DATAFRAME

Service definition

G-Cloud 15 • 2026
Lot 3 (Cloud Support)

Synthetic Data Generation & Design Sprint (SDGDS)

We Don’t Just Deliver Data and Analytics.

We Build Capability.

Our Mission is Clear

To deliver modern, cloud-ready data solutions that empower your organisation.
Every engagement is designed not only to deliver outcomes, but to upskill your people, strengthen your internal capability, and create a future-proof data ecosystem

Building Skills That Last

As we guide your data journey, we also embed internal capability at every step through collaborative design, hands-on training, and real world application. Your teams don’t just consume data, they innovate, build, and become integral to shaping your data roadmap.

Driving greater ROI and creating a sustainable data ecosystem that empowers smarter decisions and long-term growth

Service Overview

Synthetic Data Generation & Design Sprint (SDGDS)

The Synthetic Data Generation & Design Sprint (SDGDS) is a rapid, structured engagement that enables public-sector organisations to design, generate and apply high-fidelity synthetic datasets for safe AI experimentation, analytics development and product prototyping. SDGDS accelerates innovation by providing privacy-preserving data assets, clear design direction and validated feasibility while aligning synthetic-data opportunities with organisational priorities, constraints and responsible-AI expectations.

We help teams harness synthetic data confidently by refining use-cases, establishing quality criteria, mapping risks, and defining the characteristics required to support realistic modelling and rapid product iteration. Through repeatable design practices that centre on trust, compliance and technical readiness, SDGDS enables long-term capability and reduces dependency on sensitive or inaccessible operational data. Delivered through modular, cost-effective discovery, generation and validation activities that create early momentum, SDGDS adapts as needs, datasets or resource availability evolve—allowing depth, pace or domain focus to shift without losing direction or delivery flow. The sprint is outcomes-led, powered by expertise in synthetic-data engineering, AI design, analytics development, data privacy, governance and value realisation.

We strengthen data-product operating models, establish coherent dataset specifications, uplift understanding of synthetic-data utility, and develop resilient design blueprints that support scalable analytics and AI initiatives. SDGDS integrates seamlessly with existing programmes and partners, promoting transparency, co-creation and structured knowledge transfer to build durable internal capability and confident long-term stewardship of synthetic-data assets across your organisation.

Our Team

Building Skills That Last

Our team isn’t a group of consultants. We are a close-knit, **family of data specialists** who have worked together across multiple projects, grown together, and developed a strong, unified way of working.

Over the years, we’ve built **deep trust, shared delivery patterns,** and a consistent approach that allows us to collaborate seamlessly and bring clarity, momentum, and confidence to every engagement.

All members of our team hold BPSS security clearance as a minimum, and we are willing to progress to higher levels of vetting, including DV, where a call-off requires it.

We **actively invest in our people** through continuous learning, structured development pathways, and dedicated time for innovation. This means the expertise you access is always evolving, always improving, and always grounded in real-world delivery experience.

Because our team grows and learns collectively, they bring that mindset into every partnership, openly sharing knowledge, **strengthening capability within your teams,** and helping them develop, enhance, and innovate with the tools and platforms we build together.

We believe the strongest data ecosystems are **built by empowered people.** That’s why we invest deeply in our own specialists and why we’re committed to developing capability within yours.

DATAFRAME

Features & Benefits

Synthetic Data Generation & Design Sprint (SDGDS)

SDGDS - 8 STEP PROCESS

Dataframe's Synthetic Data Generation & Design Sprint (SDGDS) provides a structured, accelerated approach to designing high-fidelity synthetic datasets. Using proven modelling frameworks, fidelity benchmarks and privacy-driven validation methods, SDGDS enhances realism, safety and speed. Our eight-step process defines requirements, designs generators and delivers governed synthetic datasets, enabling trustworthy AI experimentation, analytics development and innovation readiness.

Requirements & User Stories	Captures detailed user stories, analytical needs and expected outputs to ensure the synthetic dataset aligns with real service requirements.
Output & Success Criteria	Defines dataset outputs, quality thresholds, reporting needs and AI/analytics usability expectations to establish a clear target state.
Synthetic Model Design	Translates requirements into a synthetic-data model, shaping schema, entities, distributions and constraints that preserve privacy.
Metrics & Semantic Layer	Designs KPI logic, metrics rules and semantic hierarchies to ensure analytical consistency across BI, ML and reporting workflows.
Generation Method & Prototype	Selects generation techniques and produces an initial synthetic prototype aligned to the designed model and semantic layer.
Outlier & Fidelity Testing	Tests for outliers, distribution drift, correlation strength and privacy risks to ensure AI/ML and analytics fitness.
Refinement & Tuning	Iterates parameters, schema elements and semantic rules to meet fidelity, quality and responsible-AI thresholds
Packaging & Handover	Delivers the final dataset, semantic model, metrics catalogue and validation report, enabling confident use and future regeneration.

Value Proposition

Synthetic Data Generation & Design Sprint (SDGDS)

Structured Synthetic Schema & Modelling Alignment

- Transforms analytical and AI requirements into clear synthetic-data schematics, defining entities, relationships, distributions and privacy constraints to ensure the dataset aligns with real-world modelling and decision-making needs.

