

ANTHROPIC

Anthropic Frontier AI: Agentic AI Collaboration and Workflow Integration – Claude

Lot 2b Software as a Service (SaaS)
Service Definition

January 2026

1. Introduction	3
Company Overview	3
Value Proposition	3
What the Service Provides	3
Social Value	3
Overview of the G-Cloud Service	3
Associated Services	3
2. Data Protection	3
Information Assurance	3
Data Back-Up and Restoration	3
Business continuity statement/plan	4
Privacy by Design	4
3. Using the service	4
Ordering and Invoicing	4
Availability of Trial Service	4
On-Boarding, Off-Boarding, Service Migration, Scope etc.	4
Training	5
Implementation Plan	5
Service Management	5
Service Constraints	5
Service Levels	5
Outage and Maintenance Management	5
Financial Recompense Model for not Meeting Service Levels	5
4. Provision of the service	5
Customer Responsibilities	5
Technical Requirements and Client-Side Requirements	6
Outcomes/Deliverables	6
Development life cycle of the solution	6
After-sales Account Management	6
Termination Process	6
5. Our experience	6
Case Studies	6
Clients	6
Contact Details	6

1. Introduction

Company Overview

Anthropic Ireland is a subsidiary of Anthropic PBC, an AI safety and research company founded in 2021 and headquartered in San Francisco, USA. Anthropic PBC is a Public Benefit Corporation dedicated to the responsible development and maintenance of advanced AI for the long-term benefit of humanity. Founded by former members of OpenAI, including siblings Dario Amodei (CEO) and Daniela Amodei (President), Anthropic has rapidly become a leading provider of frontier AI systems.

Anthropic has developed Claude, a family of foundation models / large language models (LLMs) designed to be helpful, harmless, and honest. Our approach, known as Constitutional AI, ensures that Claude operates according to a set of principles that prioritise safety and beneficial outcomes. This focus on robustness, safety and reliability, combined with Claude's particular strengths in fields such as coding, complex reasoning, agentic tasks and tool use, make Claude ideally suited to governmental and enterprise AI requirements.

Anthropic collaborates with civil society, government, academia, nonprofits and industry to promote AI safety industry-wide, and we are committed to transparency in our research and deployment practices.

Value Proposition

Public sector organisations face increasing pressure to deliver more efficient, accessible, and responsive services whilst managing constrained budgets and resources. Traditional approaches to a broad range of government tasks – from document processing, citizen engagement and policy analysis to information synthesis, software development and process verification – are often time-consuming, inconsistent, and difficult to scale.

Claude addresses these challenges by providing a secure, reliable AI virtual collaborator capable of processing vast amounts of information, generating high-quality content, analysing complex documents, and supporting decision-making processes. We identify three core categories of problems that Claude can help to solve:

- a. **Current human-derived process chokepoints:** For example, human processing of benefits applications, visa renewals, FOI requests, or planning permissions.
- b. **Legacy systems built on outdated technologies:** Much of government's current infrastructure relies on ageing codebases in deprecated languages. Claude can assist with code modernisation – analysing and refactoring legacy code into modern, maintainable languages while preserving business logic.
- c. **Areas where existing processes are acknowledged to be underperforming:** For example, complaint handling, procurement processes, or appointment booking systems.

Our platform enables government organisations to transform productivity, improve citizen service delivery, and unlock insights from existing data assets whilst maintaining the highest standards of security, privacy, and ethical AI use. With Claude, government can

increase productivity and performance without a corresponding increase in headcount amongst those delivering services. Claude can allow human staff to focus their efforts where they will have most impact, allowing human staff to move up the value chain – for instance to engage with complex cases or where human-to-human interaction is most needed – whilst Claude undertakes work that would otherwise be delegated to humans or, if needed, left undone.

As a case study, the European Parliament has deployed Claude to make over 2.1 million archived documents accessible to researchers, policymakers, and the public, reducing document search and analysis time by 80%. Similar impact can be achieved across UK government services.

What the Service Provides

Claude is a state-of-the-art AI assistant available as a cloud-hosted Software as a Service (SaaS) platform. The service provides access to Anthropic's family of large language models through multiple interfaces including a web-based chat interface (claude.ai), API access for integration into existing systems, and enterprise deployment options via our 'Claude Enterprise' suite.

Key capabilities include:

- Natural language understanding and generation for document drafting, summarisation, and analysis;
- Complex reasoning and problem-solving to support policy analysis and decision-making;
- Multilingual support across all official EU languages and many others, including common non-English languages within the UK;
- Code generation and technical assistance for software development and data analysis;
- Document processing including PDF analysis, data extraction, and report generation;
- Integration capabilities via APIs for embedding AI capabilities into existing workflows;
- Enterprise features including Single Sign-On (SSO), audit logging, and administrative controls.

