



G-Cloud 15

Lot 3 – Cloud Support Services (Additional Services)

Avanade Data Science and AI Services (Inc. Gen AI)



Table of Contents

1.	Service Overview	1
2.	Service Scope & Constraints	1
2.1.	Ideation.....	1
2.2.	Solution Design/Architecture.....	1
2.3.	Experiment/PoC/Pilot.....	2
2.4.	Collaborative Development / Training	2
2.5.	MVP Development	2
2.6.	Productionisation of Data Science Models	3
2.7.	Data Science Managed Services / Manage and Evolve.....	3
3.	Service Approach & Outcomes	4
3.1.	The Product approach, Agile and Dev Ops	4
4.	Tools, Assets & Accelerators	5
4.1.	Accelerators	5
4.2.	Applications of Data Science	6
4.2.1.	Customer Segmentation	6
4.2.2.	Fraud Detection.....	6
4.2.3.	Demand and Financial Forecasting	6
4.2.4.	Predictive Maintenance.....	6
4.2.5.	Churn Prevention.....	6
4.2.6.	Personalisation.....	7
4.2.7.	Product or Service Recommendation.....	7
4.2.8.	Price Optimisation.....	7
4.2.9.	Document Processing and Search	7
4.2.10.	Computer Vision	7
4.3.	Applications of GenAI	7
4.3.1.	GenAI Chatbots	8
4.3.2.	Intelligent Search and Summarisation	8
4.3.3.	Automated Email Response	8
4.3.4.	Employee or Customer Helpdesk Bot	8
4.3.5.	Entity Extraction.....	8
4.3.6.	Content Generation (language, image, video).....	8
5.	Trials & Associated Services	8
6.	Customer Responsibilities	9
7.	Avanade Experience.....	10
8.	Service Pricing	10
9.	Contact Information.....	10
10.	About Avanade	11

1. Service Overview

Avanade Data Science and AI Services (including GenAI) help organizations realize tangible benefits by applying advanced analytics, AI, and GenAI to solve business challenges.

Our services span ideation, solution design, experimentation, MVP development, and productionisation, delivering measurable outcomes such as improved decision-making, operational efficiency, and enhanced customer experiences.

Our solutions leverage Generative AI (GenAI), Retrieval-Augmented Generation (RAG), and agentic AI frameworks to deliver advanced solutions for business challenges, such as intelligent automation, conversational AI, and orchestration of complex workflows.

Use cases include GenAI-powered chatbots, RAG-based enterprise search, and agentic AI for automating multi-step business processes.

Clients can select the most relevant services based on their data science and AI maturity, ensuring a tailored approach that maximizes value.

2. Service Scope & Constraints

Which services are most relevant to you will be dependent on where you are in your journey of developing data science and AI (including GenAI) solutions. Clients may select the relevant service or services for their engagement. The following are the primary services that we offer in the data science space;

2.1. Ideation

We work with the client team to determine how and where in their organisation data science can be applied to realise benefits for the organisation, their employees, partners and customers or citizens. The ideation process is highly interactive and consists of the following steps;

- A. Collaborative discussion between Avanade and the client organisation team of the primary pain points for the organisation
- B. Identification of opportunities to address those pain points using data, technology, AI, GenAI and data science, converting ideas to use-cases
- C. Prioritisation of the use-cases identified based on a value versus complexity analysis
- D. High-level approach outline for the priority cases showing data and technology requirements and a solution approach from a data science perspective

2.2. Solution Design/Architecture

Working closely with the organisational team for whom the solution is to be developed, the solution design and architecture phase incorporates;

- A. Business process mapping to understand how the solution will embed into the operational processes it impacts

G-Cloud 15: Lot3 Cloud Support Services (Additional Services)

- B. User experience design for any personas interacting with the solution in order to ensure a positive user journey as they interact with the solution
- C. Data Science model design and deployment design to break down the functional requirements into technical approaches that will meet the requirements
- D. GenAI solution design and architecture with a keen focus on responsible AI and AI security
- E. Architectural design of the technology components, their integration, security considerations and user experience considerations (as well as other functional and non-functional requirement considerations).