Consistent KPI Logic & Semantic Design Foundation

- Establishes a robust semantic blueprint—measures, hierarchies, rules and dimensional structures—ensuring consistent reporting, repeatable analytics, and stable AI/ML behaviour across regenerated datasets.

Metric Conformance & Analytical Reliability Assurance

- Validates KPIs and derived measures against the synthetic outputs to confirm logic integrity, semantic alignment and repeatability, ensuring synthetic datasets produce stable, predictable analytical outcomes.

Governed Packaging, Access & Operational Controls

- Delivers governed, production-ready dataset packages—including dictionaries, quality reports, lineage notes, and access guardrails—enabling secure use in analytics sandboxes, test environments and model labs.

Risk-Driven Validation for Privacy, Drift & Anomalies

- Runs structured fitness tests on generated data to detect drift, outliers, mode collapse and leakage risk, ensuring synthetic datasets safely support AI experimentation, analytics workflows and model training.

Capability Enablement & Scaled Synthetic Roadmapping

- **P**rovides handover, training, runbooks and future-state roadmaps that support ongoing regeneration cycles, cross-domain expansion, and long-term operational adoption of synthetic-data practices.

Synthetic Data Generation & Design Sprint (SDGDS)

	Features	Benefits
Schema Modelling	AStructured schema modelling aligned to analytical and AI requirements	Ensures synthetic datasets reflect real-world analytical behaviour
Semantic Blueprinting	Semantic blueprints defining measures, hierarchies and calculation rules	Strengthens metric trustworthiness across dashboards, reporting and forecasting
Fidelity Benchmarking	Configurable generation methods tailored to dataset privacy constraints	Improves AI readiness using reliable, realistic training datasets
Configurable Generation	Drift, anomaly and leakage detection through automated validation tests	Accelerates prototyping by removing dependency on sensitive data
Automated Validation	AI-aligned roadmap integrates future capabilities early.	Improved long-term value and product usability.

Synthetic Data Generation & Design Sprint (SDGDS)

	Features	Benefits
KPI Logic Testing	KPI logic validation ensuring repeatable, trustworthy analytical outputs	Enhances decision confidence through stable, repeatable KPI outcomes
Governed Packaging	Governed dataset packaging with dictionaries, lineage notes and controls	Supports governed adoption across analytics teams and service areas
Secure Access Controls	Secure access patterns with role-based usage restrictions applied	Simplifies compliance alignment through consistent metadata and lineage
Regeneration Workflows	Regeneration workflows supporting continuous synthetic dataset lifecycle updates	Increases modelling scalability for future AI and automation initiatives
Cross-Department Prototyping	Synthetic datasets enable safe multi-team analytics experimentation	Reduces cross-department risk while accelerating collaborative innovation

DATAFRAME

Service Management, Performance & Transition

Synthetic Data Generation & Design Sprint (SDGDS)

Onboarding & Initial Setup

Day 1–2: Kick-Off & Alignment

- Formal service initiation meeting
- Introductions to key stakeholders, product owners, and technical leads
- Review of organisational objectives, priorities, and existing work programmes
- Agreement of communication channels and escalation paths

Day 3–4: Access & Environment Readiness

- Setup of required user accounts, identity/access permissions, and secure connectivity
- Validation of source system access (APIs, databases, file shares, cloud resources)
- Review of existing documentation, architectures, governance standards, and backlog materials

Day 5–7: Planning & Backlog Formation

- Joint working sessions to review, refine, or establish the product backlog
- Confirm capacity model (team size, roles, cadence)
- Define acceptance criteria, quality gates, and documentation expectations
- Confirm reporting, sprint structure, and performance metrics

Day 8–10: Mobilisation & Go-Live

- Squads embedded into customer teams or delivery streams
- Delivery ceremonies initiated (stand-ups, planning, reviews, retros)
- First sprint or workload cycle starts
- Early value activities (quick-win analyses, pipeline fixes, or dashboard improvements) begin immediately

(Typical Squad Activities for Illustration)

Our onboarding process provides a fast, low-risk start, giving organisations immediate clarity and delivery momentum. The Service is operational within days, using a repeatable setup process that quickly establishes access, governance, scope, and ways of working.

Onboarding Objectives

The onboarding phase ensures:

- Required access, environments, and prerequisites are in place.
- Delivery squads align to organisational priorities.
- Governance, approvals, and communication channels are confirmed.
- Risks and dependencies are understood early.
- Internal teams understand how the service operates.
- A shared value roadmap is defined from the outset.

This reduces friction and enables measurable value from the first sprint.