Claude is typically available via a range of model options, including: small lightweight models with faster inference and lower cost (currently classed as Haiku); flexible and general-purpose models with broad capabilities and a balance of speed and performance (currently classed as Sonnet); and high-performance models tailored to complex use cases but where inference requirements can be more flexible (currently classed as Opus). We continue to develop new models and expand the range of models available.

Social Value

Anthropic is committed to delivering social value aligned with the UK Government's PPN002 Social Value Model. Our initiatives address the following priority areas:

Kickstart Economic Growth:

Anthropic operates a startup support program to foster innovation and economic growth globally, including in the UK. This includes API credits, priority rate limits and exclusive access to technical experts. We have dispensed millions of dollars of credits across thousands of startups, supporting founders building the next generation of AI-first applications and developing AI skills.

In addition, Claude enables UK businesses and public sector organisations to increase productivity and efficiency. We are providing a world-leading service, free of charge via claude.ai, to support economic growth. We are the leading provider of frontier AI to enterprise, with our free tier of claude.ai already accelerating UK growth. Our core products are highly effective at automating routine cognitive tasks and enhancing knowledge work, delivering productivity gains that our estimates indicate equate to an additional 1.8% labour productivity annually in developed economies over the next decade. Our API platform enables UK-based developers and enterprises to build innovative AI-powered applications, contributing to the UK's digital economy.

Through our Anthropic Economic Index, researched and published publicly for society's benefit, we inform societies and governments about the economic potential, implications and risks of AI. This allows government to target policies to accelerate economic growth and address the impacts of AI.

Our Beneficial Deployments team works, both directly and indirectly, with educators around the world, from the UK to Rwanda and Iceland, to apply AI in education, supporting innovation, skills development and improved economic productivity from the next generation.

Make Britain a Clean Energy Superpower:

Anthropic is committed to sustainable operations and has partnered with cloud providers that utilise renewable energy sources. We aim to offset all of our operational carbon emissions each year.

In addition, Claude's scientific and research capabilities are ideal for the efficient analysis of data including climate data. This can dramatically increase efficiency, supporting sustainability and clean growth. For instance, we are working with AES, a global energy company, to use Claude in audits, reducing audit costs by 99% and time for an individual audit from 14 days to 1 hour.

Break Down Barriers to Opportunity:

Our AI assistant supports accessibility by providing tools that can transform ways for people to engage with the world and with information. Claude is accessible via text and voice modes, and our accessibility testing seeks to improve the ways in which people with disabilities can utilise AI.

Claude's abilities to synthesise, simplify, augment or expand on information are key in breaking down opportunity. For instance, academic research has shown that educational tools based on Claude are powerful in improving inclusivity in the classroom, allowing children with special educational needs to benefit. Swindon Borough Council has

implemented 'Simply Readable', a Claude-powered tool that creates accessible versions of documents for people with learning disabilities.

AI skills will be key to future success, and we seek to break down barriers to developing these skills. Our dedicated Claude Tutorials resources help people globally, from all backgrounds, access top-tier learning materials to derive the best performance from frontier AI.

Our Claude for Nonprofits programme allows NGOs working to support development, opportunity, equality and more, to access discounts of up to 75% on Anthropic products.

Our staff participate in volunteering days and social give-back events. For instance, our public sector team recently volunteered at a homeless shelter to support meal provision for c.250 homeless people.

Build an NHS Fit for the Future:

We are committed to holistic health and wellness, both for our staff and society in general. Our staff benefit from dedicated wellness stipends, health insurance and additional therapy services. This helps to ensure that they and their families are healthy and flourishing.

The success of the NHS of the future will be underpinned by modern technologies, including frontier AI. Claude can support healthcare organisations with administrative efficiency, research analysis, and patient communication. The District of Columbia Department of Health has successfully deployed a Claude-powered bilingual chatbot to improve health programme information delivery and citizen engagement. We are actively working with the NHS on opportunities to use Claude for performance improvements and efficiency. We monitor and report on social value outcomes through quarterly reviews with designated account managers.

Overview of the G-Cloud Service

Claude is Anthropic's AI assistant, offering advanced capabilities in reasoning, coding, analysis, and natural language tasks. Built on Anthropic's Constitutional AI methodology, Claude combines high performance with robust safety characteristics specifically designed for high-stakes enterprise and government deployment - including transparency in reasoning, refusal of harmful requests, and alignment with organisational policies.

The platform is available through four deployment options: Claude API for programmatic integration, Claude.ai for browser-based access, Claude Enterprise for organisation-wide deployment with advanced security controls, and Claude Code for software development workflows.

