

Request for Proposals - Data Engineering, Analytics and Governance Partnership

RFP n#: 2026-1-FS-DATAGOV

Location: Global/Remote

Start date: 1st April 2026

End date: 31st March 2027

Technical Team: IT & Data



All applications must be submitted via this [form](#).

You may submit questions to tender@bettercotton.org – RFP n# 2026-1-FS-DATAGOV until 13th February 2026, noting that the **final submission deadline for bids is the 20th February 2026**.

Questions, requests and applications sent after the deadline will only be considered in exceptional circumstances.

Important Submission Process Information:

After submitting your details through the [form](#), you will receive a separate email to upload your supporting documents to a secure platform.

PLEASE NOTE THE DEADLINE HAS NOW BEEN EXTENDED UNTIL THE 27th February 2026. IF YOU HAVE ALREADY SUBMITTED A PROPOSAL, YOUR SUBMISSION HAS BEEN RECEIVED AND YOU ARE NOT REQUIRED TO SUBMIT AGAIN. BUT IF YOU WOULD LIKE TO ADD ANYTHING EXTRA OR CHANGE YOUR SUBMISSION, PLEASE SUBMIT A NEW BID VIA THE FORM LINK AS DONE BEFORE.

AT THE END OF THIS DOCUMENT, THERE IS A QUESTION AND ANSWER SECTION WHICH YOU ARE ENCOURAGED TO GO THROUGH IN PREPARATION FOR SUBMITTING YOUR BIDS

Description

The Better Cotton Initiative (BCI) is the world's largest cotton sustainability programme. Our mission: to help cotton communities survive and thrive, while protecting and restoring the environment. In difficult times, we are meeting the challenge head on. Through our network of

field-level partners we have provided training on more sustainable farming practices to more than 2.9 million cotton farmers in 26 countries. More than a fifth of the world's cotton is now grown under the BCI Standard and our membership network includes more than 2,400 members.

More information about BCI can be found on our website: www.bettercotton.org

BCI is seeking a qualified Microsoft Solutions Partner to provide expert-level data engineering, analytics support, and data privacy auditing. The goal of this engagement is to transition our current data landscape from a Azure environment to a unified Microsoft Fabric architecture while providing continuous support for internal integration projects and ensuring robust GDPR compliance.

Background

Our current infrastructure relies on Azure SQL Data Warehouse and Power BI for reporting and analytics. We have identified Microsoft Fabric F64 as our future-state platform to take advantage of its unified "OneLake" architecture and the elimination of individual Power BI Pro licenses for content consumers. We require a partner who understands the unique financial and operational constraints of the non-profit sector to act as an extension of our internal data team.

Scope of Work

The selected partner will be responsible for the following workstreams:

- **Fabric Migration:**
 - Key Responsibilities: Planning and executing the migration of Azure SQL Data Warehouse workloads to Fabric Data Warehouse/Lakehouse, including schema conversion and asset inventory.
 - Output:
 - Detailed Migration Roadmap & Inventory of assets
 - Validated Medallion Architecture (Bronze, Silver, Gold)
 - Technical Decommissioning Report for legacy Azure DW (Also called CottonCloud at BCI)
- **Ongoing Data Integration:**
 - Key Responsibilities: Proactive monitoring of existing Salesforce-to-Fabric pipelines; BCP-to-Fabric pipelines (BCP is our Supply Chain & Traceability system) managing API limits, handling sync errors, and field mapping updates.
 - Expected Outputs:
 - Monthly Integration Health & Performance Logs.
 - Updated Data Dictionary & API Usage Report.
 - Automated error-alerting framework for daily syncs
- **New Integration Projects:**
 - Key Responsibilities: Engineering new ingestion pipelines using REST APIs and middleware (Azure Logic Apps/Functions) for other source systems (both internal and external)
 - Expected Outputs:
 - Technical Design Documents (TDD) for each new pipeline.
 - Functioning ELT/ETL pipelines with 99.5% uptime SLA.
 - Documented API authentication and pagination logic.