Offboarding & Exit Management

Day 1–2: Exit Kick-off & Planning

- Formal exit initiation meeting
- Agree exit scope, timeline, and required activities
- Confirm stakeholders, responsibilities, and communication channels
- Review remaining backlog, risks, and dependencies

Day 3–5: Knowledge Transfer & Handover

- Deliver structured knowledge-transfer sessions
- Walkthrough of pipelines, dashboards, models, and technical components
- Handover of operational runbooks, documentation, patterns, and architecture diagrams
- Review of outstanding actions and closure plan

Day 6–8: Data, Artefact & Access Transition

- Provide all code, documentation, configuration, and datasets in agreed formats
- Confirm data ownership, retention, and any required deletion steps
- Transfer or revoke access credentials, accounts, and permissions
- Validate that internal teams can operate independently

Day 9–10: Finalisation & Service Closure

- Final backlog reconciliation and handover report
- Confirmation that all artefacts, knowledge, and data have been transferred
- Sign-off from the customer that exit criteria are met
- Optional transition into short-term hypercare, if required

(Typical Squad Activities for Illustration)

Our offboarding process delivers a smooth, secure transition that gives organisations full ownership of all data, artefacts, and knowledge. It is structured and repeatable, ensuring a clean service closure with minimal disruption.

Offboarding Objectives

The offboarding phase ensures:

- All documentation, code, data, and artefacts are fully handed over.
- Internal teams gain the knowledge needed to operate independently.
- Access rights are safely transferred or revoked.
- Remaining risks, dependencies, and recommendations are captured.
- Data ownership, retention, and deletion requirements are confirmed.
- A complete exit pack is provided for future continuity.

This ensures a controlled, low-risk transition and leaves the organisation with a stable, well-understood data ecosystem.

Service Levels (SLAs)

Our service levels provide clear expectations, consistent delivery standards, and transparent performance management. We operate with defined response targets, structured service management, and predictable support to ensure continuity and quality across all engagements.

SLA Objectives

The SLA framework ensures:

- Clear response and resolution targets for all priority levels.
- Agreed support hours and escalation routes.
- Regular reporting on progress, risks, and performance.
- Consistent delivery governance through stand-ups, reviews, and retrospectives.
- Reliable communication channels for fast decision-making.
- These standards give organisations confidence in service stability, rapid issue handling, and predictable delivery outcomes.

Service Levels (SLAs)

Priority Level	Description	Response Target	Resolution Targets
P1 – Critical	Complete service outage or major failure causing severe operational impact.	1 Hour	4 Hours
P2 – High	Significant degradation or data pipeline failure impacting essential services.	2 Hours	1 Business Day
P3 – Medium	Non-critical issues, partial degradation, reporting issues, or functional defects.	4 Hours	3 Business Days
P4 – Low	Minor issues, small fixes, documentation requests, enhancement requests.	1 Business Day	5-10 Business Days (depending on complexity)

Support Hours:

Our default support hours are:

- 09:00–17:00, Monday–Friday (excluding UK public holidays)
- Extended or out-of-hours support is available by prior agreement at call-off.

Service Governance

Escalation, Reporting & Service Assurance

Escalation Paths

To ensure timely resolution, clear escalation routes are provided.

- Level 1 – Delivery Squad Lead
- Handles day-to-day issues, queries, and operational blockers.
- Level 2 – Service Manager
- Escalation point for repeated issues, service performance concerns, or delays.
- Level 3 – Head of Delivery / Leadership Team
- Used for critical escalations, service breaches, or major programme risks.
- Customer escalation contacts will be agreed during onboarding.

Reporting Cadence

We provide structured and transparent reporting throughout the engagement:

- Daily stand-ups – Progress, blockers, priorities
- Fortnightly sprint reviews – Demonstrate completed work and agree next sprint focus
- Monthly service report – SLAs, risks, performance, delivery metrics, completed work, forward plan
- Quarterly roadmap review – Re-confirm priorities, capacity, and strategic direction

Additional reporting can be added as required at call-off.

Service Credits / Financial Recompense

- No financial recompense or service credits are offered as part of this service unless explicitly agreed in the order form under Call-off Schedule 14 (Performance Levels)

Scope, Accessibility & Constraints

(Lot 3 – Cloud Support)

Scope & Constraints
This service is delivered within the customer's cloud tenant using least-privilege access and customer-approved tools. It provides cloud support and enablement activities (advisory, configuration, engineering, analytics, governance, training) and does not include software resale, hosting, long-term managed operations or staff augmentation. Timely access to environments, data sources and decision-makers is required. Dependencies on third-party vendors or data availability may require re-planning and can impact timelines and deliverables. Where relevant, exit/transition will follow the applicable Call-off Schedules in the G-Cloud 15 pack

Security & Data Protection
The customer remains data controller; Dataframe acts as data processor only during the engagement. All work is performed in the customer's environment; no data is exported outside the tenant without written agreement. Access follows least-privilege and customer security baselines; data in transit/at rest uses encryption supported by the customer platform. We can support DPIA inputs and assurance reviews where required. Security and responsibilities at call-off will be governed by the appropriate Security schedules (e.g., Call-off Schedule 9a–9e) and other relevant schedules (e.g., IPR, BC/DR, Performance Levels).

Accessibility

- We provide documentation and training materials in accessible formats on request, and we design reporting and dashboard layouts in line with WCAG 2.1 AA principles where applicable to the service, supporting inclusive public-sector delivery and buyer evaluation requirements.

DATAFRAME

Modernising Public Sector Operations
Through Trusted Expertise
When You Need It

Thank you

www.thedataframe.co.uk