Service Offerings

Claude is a suite of frontier AI models and associated capabilities. Built atop multimodal large language models, Claude is able to function as a virtual collaborator, with an AI model powering autonomous or human-directed actions ranging from: question-answer via

chatbot; code completion and analysis; data synthesis and analysis; translation; complex agentic use cases, such as those involving computer use or external tooling; and more. Claude is available through multiple deployment options, each designed for different use cases and integration requirements:

Claude API (Claude Developer Platform)

Programmatic access to Claude models for integration into customer applications, workflows, and products. The API enables organisations to build AI-powered solutions tailored to their specific requirements. Access can be obtained through purchase of usage credits for pay-as-you-go or tokens per minute priority tier for billing per month, which also provides our highest quality of service for availability in periods of peak demand.

- **Messages API:** Core conversational interface supporting text, images, and document inputs with streaming responses.
- **Batch Processing:** High-throughput processing for large-scale document analysis, data extraction, and bulk operations at reduced cost.
- **Tool Use / Function Calling:** Enable Claude to interact with external systems, databases, and APIs within customer-defined parameters.
- **Extended Context Windows:** Process long documents and expanded datasets, enabling analysis of lengthy reports, contracts, and technical documentation.
- **Model Selection:** Choose from a range of models based on task requirements, with the current model family including Opus (complex tasks requiring advanced reasoning/performance), Sonnet (balance of capability with speed and cost), or Haiku (fastest model for speed-sensitive use cases).
- **Web Search and Code Execution:** Using the power of Claude to search the internet and to execute code for coding, system automation and agentic tooling.
- **Skills:** Reusable, modular instructions that store information to enable specialized workflows.
- **Memory:** Persistent storage of preferences, terminology, procedures, and experiences to tailor Claude to your requirements.

Claude Apps

Browser, desktop, and mobile-based interface providing direct access to Claude for individual users, teams and large organizations without technical integration. Accessible via Claude Enterprise, Team, Max, Pro and Free tiers.

- **Conversational Interface:** Natural language interaction with support for document uploads, image analysis, and file generation.
- **Projects:** Organise work into persistent workspaces with custom instructions and reference materials.

- **Artifacts:** Claude can generate and render code, documents, visualisations, and interactive applications within the interface.
- **MCP Connectors:** Connect Claude to enterprise tools including Google Workspace, Slack, Microsoft 365, and other business systems for context-aware assistance.
- **Web Search and Research:** Access current information and conduct multi-source research with citations.

Pro and Max plans are enhanced services for individual users seeking improved functionality beyond our Free tier:

- **Pro:** Increased usage, Claude Code access, research, MCP connectors, memory across conversations, enhanced flexibility in model selection, and more. Includes access to Skills, which are reusable, modular instructions that store information to enable specialized workflows.
- **Max:** Significantly increased usage (either 5x or 20x), priority access at periods of peak demand, and early access to advanced features, plus features of Pro.

Claude Enterprise: Enterprise-grade deployment with advanced security, administration, and collaboration features for organisation-wide rollout.

- **Single Sign-On (SSO):** Integration with existing identity providers for seamless authentication.
- **Domain-Based Team Management:** Automatic user provisioning based on email domain.
- **Admin Console:** Centralised management of users, permissions, usage policies, and spend controls.
- **Usage Analytics:** Detailed reporting on adoption, usage patterns, and cost allocation across teams.
- **Data Governance:** Audit logs, and compliance controls.
- **Extended Context and More Usage:** Large, 500K context window and more usage for enterprise workloads.
- **Compliance API:** Programmatic access, management and audit of Claude usage data to support compliance functions.
- **Claude Code access:** Use of our leading Claude Code product, including the agentic harness for automated and augmented coding.
- **Memory:** Claude remembers your projects, preferences and patterns to maximize productivity.

Claude Team: Collaborative AI platform designed for small-to-medium teams with administration tools, identity management, and shared resources.

- **Single Sign-On (SSO):** Integration with existing identity providers for seamless authentication.

- **Central Administration:** Unified billing management and user access controls.
- **Spend Controls:** Configurable spending caps at organisation and individual user levels.
- **Projects:** Shared workspaces for team collaboration with document uploads.
- **Standard and Premium Seats:** Flexible seat types providing enhanced benefits, including increased usage, for premium users. More details available on request.
- **200K Context Window:** Support for long documents and complex multi-turn conversations.
- **Claude Code access:** Use of our leading Claude Code product, including the agentic harness for automated and augmented coding.

Claude Code Terminal-based and browser-based coding agent for software development teams, enabling AI-assisted development workflows directly within existing environments. Available via Enterprise, Team, Max, and Pro plans.

- **Codebase Understanding:** Claude analyses and understands entire repositories, not just individual files.
- **Agentic Development:** Claude can autonomously execute multi-step development tasks including writing code, running tests, fixing bugs, and managing git workflows.
- **IDE Integration:** Works alongside existing development environments via terminal.
- **Security Controls:** Configurable permissions for file system access, command execution, and network operations.