- **Data Privacy & Governance:**
 - Key Responsibilities: Conducting quarterly data privacy audits, exploring & reviewing the need for implementation of Microsoft Purview for automated data classification, and ensuring all architectures meet GDPR and other local data privacy standards.
 - Expected Outputs:
 - Quarterly RAG (Red/Amber/Green) Audit Reports.
 - Remediation & Risk Mitigation Plan.
 - Data Lineage & Sensitivity Maps in Microsoft Purview. (If Purview emerges as a potential solution for our existing data governance needs)
- **Analytics Support:**
 - Key Responsibilities: Developing advanced Power BI dashboards and providing technical recommendations for performance tuning within the F64 capacity to prevent "usage creep".
 - Expected Outputs:
 - Library of Live, governed PowerBI dashboards
 - Quarterly capacity Optimisation & tuning report
 - Hands-on skills transfer workshops for internal staff

High-level Timeline

13th February 2026	Questions deadline All questions must be sent only to tender@bettercotton.org with the RFP Reference in the Subject line.
27th February 2026	Applications deadline All applications must be submitted via this form .
2nd March – 20th March 2026	Applications review & shortlisting / Interviews
By 31st March 2026	The successful applicant will be notified Unsuccessful <u>shortlisted</u> applicants will also be notified
1st Apr 2026	Start of the consultancy

Required Skills & Knowledge

Skills, Knowledge and Experience
Essential
Microsoft Fabric Authority: Expert-level proficiency in Fabric Data Factory, OneLake, and Synapse Spark.

Integration Mastery: Proven experience maintaining complex Salesforce integrations and managing Bulk APIs.
Technical Orchestration: Advanced skills in REST API integration, Service Principal authentication, and Azure Logic Apps middleware.
Regulatory Expertise: Deep understanding of GDPR standards and governance.
Sector Insight: Minimum 5 years of experience supporting non-profits and NGOs.
Accreditations: ISO 27001 or Cyber Essentials Plus certification.
Support SLAs: Guaranteed response times for critical integration failures (e.g., P1 issues resolved within 4-8 hours).
Skills Transfer: Commitment to co-working with the internal team to ensure long-term capability building.
Optional
Fluent English: BCI's language of operation is English

Application Requirements

Please note that we have changed our RFP submission protocol, and this is now in two phases.

- **Phase 1: Initial details will be submitted on the form found in this [link](#).**
- **Phase 2: You will receive an email with live links to upload relevant documents (please check your Spam and Junk folders)**

All Bidders are recommended to follow the categorisation for pricing as below. Any additional costs, rates or fees should also be provided as part of your Financial Proposal.

Preferably Bidders should categorise costs as follows:

- **Monthly Retainer Fee:** For proactive monitoring and maintenance of existing data integrations.
- **Standard Hourly/Day Rates:** For ad-hoc engineering on new integration and analytics projects.
- **Fixed-Engagement Fees:** For initial Fabric migration execution and quarterly privacy audits.

Proposals responding to this Request for Proposals should be a maximum of 10 pages (excluding CVs), and include the following:

- **Company Credentials:** Microsoft Solutions Partner designation with specializations in Data or Infrastructure.
- **Relevant Case Studies:** Three examples of Fabric migrations and CRM maintenance projects for similar organizations.
- **Operational Methodology:** Description of the "Fractional Expert" or Managed Service model, including communication protocols (e.g., Agile sprints).
- **Assigned Team:** Profiles of the primary data engineers and privacy auditors, highlighting relevant Microsoft certifications.

- Detailed and transparent budget, in EUROS, ideally aligned to the cost categorisation mentioned above.

We thank all applicants for their interest; however, only shortlisted applicants will be contacted.

BCI is committed to good practice and transparency in the management of natural, human and financial resources. All applications will be reviewed under the principles and subject to BCI's policies on equal opportunity, non-discrimination, anti-bribery & corruption and conflict of interest.

Evaluation Criteria

Proposals will be evaluated based on the following criteria:

Technical Evaluation Criteria (70%)

- Integration Capabilities Depth (25%): Proven ability to manage existing and new API-led connections.
- Fabric Architecture & Migration Success (25%): Quality of proposed migration and optimization strategy.
- Audit & Privacy Frameworks (20%): Strength of the quarterly audit methodology.

