ENTERPRISE PROGRAM BRIEF

AI-Ready Data and Knowledge Engineering with Fabric

Prepare trustworthy structured and document context for downstream AI consumers.

VNode ITeS | D03

DURATION	LEVEL	FORMAT
48 hours	Practitioner / specialist	Private cohort

AUDIENCE PROFILE

Built for these roles

Data engineers, Analytics engineers, Knowledge owners

OVERVIEW

Executive overview

Prepare trustworthy structured and document context for downstream AI consumers. Proposed instructor-led duration; delivery subject to environment and lab validation.

READINESS

Prerequisites

- Fabric foundations and practical PySpark
- AI and Generative AI Foundations, or equivalent demonstrable experience
- Microsoft Fabric Engineering Foundations, or equivalent demonstrable experience
- PySpark for Data Engineering, or equivalent demonstrable experience

PROGRAM OUTCOMES

Capabilities your teams will gain

- Prepare traceable structured and document context in Fabric
- Prove coverage, quality, freshness and consumer reference checks
- Test permitted data visibility and source withdrawal propagation
- Provide a semantic and index-handoff contract with operating ownership

ENTERPRISE CUSTOMIZATION

Enterprise customization

Tailor this program to your organization's priorities: Prepare trustworthy structured and document context for downstream AI consumers.

- Private cohort scheduling
- Equivalent-experience entry review

CURRICULUM

Curriculum roadmap

1

MODULE 1 · 8 HOURS

AI information requirements

Map what an AI consumer needs and identify unsupported assumptions.

- Consumer tasks
- Source inventory
- Access boundaries
- Freshness and answer contracts

2

MODULE 2 · 8 HOURS

Governed ingestion and provenance

Ingest synthetic content while retaining provenance and access metadata.

- Documents and tables
- Metadata
- Stable identifiers
- Source ownership

3

MODULE 3 · 8 HOURS

Content preparation and quality

Create validated context units and measure coverage and freshness.

- Parsing
- Normalisation
- Chunk boundaries
- Duplicate and stale content

4

MODULE 4 · 8 HOURS

Semantic context and permissions

Define business context and test visibility using separate identities.

- Business definitions
- Relationships
- Semantic models
- Persona tests

5

MODULE 5 · 8 HOURS

AI consumption and retrieval validation

Validate structured grounding and a documented document-index handoff; use Azure AI Search only if provisioned.

- Structured query grounding
- Document retrieval export
- Reference answers
- External index boundaries

6

MODULE 6 · 8 HOURS

Knowledge engineering capstone

Demonstrate updated and withdrawn content propagating through the data product and consumer tests.

- Refresh
- Deletion propagation
- Quality gates
- Operational ownership

PREPARATION

Environment prerequisites

Fabric capacity and isolated workspace Author and consumer personas Source files and notebook runtime Data-agent labs: paid F2+ or qualifying P1+ capacity, tenant AI settings and source access as documented; cohort capacity sized separately Delivery depends on environment and lab validation. This outline does not include step-by-step lab instructions.

CAPSTONE

Capstone

Demonstrate updated and withdrawn content propagating through the data product and consumer tests. Assessed AI-ready Fabric knowledge product Evidence of relevant access, quality and risk controls Demonstration and limitations record

ASSESSMENT

Assessment

Scenario-based practical assessment Prepare traceable structured and document context in Fabric 30% Prove coverage, quality, freshness and consumer reference checks 25% Test permitted data visibility and source withdrawal propagation 25% Provide a semantic and index-handoff contract with operating ownership 20% Proposed VNode outline-level assessment; no vendor accreditation or completed lab validation implied.

DELIVERABLES

Participant deliverables

- Assessed AI-ready Fabric knowledge product
- Evidence of relevant access, quality and risk controls
- Demonstration and limitations record

SCOPE

Scope boundaries

Step-by-step lab instructions; Runnable solution code; Packaged source datasets; Executed environment validation

DELIVERY

Delivery options

Count teaching hours, not calendar days. Exclude breaks. Split into 4-hour or 8-hour teaching sessions; schedule complex project work in contiguous blocks. Proposed hours exclude breaks and are subject to pilot validation.

PRIVATE DELIVERY

Plan a private cohort for your team

Scope delivery format, role mix, and rollout timeline with our solutions team.

Plan Private Delivery

Book advisory call