Service - Realtime Migrating DBs using Change Data Capture at Scale

Change Data Capture at Scale is a Lot 3 Cloud Support service that helps public sector organisations design, implement, and operate reliable low-latency CDC pipelines. It enables near real-time replication of database changes into analytics, integration, and event-driven platforms (for example, IBM DB2 or Oracle into Azure or GCP) without disrupting core transactional systems.

---

1. Service description

Change Data Capture at Scale provides consultancy and implementation services to move from batch-based data movement to streaming, log-based CDC pipelines. The service covers discovery of existing databases and integrations, CDC architecture design, technology selection, build, testing, and handover into operations. It is suitable for organisations building modern data platforms, real-time analytics, event-driven systems, or undertaking large-scale migrations where minimal downtime and reduced data-loss risk are critical.

The service uses open-source and cloud-native CDC tooling where appropriate and integrates with major cloud data warehouses, data lakes, and streaming platforms. Engagements are delivered by data and platform engineers working closely with internal database, integration, and operations teams.

---

2. Service features and benefits

2.1 Service features

- **Discovery and assessment of existing estates** – Analysis of current databases, integrations, and change patterns to understand sources, targets, volumes, and constraints.
- **CDC architecture design** – Design of CDC architectures for transactional, analytical, and event-driven workloads, including patterns for fan-out, replay, and backfill.
- **Technology selection and implementation** – Selection and configuration of appropriate open-source and cloud-native tools (for example, log-based CDC connectors and streaming platforms).

- **Design and configuration of log-based and streaming pipelines** – Implementation of CDC pipelines using database logs or equivalent mechanisms to capture and propagate changes reliably.
- **Integration with downstream platforms** – Integration of CDC streams with data warehouses, data lakes, messaging systems, and other endpoints.
- **Performance, scalability, and resilience testing** – Testing of end-to-end CDC pipelines for throughput, latency, failure handling, and recovery characteristics.
- **Observability design** – Design of metrics, logs, and CDC-specific monitoring (for example, lag, error rates, and throughput) for operational visibility.
- **Runbooks and operational playbooks** – Creation of runbooks for incident response, backfill, re-sync, and schema change handling.
- **Knowledge transfer and enablement** – Training and enablement for in-house data and platform teams to operate, troubleshoot, and extend CDC pipelines.
- **Governance and guardrails** – Guidance and patterns for schema evolution, retention, and data-quality controls aligned with governance and regulatory requirements.

2.2 Service benefits

- **Near real-time replication without disruption** – Delivers fresh data into downstream systems while leaving transactional databases focused on core workloads.
- **Reduced integration complexity** – Standardises on reusable CDC patterns instead of bespoke batch jobs and file-based interfaces.
- **Timely analytics and reporting** – Enables dashboards and reports based on up-to-date operational data, improving decision-making.
- **Improved resilience of data flows** – Uses tested failover, backfill, and recovery processes to maintain data continuity.
- **Less manual data movement** – Minimises manual exports and fragile batch file transfers, reducing operational overhead and error risk.
- **Faster delivery of modern data platforms** – Accelerates implementation of data lakes, warehouses, and event-driven architectures fed by CDC streams.

- **Enhanced observability** – Provides clear visibility of pipeline health, lag, and end-to-end performance, supporting proactive operations.
- **Lower risk of data loss during change** – Reduces risk during migrations and system modernisation by maintaining continuous change capture.
- **Stronger internal capability** – Builds internal skills to run and extend CDC solutions sustainably.
- **Alignment with governance and regulation** – Ensures CDC implementations respect governance, security, and regulatory requirements.

---

3. Service constraints

- Support and delivery are primarily remote; on-site engagements are available by agreement and may be scoped separately.
- Source and target systems must expose required logs, change feeds, or connectors to support CDC patterns.
- CDC scope excludes broader application feature changes that sit outside agreed migration or data-platform objectives.
- Cutovers and backfills are scheduled only within buyer-approved maintenance windows to control operational risk.
- Successful delivery depends on timely access to database, security, and operations staff within the buyer organisation.

---

4. User support

- **Email and ticketing support:** Available, with responses to tickets and questions within four working hours Monday–Friday, 9am–5pm GMT; outside these hours responses are best-effort and triaged on the next business day.
- **Phone support:** Available 9am–5pm UK time, Monday to Friday for active engagements.
- **Web chat support:** Available 9am–5pm, seven days a week with an AI chatbot, delivered on a platform designed to meet WCAG 2.2 AA accessibility standards (keyboard-only navigation,

screen-reader checks, high-contrast verification).

Change Data Capture at Scale uses a single standard support level included in day-rate or fixed-price engagement costs; there are no separate per-user or per-ticket fees. For larger programmes, an assigned technical account manager or lead cloud support engineer can act as a primary contact for planning, escalation, and service reviews, scoped and priced in the statement of work.