2.3. Experiment/PoC/Pilot

The rapid development of a simple pilot or PoC for data science models or solutions to demonstrate the feasibility of the proposed solution and demonstrate that to the organisation's stakeholders to gain buy-in and grow confidence in the ability of data science to address the challenge. These projects are run under Agile methodology with the relevant Agile ceremonies. At a high level, the project covers the following elements;

- A. Requirements refinement with the organisation's stakeholders to decide on the focus of the scope of the short PoC or pilot project.
- B. Rapid development of the solution required, including training, testing and validation of the data science model(s) and GenAI solutions included in the scope of the project.
- C. Communication of the results of the PoC or pilot and validation of the performance of the solution relative to the expectations of the stakeholders.

2.4. Collaborative Development / Training

Where organisations are interested in enhancing the skills within their own team, we can offer a collaboration based training opportunity. The structure of this is through the creation of a joint team of Avanade and client team members to work together on a specific project for a period of time (recommended 10 weeks minimum).

- A. Technical training is focused on data science, AI and GenAI on Azure. The approach works well for team members who already have knowledge of coding, AI, software engineering and data science from previous work.
- B. Shadowing of Avanade resources and collaborative sharing of workload for development activities with multiple daily checkpoints.
- C. Recommendations for additional training in areas relevant to the user's specific role, where gaps are identified, or the specific project on which the team is collaborating do not cover all elements recommended for that role type in the client organisation.

2.5. MVP Development

The development of a minimum viable product and deployment of that to an environment where it can be used by a group of end-users or be embedded in a business process. Typically, an MVP will require more software engineering and automation on top of data science development activities and, as such, contains a broader scope than found in the PoC or pilot projects. This process also uses Agile methodology, and the role of the client product manager becomes more important.

G-Cloud 15: Lot3 Cloud Support Services (Additional Services)

- A. Solution Design, including user experience design, is a collaborative process working with the organisation's stakeholders to ensure the product matches the needs of the users.
- B. DevOps, AIOps and MLOps best practices are followed to ensure, from a software engineering perspective, that the product can be iteratively developed by a multi-disciplined team, including data scientists, engineers and front-end developers.
- C. We follow an approach of "develop iteratively, release often" so that the users and stakeholders have the opportunity to feedback on features released each sprint, enabling the team to refine the solution based on that feedback.
- D. For products which have a user interface, we follow a "user-centred design" approach keeping frequent collaborative contact with the relevant user representatives.
- E. Accelerators for GenAI solution deployment, data science model development, deployment and any other component of the solution will be used by the delivery team where possible to accelerate the MVP development process.

2.6. Productionisation of Data Science Models

When models are proven to be effective for the intended purpose, either through experimentation or pilot/PoC development processes, our team can productionise the models and related applications and integrations required to realise the benefit they are designed to enable. Productionisation uses the same Agile methodology, and DevOps best practices discussed in MVP development but also meets some additional requirements;

- A. Automation of the data pipelines and execution of the models in order to ensure intervention-free operation of the models under normal conditions.
- B. Monitoring of the solution as a whole is enabled so that support or managed services teams will be alerted if the solution stops working or the performance of the model degrades.
- C. Logging mechanisms are put in place to ensure that all scoring of the model can be recreated in the future. This is particularly important for regulated industries where an audit train of data used in decision making is required.
- D. Responsible AI assessment of every use-case forms a key part of our entire process through which each use-case is scored against a framework covering impact on people, society and the environment.
- E. Code management is key to all of our development activities, whether it is a PoC, MVP or production solution. The team will either work with the client's CI/CD and code management practices or define them for the project.
- F. Security from a network, data, storage and access perspective is also a key component of all solutions, whether PoC, MVP or production solution.