Key Capabilities Across All Offerings

Reasoning and Analysis Claude excels at complex reasoning tasks including document analysis, research synthesis, strategic planning, and decision support. The model can break down complex problems, consider multiple perspectives, and provide well-structured recommendations with clear reasoning chains.

Writing and Communication Professional content generation spanning policy documents, briefing papers, correspondence, reports, and public communications. Claude adapts tone and style to audience requirements while maintaining accuracy and appropriate formality for government contexts.

Coding and Technical Tasks Full-stack software development support including code generation, review, debugging, documentation, and modernisation of legacy systems. Supports all major programming languages and frameworks.

Data Processing and Extraction Structured data extraction from unstructured documents, spreadsheets, PDFs, and images. Suitable for processing Freedom of Information requests, digitising records, and automating data entry workflows.

Multilingual Support Proficient across major world languages for translation, localisation, and multilingual citizen engagement.

System Advantages

Cloud-Native Architecture: Claude is delivered as a fully hosted cloud service, eliminating the need for on-premises infrastructure or software installation. The service is accessible via any modern web browser or through API.

Multi-Cloud Availability: Claude is available via AWS / Amazon Bedrock, Google Cloud Platform (GCP) Vertex and Microsoft Azure / Microsoft Foundry.

Multi-Tenancy with Data Isolation: Enterprise deployments benefit from logical data isolation ensuring that customer data is segregated and protected. Commercial and Enterprise customers benefit from data not being used for model training.

Elastic Scalability: The platform automatically scales to accommodate varying workloads, from individual users to enterprise-wide deployments.

Continuous Updates: New features and model improvements are delivered automatically without requiring customer action or downtime.

Public Sector Experience: Claude is already trusted by government organisations including the European Parliament (document archive access), UK Department for Science, Innovation and Technology (ongoing collaborations, including on AI safety), District of Columbia Department of Health (citizen engagement chatbot), and Swindon Borough Council (accessibility tools).

Associated Services

The following services are available in addition to core Claude AI access:

API Access: Programmatic access to Claude models for integration into customer applications and workflows. Includes support for Messages API, embeddings, and batch processing. Available via Tokens per Minute priority tier plan or consumption-based credits, as per Pricing Document.

Claude Enterprise / Claude Team: Depending on service selection, advanced enterprise features include SSO integration, domain-based team management, extended context windows, and enhanced administrative controls. Pricing details are available in our Pricing Document.

Priority Tier (API): Committed capacity with target 99.5% uptime and prioritised access during peak demand periods. Priced by model as per Pricing Document.

Training and Enablement: Library of online training, supplemented with user training programmes and technical workshops depending on nature of engagement, use case and customer. Provided at no additional cost, by agreement only per engagement - please discuss with us when seeking to procure.

Pro: The entry-level paid individual subscription tier on claude.ai that provides higher usage limits than free. Priced per seat as per our Pricing Document.

Max 5x: A premium individual tier offering higher usage capacity than Pro for power users with heavier workloads on claude.ai - like Pro, this is a personal subscription, priced per seat as per our Pricing Document.

Max 20x: The highest individual tier currently offered, providing the most usage capacity for demanding individual users on claude.ai - like Pro, this is a personal subscription, priced per seat as per our Pricing Document.

2. Data Protection

Information Assurance

Anthropic maintains a comprehensive information security programme with the following certifications and attestations:

- **ISO 27001:2022** – Information Security Management System certification;
- **ISO/IEC 42001:2023** – AI Management Systems certification (demonstrating responsible AI governance);
- **SOC 2 Type II** – Independent attestation covering Security, Availability, Confidentiality, and Privacy trust services criteria;
- **HIPAA Configurable** – Business Associate Agreements available for healthcare data (US accreditation but demonstrating the significant care with which Anthropic handles sensitive data);
- **GDPR Compliance** – Data Processing Agreements available to support EU/UK data protection requirements.

Our security controls include: role-based access control with just-in-time privileged access; multi-factor authentication for all production system access; quarterly access reviews; encryption of data in transit (TLS 1.2+) and at rest (AES-256); formal change management with security review; comprehensive audit logging; and 24/7 security monitoring.

Copies of our compliance documentation are available through our Trust Portal at trust.anthropic.com upon request.

Data Back-Up and Restoration

Anthropic implements comprehensive data protection measures to ensure customer data recovery and operational continuity across all service offerings including first party/IP API, Claude for Enterprise, and third-party API integrations via AWS Bedrock and Google Vertex AI. Customer data is encrypted at rest using AES-256 GCM encryption and automatically distributed across multiple geographic locations and data centers hosted by our cloud service providers (Google Cloud Platform and Amazon Web Services) to ensure recoverability in the event of data center outages or natural disasters. Automated backups are performed continuously with multiple redundancy layers in place.

Commercial customers can export their conversation data and have control over data retention settings within their account. Specific backup and recovery requirements can be discussed prior to an order being placed and documented in the Order Form as agreed between both parties.