Financial Evaluation Criteria (30%)

- Quality and clarity of budget provided, and level of detail included
- Alignment of the budget to the activity timeline detailed in the technical proposal
- Value for money
- Adherence to the available budget

Questions & Answers For RFP 2026-1-FS-DATAGOV

1. Question 1

Number of data sources?

Answer

5 current data sources, with an average of 2 new sources added each year.

2. Question 2

No of existing ETL/ELTs

Answer

ETL/ELT processes are currently implemented using Azure Data Factory (ADF), Azure Web Apps, and some Fabric Dataflows (planned for migration).

3. Question 3

Daily Ingestion File/Data Size and weekly/monthly growth

Answer

All source data is fully ingested into the staging area daily. There is currently no change-detection-based loading mechanism implemented. Growth is not actively monitored, but only a small number of tables contain data in the millions of rows, and overall data volume remains small at this time.

4. Question 4

Existing DWH size, no of tables etc.

Answer

The current data warehouse is approximately 22 GB in size and contains 456 tables.

5. Question 5

Is Data cleaning and transformation implemented in Azure SQL, ADF or any 3rd party tool?

Answer

Data cleaning and transformation are performed using Azure Data Factory and SQL Stored Procedures.

6. Question 6

No of SQL Stored Procedures/Views

Answer

There are 57 stored procedures and 277 views currently in use.

7. Question 7

No of existing Power BI Measures, KPIs & Dashboards

Answer

There are approximately

- 32 PBI reports
- 7-10 Pages in each

8. Question 8

Number of business users

Answer

There are approximately:

- 250 internal Power BI users
- 500 external shared-report users
- Fewer than 10 data analysts accessing the SQL database directly

9. Question 9

Existing performance/architecture issues

Answer

The current platform is not prepared to scale for large incoming datasets from new sources. Batch transformation capabilities may be required going forward. There are also challenges with external dashboard-sharing strategy and performance.

10. Question 10

What are Data masking/encryption requirements?

Answer

No masking or encryption requirements exist at this time because no PII is stored in the data warehouse.

11. Question 11

The deadline for this RfP is 20th of February (page 1) however the forwarded links to upload the RfP documents state that the links will expire on 02/02/2026. We are assuming that these links **remain open until the stated deadline of 20/02/2026**. Could you confirm that?

Answer

The forwarded links to upload documents are auto generated and last for only 12 days. However, if the links expire before the application deadline and before you can upload your documents, do generate new links by clicking on the forms to apply again.

12. Question 12

There are currently two source systems in use at BCI: Chainpoint and Salesforce. Could you please share the remaining three source systems involved in the process?

Answer

In addition to Chainpoint and Salesforce, the remaining three source systems are:

- Certification-related data hosted in Azure SQL DB.
- API integration with the GoodGrants platform
- CottonCloudPortal Azure WebApp – used to upload Excel files, which are stored in SharePoint

13. Question 13

To properly assess the migration path to Microsoft Fabric (F64), please describe the current state of your Azure data environment. Specifically, can you clarify the deployment model of your Azure SQL estate?

Answer

Our Azure SQL estate currently operates with two environments—Test and Production—each deployed as a standalone Azure SQL Database. Deployments to both environments follow a

GitHub Actions–based CI/CD workflow: one dedicated to Azure Data Factory pipelines and assets, and another dedicated to the Azure SQL Database schema.

We have identified the need for a third Development environment to support isolated build and validation cycles, which we plan to introduce as part of our migration to Microsoft Fabric.

14. Question 14

- Is the environment based on:
 - Azure SQL Managed Instance,
 - Azure SQL Single Database / Elastic Pool,
 - Azure Synapse Analytics (Dedicated SQL Pool or Serverless SQL)?

Answer

Our solution is deployed on Azure SQL Database using the Single Database (Serverless) model.

15. Question 15

Can you share information about the APIs that BCI is currently using? Also, are there any issues or limitations with how these APIs work, such as slow performance, connection constraints, or missing features?

Answer

BCI currently uses two primary integrations:

Good Grants API: Stable and reliable over the past year with no technical performance issues. The only challenges relate to upstream data quality and governance, not the API itself.

CommCare: We are currently developing an integration using the OData feed, which has limitations in query capability, filtering, and reliability. This approach supports rapid delivery for the current use case, but we plan to transition to the CommCare REST API for more robust and reliable integration. We have not yet identified issues beyond the known OData constraints.