2.7. Data Science Managed Services / Manage and Evolve

Avanade offers customers the option to avail of our best shore model for management and operation of Azure-based data science platforms and solutions developed by our project teams. This service can be delivered from multiple geographies depending on the needs of the customer. The service consists of;

- G. Monitoring and Alerting of the infrastructure, data pipelines, models and applications to ensure continuous operation within agreed SLAs.

- H. Maintenance of the solutions, including any patching required to ensure continued support and licencing as required.
- I. Tuning of the solutions in order to ensure continued performance of the solution within target SLAs related to performance.
- J. Application of fixes required where any element of the solution stops working under operational use under the conditions of use agreed in the service agreement.

3. Service Approach & Outcomes

3.1. The Product approach, Agile and Dev Ops


While in traditional business intelligence and analytics, we generally develop data models, reports and visualisations, in the field of data science and AI, we are more frequently developing an integrated solution or product. At times we may carry out investigative analytics and calculate probabilities as part of an experiment to find value or test a hypothesis. Typically, however, the aim is to build something that can be consumed; a product. As such, the structure of data science projects is different, as are the skills required and the way in which the output is consumed. While reports may be used to analyse the performance of the model in the process in which it operates, it is generally more likely that the output will be consumed by another application to assist in making an automated or manual decision.

It is best to illustrate this using an example, let's say Fraud Detection in an Insurance Claims department. Data science experiments will let you test the hypothesis of whether you can use statistics and data to find fraud. The product, however, is some combination of components, including, at a minimum, the engine (executable data science model that has evolved from your experiment and the execution framework and resources), which processes a new claim to determine whether it is likely to be fraudulent and warrants investigation or some other action. To support the training of the model, data pipelines, including transformations and the generation of complex features, are required.

In order to feed a new case to the engine for scoring, there is a requirement for an ability to send a message with the appropriate case data to the engine and to receive a response with the appropriate score (typically using an API framework). Then depending on the score, there may be a workflow to decide if a specific action is required and to automatically take that action. Where that action requires a person to interact with the case, perhaps for investigation of whether it is truly a fraud, there may be a new interface that is required as part of the product. These components together give you a solution that can be used, and that will deliver value. As such, the data science model or GenAI model is just a piece of your final product. Alone it does not necessarily give you all the desired value.

We typically use the Agile methodology in the delivery of our projects and follow DevOps best practices from a software engineering perspective.

G-Cloud 15: Lot3 Cloud Support Services (Additional Services)



MVP Development

- Follow Agile method to support rapid, cyclical development of the initial product.
- The “product” approach ensures alignment to developing something that can be integrated with business processes even as an MVP (supporting AB testing).
- Understanding of how the product will be used in autonomous, semi-autonomous or manual business processes, and designing the MVP accordingly.
- Technical and product teams work closely together to ensure alignment and retain the ability to adjust the product as needed.
- The product owner prioritises development activities of the joint team.
- Evaluate appropriate data sets, features, algorithms and models to obtain the appropriate performance.
- Test plan evolves with the MVP to ensure full testing is complete before release 1.

Dev Ops

- Continuous integration of the solution from the initial release.
- Release management in an agile and flexible manner.
- Monitoring and operation of the product feeding back to the development cycle.

Figure 1: Agile Development and Dev Ops

Responsible AI and AI security are embedded throughout our delivery lifecycle. Every use case is assessed against a Responsible AI framework, ensuring ethical, secure, and compliant deployment of AI technologies.

We prioritise AI security, data privacy, and regulatory compliance, leveraging Microsoft Purview and industry best practices.

4. Tools, Assets & Accelerators

4.1. Accelerators

From our many years of experience in data science, we have developed service accelerators to support more rapid development and deployment of data science solutions in cloud environments. These give us consistency in our approach as well as ensure key elements are included in the architecture from the outset. All of these are deployed through Azure DevOps.