Business continuity statement/plan

Anthropic maintains a Business Continuity Plan (BCP) covering production services and supporting infrastructure for the Claude platform, with documented Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) established for critical dependencies to ensure minimal service interruption. The Business Continuity Plan includes automated failover to redundant systems, incident response procedures, communication protocols, and regular testing through tabletop exercises and disaster recovery simulations. The plan is reviewed and updated annually, with testing to validate recovery capabilities and identify improvement opportunities. For Enterprise customers requiring specific recovery metrics or customized continuity arrangements, these requirements can be documented in the Order Form and reflected in applicable service level agreements. A more detailed plan can be provided to buyer on request.

Privacy by Design

Anthropic has implemented Privacy by Design principles throughout the development and operation of Claude:

- **Data Minimisation:** We collect only the data necessary to provide the service. By default, Anthropic does not use commercial account data to train generative models.
- **Purpose Limitation:** Customer data is used solely to provide the requested service and improve service quality (with appropriate consent).
- **Retention Controls:** Zero data retention agreements can be available for enterprise customers requiring enhanced privacy controls.
- **Transparency:** Our [Privacy Policy](#) clearly explains data handling practices, user rights, and how to exercise them. The policy outlines categories of personal data collected (identity and contact data, payment information, inputs and outputs, technical information), the purposes for which data is used, and how data is disclosed to affiliates, service providers, and business partners. Users have rights including access and data portability, deletion, correction, objection, restriction, and withdrawal of consent, which can be exercised by contacting privacy@anthropic.com. Users can manage their data preferences through Privacy Settings, including opting out of data use for model training. Anthropic does not sell personal data, and implements appropriate technical and organisational security measures to protect personal data from loss, misuse, and unauthorised access.
- **User Controls:** Users can delete conversations, export their data, and manage privacy settings.
- **GDPR Compliance:** Data Processing Agreements are available. We support data subject rights including access, rectification, erasure, and portability.

3. Using the service

Ordering and Invoicing

To order Claude services under G-Cloud 15:

- **Initial Contact:** Email gcloud@anthropic.com or, if preferred, pubsec-uk@anthropic.com with your requirements. For government, we are content to discuss more exploratory use cases but concrete opportunities for Claude will usually include details of estimated user numbers, intended use cases, and any specific security or compliance requirements.
- **Requirements Discussion:** An Anthropic account representative from our experienced and focused public sector team will arrange a call to understand your needs and recommend the appropriate service tier and configuration. If appropriate, we may be able to facilitate some light-touch initial testing through our Applied AI team to define your requirements further.
- **Order Form:** We will assist you in completing the G-Cloud Call-off Contract Order Form, documenting agreed service levels, pricing, and any specific terms.
- **Invoicing:** Invoices are usually issued monthly in arrears (for usage-based pricing) or annually in advance (for subscription tiers). We can accept payment via BACS, purchase order, and Government Procurement Card where applicable.

Availability of Trial Service

Anthropic offers several options for evaluating Claude before committing to a full deployment:

- **Free Tier:** Limited access to Claude via claude.ai, allowing basic evaluation of capabilities.
- **API Credits:** Complimentary API credits may be provided to qualified public sector organisations for proof-of-concept testing.
- **Pilot Programme:** Trial periods with enterprise features can be arranged for significant deployments, subject to agreement.
- **Applied AI support:** Our expert applied AI engineers can support evaluation, testing and optimisation as part of a trial service, depending on nature of engagement, customer and use case.

Please refer to our Pricing Document for applicable pricing around our services. Contact gcloud@anthropic.com to discuss options for your organisation.

On-Boarding, Off-Boarding, Service Migration and Scope

On-Boarding Process:

- Kick-off meeting with your team and dedicated Account Executive and/or Customer Success Manager
- Technical configuration including, if applicable, SSO setup, API key provisioning, and workspace configuration
- User provisioning and access management setup

- Documentation and knowledge library provided including quick-start guides, API documentation, and best practice materials
- Training sessions for administrators and end users
- 90-day enhanced support period with dedicated Account Executive / Customer Success Manager assistance, potentially including Applied AI / Forward Deployed Engineer Support

Off-Boarding Process:

- Written notice of termination as per contract terms
- Transition planning meeting to agree timeline and data export requirements
- Data export in standard formats (JSON, CSV)
- User access revocation on agreed date
- Confirmation of data deletion per retention policies and applicable regulations

Training

Anthropic provides comprehensive training resources:

- **Self-Service Resources:** Comprehensive documentation at docs.anthropic.com including tutorials, best practices, and API reference
- **Prompt Engineering Guides:** Detailed guidance on crafting effective prompts for optimal results
- **Video Tutorials:** On-demand training content covering common use cases and features
- **Live Training:** Instructor-led sessions for enterprise customers (included in enterprise tiers or available for purchase)
- **Administrator Training:** Technical training for IT administrators on integration, security, and management features

Implementation Plan

Implementation timelines typically range from immediate access (for standard configurations) to 2-4 weeks (for enterprise deployments with SSO integration and custom configuration), but where our engineers are collaborating with the customer on a complex agentic system based on Claude, development and implementation timelines can range up to 6-12 months.