16. Question 16

Can you clarify the scope and expectations for managed services (covering monitoring, incident handling, and maintenance) of SalesforcetoFabric and BCPtoFabric integrations?-to-Fabric and BCP-to-Fabric integrations?

Answer

Integration between Salesforce to Cotton Cloud and BCP-Cotton Cloud already exist. For monitoring, we would like to focus on API consumption to prevent any service interruptions. We would also like to monitor the health of Fabric pipelines which includes monitoring the success of Copy data activities. As for incident handling, we would like high priority issues such as total failure of pipeline affecting a specific foundation data/critical data to be addressed within 1-2 hours. The incident handling would include technical troubleshooting of handling specific API failures or authentication failures. Ongoing maintenance tasks would include performance tuning, schema and mapping updates, credential and token management & data integrity which includes syncing deleted records from source system.

17. Question 17

Is there a preference for specific coding languages (e.g., Python, SQL, Java, etc.)?

Answer

Our current codebase primarily uses SQL for data operations and C# for application-level logic and integration work. We do not have a strict language preference and are open to incorporating Python, particularly for data transformation, and analytics use cases where it offers clear advantages.

18. Question 18

Is there a preference to reuse existing stored procedures, or can the partner rebuild logic using the best available techniques (for example, notebooks with PySpark)?

Answer

We prefer to retain existing stored procedures and integration logic where they are stable, reliable, and aligned with current business rules. However, we are open to rebuilding or modernizing logic where it provides clear benefits—especially for larger datasets such as inventory or field-digitization workloads that may be better suited for Spark-based batch processing.

We also recognize areas of technical debt and are open to working with the partner to identify where refactoring or re-engineering would be most beneficial.

19. Question 19

Is there a preference for batch processing or near–real-time or streaming ingestion?

Answer

We do not currently require real-time or streaming ingestion, we don't anticipate this need in the near future. Our workloads are well served by batch processing, and we see opportunities to re-engineer some logic using modern batch-processing techniques where it provides clear value. Any such changes have to be evaluated on a case-by-case basis.

20. Question 20

What is the expected frequency of daily or monthly loads?

Answer

Our current workloads run on daily scheduled pipeline executions, and we expect to maintain this same frequency going forward. At this stage, we do not anticipate a need for higher-frequency intraday or monthly-only load patterns, but we are open to evaluating adjustments if future use cases require them.

21. Question 21

Is there a requirement for data archiving and purge mechanisms?

Answer

We currently do not have a policy around data archiving. However, we will be working on it this year and will need this to be implemented on Cotton Cloud as well.

22. Question 22

Are user administration activities in scope for this project?

Answer

User administration and security is managed internally within BCI as a shared responsibility between the IT and Data teams. As such, we do not expect any user administration activities to fall within the partner's scope.

23. Question 23

Are security controls and security administration activities (at tenant level) in scope for this project?

Answer

User administration and security is managed internally within BCI as a shared responsibility between the IT and Data teams. As such, we do not expect any user administration activities to fall within the partner's scope.

24. Question 24

Is there a distinction in your evaluation between: - a supplier that operates fully on the Microsoft technology stack, has demonstrable experience with Microsoft Fabric, multiple relevant case studies, and strong technical capability, **versus** - a supplier that has paid for and obtained the formal Microsoft Solutions Partner designation?

Answer

We will prioritise a supplier with demonstratable Microsoft technology stack & integration experience first, but use the Solutions Partner designation as a secondary filter to ensure baseline of staff and organisational stability.

25. Question 25

Is this a binary requirement? Will proposals be assessed on a pass/fail basis regarding the Microsoft Solutions Partner status?

Answer

We will prioritise a supplier with demonstratable Microsoft technology stack & integration experience first, but use the Solutions Partner designation as a secondary filter to ensure baseline of staff and organisational stability.

26. Question 26

Can a submission still be considered if we do not currently hold the Microsoft-validated status (Performance, Skilling, Customer Success), with the understanding that this can be addressed if required? Notably, this investment has not previously been requested by other clients.?

Answer

We will prioritise a supplier with demonstratable Microsoft technology stack & integration experience first, but use the Solutions Partner designation as a secondary filter to ensure baseline of staff and organisational stability.