Architecture Blueprints: Standardised architectural blueprints for data science projects incorporating best design principles for data science experimentation, model deployment and model monitoring and maintenance.

Infrastructure as Code: Terraform based deployment of architectural blueprints to accelerate the process of getting from PoV or MVP to production.

Data Science Project Templates: Python and R templates for project code to ensure modularisation as well as ease of deployment of code to the relevant point in the corresponding architecture blueprint.

GenAI Project Templates: Software and Infrastructure templates for common GenAI use cases such as RAG, chatbots, summarisation, entity extraction, and language similarity.

Solution Accelerators: Specific accelerators for solutions such as price optimisation, demand forecasting, recommender engines, predictive maintenance, fraud detection, document processing, NLP and search.

4.2. Applications of Data Science

4.2.1. Customer Segmentation

Through the use of advanced analytics, we enable businesses to segment their customer base in order to get a better grasp of distinct customer profiles that exist naturally within a population. This allows businesses to customise their products and services to better address the needs of distinct customer profiles. Furthermore, it supports the improvement of the customer experience through tailored interaction from digital or other perspectives.

4.2.2. Fraud Detection

Our predictive analytics capabilities detect fraudulent and risky transactions from legitimate ones by applying machine learning methods to historical transaction data. This allows businesses to reduce risk and associated costs by adjusting their operational procedures in detecting and dealing with fraud. This model uses features from a wide selection of internal and external (publicly available) data sources. It also considers features extracted from free-text using natural language processing. This broad feature space enables us to remove features which may perpetuate bias and protect us from bias that may exist in particular features of the historic true positive cases, whether we are aware of it or not.

4.2.3. Demand and Financial Forecasting

Powered with Artificial Intelligence algorithms and historical data, we have been able to forecast customer demand and participation in specifically tailored promotional programmes empowering businesses to optimise their supply chain operations. In this endeavour, to really understand customer demand and shopper behaviour, we have been able to draw on an array of sources and integrate country-specific data as well as global macroeconomic metrics to elucidate cultural and regional differences in how customers approach shopping. Similarly, we have used statistical modelling and predictive analytics to forecast profit and help our clients fine-tune their revenue streams and sales strategies.

4.2.4. Predictive Maintenance

At Avanade, we develop predictive maintenance solutions using sensor data from IoT (Internet of Things) devices in machines. Typically, we use two approaches to do this. Firstly, by analysing the patterns in the data leading up to previous failures, we can use supervised machine learning to teach the solution to identify similar failures likely to occur in the future. This approach captures the previously known failures but misses new failures not previously observed in the data used to train the model. For "new" failures, we use anomaly detection. In this scenario, we look for patterns of behaviour that are not deemed "normal" for the machine. Combining these two, we can help businesses to decrease the downtime of machines, optimise maintenance regimes, re-design products and improve customer experience.

4.2.5. Churn Prevention

Working with a variety of customers, we have developed algorithms that predict customer churn and employee attrition. The models combine customer data and usage data with a variety of external data sets to forecast the propensity of each individual to churn. The solution allows marketing, customer retention and HR activities to be targeted at individuals with the highest return rates. The approach we follow results in an output which indicates which elements in the data are mostly influencing the score predicting the churn so that the retention team can understand causality. An extension to the churn model is the best action model. This is a recommendation engine which informs the retention team what action is best used to retain a customer like the one you are dealing with who is interacting in a similar way with your service. This solution prevents churn and therefore improves revenue, customer experience and margin (through a reduction of the cost of customer acquisition).

4.2.6. *Personalisation*

In order to improve the experience of customers, patients, and citizens engaging with an organisation, we assist clients with developing personalisation solutions to enable the surfacing of the most relevant content, products, services or experiences to each individual. This personalisation can have different focuses in order to support the user journey, whether it is about expediting the completion of a task, accessing the most relevant information, sharing the most relevant, interesting or novel product or service or creating an experience that is positive, easy and beautiful for the particular user.