A detailed implementation plan tailored to your organisation's requirements can be provided to the buyer on request.

Service Management

We follow a wide variety of internally & externally developed standards for ensuring a high service bar across our platform. Our practices align with industry-recognised frameworks and are reinforced by our compliance portfolio. We have procedures for incident management, change management, and customer support.

- **Industry Standards Followed:** Anthropic does not currently hold ITSM-specific certifications. However, our operational practices align with ITSM principles and

other widely known operating standards and are supported by our broader compliance portfolio, which includes ISO 27001, SOC 2 Type II, ISO 42001 (AI Management System), CSA STAR, and Cyber Essentials Certification. More information about our compliance artifacts and certifications can be found in our [Trust Center](#). The Trust Center provides access to compliance documentation including SOC 2 Type II attestation reports (available under NDA for Enterprise customers), ISO 27001:2022 and ISO/IEC 42001:2023 certificates, and a publicly accessible SOC 3 summary report. Anthropic implements role-based access control, multi-factor authentication for production systems, quarterly access reviews, and a formal change management process with technical and security review requirements before deployment.

- **Incident Management Practices:** Anthropic operates a formal incident management system with priority-based classification (P0-P3) and structured response procedures including defined roles and post-incident retrospectives. We provide communications regarding incidents via [claude.status.com](#) (where we provide real-time updates on scope, impact, and resolution progress about ongoing incidents). For subscribed customers, we provide communication via email, and Slack. Root cause analyses are conducted for significant incidents.
- **Change Management Practices:** Governed by internally documented procedures with deployment controls, and model deprecations follow a defined procedure that can be found [here](#). Models progress through a defined lifecycle: Active, Deprecated, and Retired. Customers with active deployments are notified of upcoming retirements, with a minimum of 60 days' notice before model retirement for publicly released models. Customers can access an audit of their API usage via the Console to identify any instances where applications are still using deprecated models. Anthropic has also committed to preserving the weights of all publicly released models for the lifetime of the company to ensure optionality for future availability. The model deprecation process is designed to ensure ample notification & time are provided to users before a model is fully deprecated. Support tickets are handled via email, with regular service reviews available for enterprise customers. Customers can also access a chat bot on our Help Center to ask support questions.
- **Customization Available:** Claude is provided as a standardized SaaS offering. Enterprise customers have access to configuration options, but the underlying models and platform infrastructure are shared across customers. Available customization capabilities include:
 - **Administrative configuration:** Enterprise administrators can manage workspace settings, user access, and organizational controls via the Administration API. Documentation is available at: [platform.claude.com/docs/en/build-with-claude/administration-api](#)
 - **Product configuration:** Customers can configure API parameters, model selection, and integration settings as documented at: [platform.claude.com/docs/en/home](#)

Service constraints, service levels, and outage management procedures are detailed in the relevant sections below.

Service Constraints

- **Maintenance Windows:** Anthropic's Claude services operate without scheduled maintenance windows as the platform is designed for continuous availability. Any service disruptions are unplanned incidents during which real-time service status and incident notifications are available at status.claude.com.
- **Customisation:** Claude is provided as a standardised SaaS offering. Enterprise customers have access to configuration options but the underlying models and platform are shared. Available customization capabilities include:
 - **Administrative configuration:** Enterprise administrators can manage workspace settings, user access, and organizational controls via the Administration API. Documentation is available at: platform.claude.com/docs/en/build-with-claude/administration-api
 - **Product configuration:** Customers can configure API parameters, model selection, and integration settings as documented at: platform.claude.com/docs/en/home
- **Rate Limits:** Rate/usage limits apply based on service tier. Priority Tier customers receive guaranteed capacity. Those buying additional services such as Pro/Max seats benefit from additional capacity.
- **Context window/maximum output:** The maximum input and/or output that Claude can accommodate in any single interaction varies by model.
- **Feature Deprecation:** Minimum 90 days notice can be provided for any feature deprecation or significant changes excluding beta features.
- **Regional availability:** Services vary by region depending on data residency requirements.
- **Browser requirements:** Although Claude is flexible across browsers, a modern browser is required for web interface (e.g. Chrome, Firefox, Safari, Edge)

Service Levels

- **Availability Target:** With an SLA exhibit, 99.5% is the default target uptime for our 1P API.
- **Support Hours:** Standard support available during UK business hours (09:00-17:00 GMT). Enhanced support available 24/7 for enterprise customers.
- **Response Times:** Anthropic tiers incidents based on business impact and has calibrated target response times. Further detail is available on request. Contractual SLAs are available through our paid Enterprise Support tier. Human and AI support is available 24/7.