27. Question 27

What level of support coverage is expected, required, or considered desirable? Specifically, can you clarify the difference between your expectations for: - (1) *Global 24/7/365* support coverage, versus - (2) response and action aligned to specific time zones during standard working hours/days?

Answer

We are not looking for *Global 24/7/365* support coverage but business hours support (standard working days)

- **P1 (Critical):** Complete system outage (e.g., Salesforce sync down). **Target:** 15 minutes – 1 hour
- **P2 (High):** Function-wide application failure (e.g., Power BI dashboards not refreshing). **Target:** 1 hour-3hours
- **P3 (Medium):** Single-user issue with a workaround. **Target:** 1 – 2 business days
- **P4 (Low):** General inquiries or minor enhancements. **Target:** 2-5 business days

28. Question 28

As there is a significant difference in scope and cost between 24/7 global coverage and regionally based (e.g., UK/Europe) business-hours support, guidance on the minimum acceptable requirement would be appreciated?

Answer

We are not looking for Global 24/7/365 support coverage but business hours support (standard working days)

- **P1 (Critical):** Complete system outage (e.g., Salesforce sync down). **Target:** 15 minutes – 1 hour
- **P2 (High):** Function-wide application failure (e.g., Power BI dashboards not refreshing). **Target:** 1 hour-3hours
- **P3 (Medium):** Single-user issue with a workaround. **Target:** 1 – 2 business days
- **P4 (Low):** General inquiries or minor enhancements. **Target:** 2-5 business days

29. Question 29

For the 5 existing data source, how many entities (tables) are ingested from each?

Answer

- CommCare – 1
- ChainPoint – 25
- Intact platform – 35
- GoodGrants – 2
- SalesForce – 16

30. Question 30

Roughly how much data cleansing, transformation or business logic is applied to the ingested entities?

Answer

A moderate level of cleansing and transformation is applied. Simple idempotency checks, cross-validation logic are executed when data moves from Staging → Domain schemas.

31. Question 31

How do you cleanse and transform data? All through Stored Procedures?

Answer

All core cleansing, validations, and transformations are implemented using SQL Stored Procedures during the Staging-to-CottonCloud load. Lightweight business-rule filtering and computed fields are applied through SQL Views.

32. Question 32

Of the 456 tables in the data warehouse, what number are 'staged' tables, 'cleansed/transformed' tables and 'warehouse tables'?

Answer

Staging: 88
Cleansed: 184
Views: 279

33. Question 33

How many fact tables and dimension tables do you have? Are they all stored in a single dimension model?

Answer

We maintain a domain-oriented warehouse with 12 schemas, each representing a specific business domain.

Each domain contains its own fact tables and supporting dimension tables.

These are organized as multiple domain-specific dimensional models rather than a single enterprise-wide model.

Exact counts can be provided through metadata analysis if required.

34. Question 34

Are the 277 views all for individual Power BI report pages?

Answer

Not necessarily. While some of the 277 views directly support individual Power BI report pages, a portion of them are data-quality or intermediate views that are combined to produce final results.

For non-data-quality reports, we generally maintain a 1:1 relationship between views and report pages.

35. Question 35

Do you have a semantic model, separate to the data warehouse in Power BI for reporting?

Answer

We do not maintain a separate enterprise-wide semantic model in Power BI.

Our reports primarily connect directly to curated warehouse views, which already contain the required business logic.

In a few cases, we use small semantic models that are tied together solely to support shared filtering or lightweight modeling needs.

36. Question 36

Is the plan for the Power BI reports to stay exactly the same and be re-pointed from the migrated Fabric data warehouse?

Answer

We are working towards consolidating our PowerBi dashboards

37. Question 37

Do you have any pressing deadline or expected end-date for the engagement?

Answer

No, we consider this an ongoing engagement & partnership as we prepare to strengthen & scale our existing architecture for BCI's 2030 strategic goals.

38. Question 38

Please may you confirm whether proactive monitoring requires real-time automated alerting and immediate incident detection, or whether scheduled validation checks (e.g., daily execution and API health review with automated failure alerts) would meet the requirement.