4.2.7. *Product or Service Recommendation*

At Avanade, we have significant experience in developing personalised product recommendation solutions in the retail space. Product Recommenders are used in consumer business for a variety of reasons, including to increase basket size, to increase basket revenue, to improve basket margin and to improve customer experience. In developing these solutions, we use a mix of approaches, including market basket analysis, matrix factorisation and collaborative filtering. As the number of historical transactions of an individual increase, the recommendations become more personalised. In fact, the solution selects the appropriate model in the engine to execute depending on the individual shopping and their shopping history.

4.2.8. *Price Optimisation*

Our teams have been developing price optimisation solutions across industries (mainly retail, travel, transport and hospitality) for many years and have published papers sharing thought leadership in this space. We can support the development of price optimisation solutions and can bring accelerators to support more rapid development, testing and productionisation of pricing engines to the table. Depending on the environment, price optimisation needs to consider constrained versus unconstrained supply or demand, price sensitivity, elasticity and other factors in determining the optimal approach to solving the problem.

4.2.9. *Document Processing and Search*

Our teams work with the processing of scanned documents to extract the information required for analytical or operational processes. Using techniques such as optical character recognition (OCR), entity extraction, natural language processing (NLP), image recognition and others, we can create a digital version of the important information from documents. We then are able to build applications to enable users to interact with, search or analyse the data extracted to support relevant business processes, decision making or automated actions.

4.2.10. *Computer Vision*

At Avanade, we have a specific Azure-based architecture for the processing of videos and images to support entity identification, tracking and information extraction. This can be used in many different ways, including tracking wayfinding in buildings, identifying people or objects, analysing movement or extracting specific information with regard to an item in the video or image, such as size, direction, counts or extraction of text. Using different types of sensors or cameras, different versions of images can be processed to support different use cases.

4.3. Applications of GenAI

Avanade provides accelerators and templates for rapid deployment of GenAI solutions, including chatbots, summarisation engines, entity extraction, and language similarity models. These are built on Azure and integrate with enterprise data sources

4.3.1. *GenAI Chatbots*

Self-service bots with “ChatGPT”-like interfaces to enable users to find relevant information from internal and external sources and summarise the relevant information for the user. This includes the ability to provide citation or references for where the information was retrieved from.

4.3.2. *Intelligent Search and Summarisation*

Searching across the enterprise knowledge base for information to support your tasks. This solution improves productivity of the user where tasks could be accelerated by more rapid retrieval of information across the organisation. This can be extended to generate content in a specific format e.g. pitch decks, contracts, reports.

4.3.3. *Automated Email Response*

Suggested email responses that are personalized and context specific, automated client read-outs, meeting preparation information, meeting summaries, and action items. Auto-generated next-best action, upsell, cross-sell opportunities

4.3.4. *Employee or Customer Helpdesk Bot*

Cross functional front line helpdesk enabling employees or customers to rapidly find relevant information from a variety of sources including information related to services, pricing, policy, process, how-to guides. This can be extended to automate actions such as leave booking, service details, contact request or password resets by triggering automated workflows behind the scenes.

4.3.5. *Entity Extraction*

Extracting key information from complex documents to support financial or other processes which require the data in a structured format. The AI automates the process accelerating an often tedious, manual process.

4.3.6. *Content Generation (language, image, video)*

Generation of new content based on existing examples with additional information provided by the user. This can include language based content or image and video based content. This content can be aligned to company brand and tone of voice to ensure a consistent brand experience.