Specific service levels can be discussed and agreed prior to an order being placed and documented in the Order Form.

Outage and Maintenance Management

- **Status Page:** Real-time service status available at status.anthropic.com with subscription options for updates
- **Planned Maintenance:** Advance notification via email and status page

- **Incident Communication:** Updates provided via status page, email, and Slack (for subscribed customers)
- **Post-Incident Reviews:** Root cause analysis conducted for significant incidents. We notify status publicly on our website and seek to publish reports and/or respond to customer requests for information where feasible for significant incidents.
- **Redundancy:** Multi-region architecture with automatic failover to minimise service disruption
- **Rate limiting and Quality of Service (QoS) tiers:** We use these to ensure fair resource allocation. Each organisation has defined rate limits (requests per minute, tokens per minute) that prevent any single customer from consuming excessive capacity. Traffic is isolated by organisation with separate quotas. Enterprise customers can negotiate custom rate limits.

Financial Recompense Model for not Meeting Service Levels

Service Credits are available for customers with a 1P API Availability SLA Exhibit in contract. This does not cover Claude for Enterprise or Customer's access via third-party environments (AWS Bedrock, GCP Vertex).

Anthropic will apply Service Credits against future payments otherwise due from Customer for use of the applicable Covered Service. Any unused portion of Service Credits remaining as of the expiry or termination of the Terms will be forfeited without compensation.

<u>Monthly Availability</u> <u>Percentage</u>	<u>Service Credit</u>
< 99.5% to 99.0%	10%
< 99.0% to 95.0%	15%
< 95.0%	25%

In order to receive eligible Service Credits, Customer must submit a request to notices@anthropic.com (or another email Anthropic may provide) within 30 days of the end of the applicable calendar month. The request must include appropriate logs or other materials that document the errors and corroborate a claimed breach of the relevant Service Level Objective (and any confidential or sensitive information in these logs should be removed or replaced with asterisks). Anthropic will verify Customer's report through any available system logs and records.

The maximum aggregate Service Credits issued by Anthropic to Customer: (i) for a single billing month, will not exceed 25% of the amount due from Customer for the applicable Covered Service, (ii) for Covered Services under an active order form, will not exceed 5% of Customer's initial term spend commitment as defined in the applicable order form, and (iii) in the absence of any applicable order form, will not exceed five percent (5%) of Customer's total spend, calculated as: (x) fees paid for the applicable Covered Service in the previous 12-month period, or (y) if Customer has used the Covered Service for less than 12 months, the annualized amount based on Customer's actual usage period.

4. Provision of the service

Customer Responsibilities

Use of Services:

The Customer is responsible for complying with all laws applicable to the use of the Services, including any applicable data privacy laws.

Customer and users may only use Services in compliance with the [Anthropic Usage Policy](#) aka Acceptable Usage Policy, the policy on countries and regions Anthropic currently supports (Supported Regions Policy) and Anthropic Service Specific Terms. The Customer will cooperate with reasonable requests for information from Anthropic to support compliance with the Usage Policy, including verification of Customer's identity and use of Services.

It is the Customer's responsibility to evaluate whether the Outputs (defined in Anthropic Terms of Service) are appropriate for Customer's use case, including where human review is appropriate, before using or sharing Outputs. The Customer must acknowledge, and must notify Users, that factual assertions in Outputs should not be relied upon without independently checking their accuracy, as they may be false, incomplete, misleading or not reflective of recent events or information.

The Customer may not and must not attempt to (a) access the services to build a competing product or service, including to train competing AI models or resell the Services except as expressly approved by Anthropic; (b) reverse engineer or duplicate the Services; or (c) support any third party's attempt at any of the conduct restricted in this sentence.

The Customer is responsible for all activity under its account and promptly notify Anthropic if it believes the account it uses to access the Services has been compromised, or is subject to denial of service or similar malicious attacks that may negatively impact the Services.

In addition, the Customer is responsible for:

- Providing accurate information during onboarding and maintaining up-to-date contact details;
- Securing keys and credentials;
- Ensuring compliance with internal governance and approvals for AI usage;
- Paying invoices within agreed terms.

Technical Requirements and Client-Side Requirements

- **Internet Connectivity:** Reliable internet connection required
- **Web Browser:** Modern browser (Chrome, Firefox, Safari, Edge - latest two versions)
- **API Integration:** HTTPS capability for API calls; JSON parsing capability

- **SSO Integration:** SAML 2.0 or OIDC-compatible identity provider (for enterprise SSO)
- **Network:** Firewall rules permitting HTTPS traffic to *.anthropic.com and *.claude.ai domains

Specific requirements can be discussed prior to an order being placed and documented in the Order Form.