Answer

Our current approach uses Teams-based automated notifications for pipeline run failures and dashboard refresh failures. We can review this setup and enhance it as needed.

39. Question 39

Please may you confirm what specifically must be reviewed as part of the quarterly privacy audit (e.g., architectural control validation, access reviews, data-level inspection, GDPR control mapping)?

Answer

The quarterly privacy audit is a comprehensive evaluation of the technical, operational, and regulatory controls governing our data ecosystem.

The scope of the audit is summarized across five critical pillars:

Architectural & Infrastructure Validation: This confirms the environment remains "secure by design" by verifying workspace isolation via Private Links and ensuring firewall rules restrict access to approved IP ranges. It includes a review of encryption for data at-rest and in-transit, as well as the security of integration points, specifically ensuring that middleware like Azure Logic Apps uses managed identities and that secrets are rotated within Azure Key Vault.

Identity & Access Management (IAM): The audit validates the Principle of Least Privilege (PoLP) by reconciling roles within Fabric and Power BI.

Data Integrity & Classification: Using Microsoft Purview, the audit confirms that sensitive data is correctly tagged with sensitivity labels. It also includes technical testing of Row-Level Security (RLS) to ensure they effectively restrict data access based on user roles.

GDPR & Regulatory Mapping: This focuses on legal accountability by updating "as-built" data flow maps from Salesforce through the Fabric architecture. The audit verifies the lawful basis for all processing, confirms readiness to fulfill Data Subject Access Requests (DSARs) within the mandatory 30-day window, and ensures appropriate safeguards are in place for any data transferred outside the EEA.

Monitoring & Log Review: The final pillar examines "footprints" in the system to detect anomalies. This involves auditing high-risk actions like mass data exports to Excel or CSV, performing targeted

Purview audit searches for sensitive report access, and verifying that incident response procedures are updated based on signals detected by Azure Monitor.

The output of this audit provides a RAG (Red/Amber/Green) status report, a prioritized risk remediation plan, and a documented accountability trail for regulatory bodies like the ICO.

40. Question 40

Please may you confirm the intended scope of GDPR compliance within this engagement. ?

Answer

The quarterly privacy audit is a comprehensive evaluation of the technical, operational, and regulatory controls governing our data ecosystem.

The scope of the audit is summarized across five critical pillars:

Architectural & Infrastructure Validation: This confirms the environment remains "secure by design" by verifying workspace isolation via Private Links and ensuring firewall rules restrict access to approved IP ranges. It includes a review of encryption for data at-rest and in-transit, as well as the security of integration points, specifically ensuring that middleware like Azure Logic Apps uses managed identities and that secrets are rotated within Azure Key Vault.

Identity & Access Management (IAM): The audit validates the Principle of Least Privilege (PoLP) by reconciling roles within Fabric and Power BI.

Data Integrity & Classification: Using Microsoft Purview, the audit confirms that sensitive data is correctly tagged with sensitivity labels. It also includes technical testing of Row-Level Security (RLS) to ensure they effectively restrict data access based on user roles.

GDPR & Regulatory Mapping: This focuses on legal accountability by updating "as-built" data flow maps from Salesforce through the Fabric architecture. The audit verifies the lawful basis for all processing, confirms readiness to fulfill Data Subject Access Requests (DSARs) within the mandatory 30-day window, and ensures appropriate safeguards are in place for any data transferred outside the EEA.

Monitoring & Log Review: The final pillar examines "footprints" in the system to detect anomalies. This involves auditing high-risk actions like mass data exports to Excel or CSV, performing targeted Purview audit searches for sensitive report access, and verifying that incident response procedures are updated based on signals detected by Azure Monitor.

The output of this audit provides a RAG (Red/Amber/Green) status report, a prioritized risk remediation plan, and a documented accountability trail for regulatory bodies like the ICO.

41. Question 41

Does the organisation have an idea on desired baseline quantity of ongoing consultancy days/heads per month/quarter, or are there also deliverables to be achieved under a set timeframe or set budget?

Answer

We are looking for a partner for our ongoing data engineering partners since our in-house team is lean. The focus for the next FY with the partner is to migrate to Fabric but scope will include ongoing support & development work from time to time. Our current engagement looks at 2 days per month for support and 40 days for the entire year budgeted for any platform development.