ENTERPRISE PROGRAM BRIEF

Microsoft Fabric Data and AI Solution Engineering

Deliver a governed service-operations analytics solution with a bounded conversational analytics use case.

VNode ITeS | D01

DURATION	LEVEL	FORMAT
48 hours	Practitioner / specialist	Private cohort

AUDIENCE PROFILE

Built for these roles

Data engineers, Analytics engineers

OVERVIEW

Executive overview

Deliver a governed service-operations analytics solution with a bounded conversational analytics use case. 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

- Deliver Fabric source-to-insight and conversational analytics
- Reconcile data, measures and data-agent reference answers
- Prove persona access and bounded read-only AI consumption
- Demonstrate recovery and provide a data-product runbook

ENTERPRISE CUSTOMIZATION

Enterprise customization

Tailor this program to your organization's priorities: Deliver a governed service-operations analytics solution with a bounded conversational analytics use case.

- Private cohort scheduling
- Equivalent-experience entry review

CURRICULUM

Curriculum roadmap

1

MODULE 1 · 8 HOURS

Business requirements and data product architecture

Translate operational questions into an architecture and measurable acceptance tests.

- KPIs and grain
- Source contracts
- Lakehouse versus warehouse
- Ownership

2

MODULE 2 · 8 HOURS

Incremental ingestion and orchestration

Build ingestion with controlled reruns and source reconciliation.

- Pipelines
- Watermarks
- Idempotency
- Failure recovery

3

MODULE 3 · 8 HOURS

Transformations and engineering quality

Build curated tables and test late, duplicate and invalid records.

- PySpark
- Curated tables
- Data quality
- Performance and cost awareness

4

MODULE 4 · 8 HOURS

Semantic and analytical consumption

Publish agreed metrics and test results across consumer personas.

- Semantic models
- Measures
- Data access
- Lineage

5

MODULE 5 · 8 HOURS

Governed conversational analytics

Configure a data agent and compare answers against reference queries.

- Fabric data agent
- Instructions and examples
- SQL/DAX answer verification
- Read-only boundaries

6

MODULE 6 · 8 HOURS

Solution capstone and operations

Demonstrate source-to-insight processing and recover from an injected failure.

- End-to-end runs
- Monitoring
- Quality gates
- Runbook

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 source-to-insight processing and recover from an injected failure. Assessed Fabric data and AI solution Evidence of relevant access, quality and risk controls Demonstration and limitations record

ASSESSMENT

Assessment

Scenario-based practical assessment Deliver Fabric source-to-insight and conversational analytics 30% Reconcile data, measures and data-agent reference answers 25% Prove persona access and bounded read-only AI consumption 25% Demonstrate recovery and provide a data-product runbook 20% Proposed VNode outline-level assessment; no vendor accreditation or completed lab validation implied.

DELIVERABLES

Participant deliverables

- Assessed Fabric data and AI solution
- 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