Outcomes/Deliverables

Customers can expect the following outcomes from Claude deployment:

- Access to state-of-the-art AI virtual collaborator, suited, for example, for text generation, multimodal and text analysis, and reasoning tasks
- Productivity improvements in document processing, research, and content creation
- Enhanced accessibility of information and services for citizens
- Secure, auditable AI usage with enterprise governance controls
- Continuous access to model improvements and new capabilities

Development life cycle of the solution

Anthropic follows an Agile development methodology with continuous deployment:

- **Research:** Continuous AI safety and capability research
- **Development:** Iterative development with automated testing
- **Safety Testing:** Rigorous red-teaming and safety evaluations before release
- **Staged Rollout:** Progressive deployment with monitoring
- **Continuous Improvement:** Regular model updates and feature enhancements

After-sales Account Management

Anthropic has and continues to build a dedicated go-to-market function, including providing dedicated support in the UK for Public Sector customers; on-going support would typically include a dedicated & named Account Manager with deep expertise in the public sector, alongside a technical expert to ensure customers make the most out of the platform.

Customers often fall into two camps: those that are new to Claude and how it can improve their workflows; and power users who are used to using Claude Developer Platform in their personal lives to develop applications but are potentially yet to adopt it for their daily work.

Typical engagements include:

- 1) **Workshops** - Introduction to Claude/Developer Platform, prompting masterclass, Claude code masterclass, hackathons, roadmap sessions and product deep-dives;
- 2) **Art-of-the-possible** - building demos based on customer pain points;
- 3) **Evaluations** - helping customers develop evaluations for model outputs to ensure it meets their specific needs;
- 4) **FinOps** - helping customers understand their usage patterns, assessing spend and evaluating different commercial models for predictable budgeting

- 5) **Custom Agents** - in certain situations engagements could include the deployment of Anthropic engineers to build out complex workflows likely as a result of a multi-year commitment to the Anthropic platform

In addition to named Account Management support, as a paid user of Claude or the Console, customers have full access to:

- All help documentation
- Fin, our AI support bot
- Further assistance from our Product Support team

Termination Process

Either party may terminate the agreement in accordance with the G-Cloud Call-off Contract terms. Key termination provisions include:

- **Notice Period:** As specified in CCS General Terms
- **Data Export:** Assistance provided to export customer data prior to termination
- **Transition Support:** Reasonable assistance for migration to alternative services
- **Data Deletion:** Confirmation of customer data deletion following termination
- **Early Termination:** The Customer acknowledges that it has purchased the Services for the Minimum Period and any Renewal Term(s) as defined in the Order Form and in accordance with CCS General Terms.

5. Our experience

Case Studies

1. European Parliament – Archive Accessibility

- a. **Challenge:** The European Parliament needed to make over 2.1 million archived documents accessible to researchers, policymakers, educators, and the public across multiple languages.
- b. **Solution:** Leveraging Claude via Amazon Bedrock, the EU Parliament created 'Ask the EP Archives' (Archibot), an AI assistant providing instant access to documents including resolutions, positions, policies, and inter-institutional negotiations.
- c. **Outcome:** 80% reduction in document search and analysis time; multilingual access in all EU languages; enhanced transparency for citizens; enabling researchers and staff to analyse data and create comprehensive reports efficiently.

2. Swindon Borough Council – Simply Readable

- a. **Challenge:** Making complex local government documents accessible to residents with learning disabilities.
- b. **Solution:** Claude-powered tool that transforms complex documents into accessible formats with larger fonts, increased spacing, and supporting images.
- c. **Outcome:** Free public tool improving access to important information for vulnerable residents.

3. District of Columbia Department of Health

- a. **Challenge:** Improving accessibility and engagement with health programme information for DC residents.
- b. **Solution:** Claude-powered bilingual chatbot providing health programme information delivery.
- c. **Outcome:** Enhanced accessibility, improved information delivery, and increased resident engagement supporting the city's goal of becoming America's healthiest.

Clients

Anthropic's Claude is trusted by government and public sector organisations worldwide. It is the world's leading AI system for enterprise use cases. Specific public sector clients include:

- European Parliament
- UK Department for Science, Innovation and Technology (DSIT)
- Swindon Borough Council
- District of Columbia Department of Health
- US Government (available extensively across all three branches, including major partnership with Department of Defense/Department of War)
- Lawrence Livermore National Laboratory
- Icelandic Ministry for Education

Enterprise customers include: Pfizer, Intuit, Databricks, and over 130 publicly referenced organisations.

Contact Details

For G-Cloud enquiries:

Named Point of Contact: Iain Mackay, Account Executive

Email: gcloud@anthropic.com or pubsec-uk@anthropic.com

Website: www.anthropic.com

Our UK public sector team is available to discuss your requirements and arrange demonstrations of Claude's capabilities.