5. Trials & Associated Services

Avanade offers a suite of services to help clients rapidly evaluate, validate, and realize the benefits of Data Science, AI, and GenAI. These services are designed to minimize risk and maximize value during the early stages of engagement:

Proof of Concept (PoC):

- Rapid development of PoC and pilot projects for targeted data science, AI, and GenAI use cases
- PoCs are structured to demonstrate technical feasibility, secure stakeholder buy-in, and build confidence in the effectiveness of advanced analytics, GenAI, and intelligent automation

G-Cloud 15: Lot3 Cloud Support Services (Additional Services)

- Delivered using Agile methodology, PoCs include collaborative workshops, analysis, and design sessions

Discovery Phases & Workshops:

- **Insight Discovery Workshops:** Used to define engagement scope, clarify challenges, and establish architectural approaches and roadmaps for data science and AI solutions
- **Prioritisation workshops:** Identify and quantify use cases that deliver value, supporting effective project selection before or during engagement
- **Collaborative Development/Training:** Joint teams of Avanade and client personnel work together on specific data science, AI, or GenAI projects, enabling clients to experience Avanade's delivery approach and build internal skills for advanced analytics and GenAI adoption

Hero Offers:

- Hero offers are structured, short, focused engagements (often 4–6 weeks) to demonstrate tangible business outcomes early
- Scoped to provide clarity and defined outcomes for a fixed fee, helping clients rapidly evaluate and realize the benefits of Avanade's Data Science, AI, and GenAI

6. Customer Responsibilities

- **Provision of Resources:** Customers are expected to provide necessary resources, such as access to systems, data, and relevant personnel, to enable Avanade to deliver the service effectively.
- **Collaboration and Engagement:** Active participation in collaborative workshops, ideation sessions, and ongoing feedback cycles is crucial. This includes engaging with Avanade during discovery, architecture design, and MVP development phases.
- **Governance and Roles:** We'll work with you to identify the most appropriate governance structures to ensure quick decision making, issue resolution and to ensure appropriate quality and financial controls
- **Training and Skill Development:** Where collaborative development or training is part of the engagement, customers should allocate team members to participate in joint development and training activities.
- **Operational Support:** For managed services, customers may need to support monitoring, maintenance, and data stewardship activities, ensuring that operational dependencies and SLAs are met.
- **Timely Decision-Making:** Customers are expected to provide timely decisions and approvals, especially during agile development cycles and migration projects, to avoid delays.
- **Compliance and Security:** Customers must ensure compliance with internal policies and regulatory requirements, particularly regarding data access, security, and GDPR.

7. Avanade Experience

AI-Driven Analytics in Regulated Industries

Healthcare Regulator: Built an Azure-powered data platform with Power BI, enabling real-time monitoring, risk-based inspections, and a shift to a product-driven, data-centric culture.

Generative & Agentic AI in Recruitment

Global recruitment firm: Developed "AI ASSIST" using agentic AI to automate CV summarization, PII redaction, candidate pitch generation, interview question creation, and client intelligence gathering. The solution orchestrates complex tasks across multiple systems, improving consultant productivity, compliance, and client experience.

Further details of Avanade's experience, client references, and additional case studies are available and can be shared upon request to demonstrate our proven track record and the breadth of our expertise in delivering successful solutions for both sector-specific and cross-industry engagements.

8. Service Pricing

Please refer to the associated Pricing Document relevant to this Service.

9. Contact Information

Sangeeta Pierce

Head of UK Health & Public Sector

Email: uk.hps.support@avanade.com

Tel: 020 7025 1000

10. About Avanade

Avanade, a joint venture between Accenture and Microsoft, is a privately held company was founded in 2000 with the goal of delivering innovative services and solutions to enterprises worldwide using the Microsoft platform. Avanade's main business focus is to purely deliver innovative services and solutions to enterprises worldwide on Microsoft technology. Avanade is a global organisation with over 56,000+ professionals worldwide, serving our clients in major geographic business areas in 26 countries.

This vast network of highly skilled resources is further complemented by our network of delivery centres that we refer to as Advanced Technology Centres (ATCs). This complementary capability provides the agility, cost efficiency and diversity of skills that today's businesses demand. This construct underpins the results we generate for our clients and forms the foundation for our long-term relationships with them.