to ensure that when advanced analytics and AI are deployed, they are defensible, ethical, and genuinely impactful to patient care.