ANTHROP\C

Anthropic AI via AWS Bedrock: Agentic AI Virtual Collaborator - Claude

Subcontractor of Softcat

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

January 2026

1. Pricing

Company Overview

This listing is for customers to access Anthropic's AI models via AWS Bedrock.

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b. **Legacy systems built on outdated technologies:** Much of government's current infrastructure relies on ageing codebases in deprecated languages.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 previously 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. **For this service offering, access to Claude is via AWS Bedrock, allowing government users to benefit from AWS's flexibility in regional data hosting.**

This service provides access to Anthropic's family of large language models via API, delivered through AWS's managed Bedrock service with enterprise-grade security and compliance controls.

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 functionality 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 products 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 API via AWS Bedrock, including access to our state-of-the-art coding platform, Claude Code, ideally suited to software development workflows but adaptable to broader computer use agentic AI tasks.

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 on Amazon Bedrock

API-based access to Anthropic's Claude models through Amazon Web Services' fully managed foundation model service. Enables enterprises to build, scale, and deploy generative AI applications with Claude's advanced reasoning, coding, and agentic capabilities within AWS infrastructure.

Model Access

- **Claude Model Family**: Access to the Claude model lineup via unified API.
- **Extended Context Windows**: Support for up to 200,000 tokens standard, with select models supporting up to 1 million tokens for complex multi-turn conversations and extensive document processing.
- **Core Capabilities**: Advanced reasoning, vision analysis, code generation, multilingual processing, and tool use across all model tiers.
- **Agentic Features**: Tool search, tool use examples, and configurable effort parameters to balance performance with latency and cost.

Security & Identity

- **Single Sign-On (SSO)**: Integration via AWS IAM Identity Centre supporting SAML 2.0 and OIDC protocols, with compatibility for enterprise identity providers including Okta, Azure AD, and Ping Identity.
- **IAM Integration**: Granular role-based access controls through AWS IAM policies, eliminating the need to distribute API keys.
- **Encryption**: TLS 1.2+ for all data in transit; AES-256 encryption for data at rest.
- **Private Network Deployment**: Zero data egress options for regulated workloads through VPC endpoints and private connectivity.
- **Regional Compliance**: Available across commercial AWS regions and AWS GovCloud (US) for public sector requirements.

Cost Management & Usage Analytics

- **Application Inference Profiles**: Custom cost allocation tags enabling usage tracking by cost centre, business unit, application, or tenant.

- **AWS Cost Explorer Integration**: Tagged usage flows directly into Billing and Cost Management for detailed expense reporting across organisational hierarchies.
- **Budget Controls**: Configurable spending thresholds via AWS Budgets with tag-based alerts and anomaly detection.
- **Cost Optimisation Features**: Prompt caching, Model Distillation (up to 75% cost reduction), Intelligent Prompt Routing (up to 30% savings), and batch processing (up to 50% lower overall costs, as per our Pricing Document).

Governance & Audit

- **CloudTrail Logging**: Complete audit trail of model invocations including user/role, timestamp, and parameters.
- **CloudWatch Integration**: Usage metrics collection and monitoring with customisable dashboards and alarms.
- **SIEM Compatibility**: Audit logs can feed into existing security information and event management systems.
- **Managed Permissions**: Organisation-wide policy enforcement through managed configuration that individual users cannot override.

Developer Experience

- **Unified API**: Single endpoint for all Claude models with consistent request/response patterns.
- **Claude Code Access**: Support for Anthropic's agentic coding harness for automated and augmented software development workflows.
- **Amazon Bedrock AgentCore**: Production-grade infrastructure for autonomous agents including Runtime, Memory, Identity, Gateway, Code Interpreter, Browser Tool, and Observability services—enabling agents that can run for up to eight hours with automatic scaling.
- **Memory Tools**: Store and consult information outside the context window to boost accuracy and performance on long-running tasks.
- **Context Editing**: Automatic clearing of stale information from past tool calls to maximise effective context utilisation.

Enterprise Deployment

- **AWS Marketplace Availability**: Streamlined procurement and consolidated billing through existing AWS agreements.
- **No Development Required Option**: Claude for Enterprise available via AWS Marketplace with native integrations and UI access for non-technical team members.
- **Just-in-Time Provisioning**: Automatic user provisioning tied to identity provider authentication.

- **Multi-Tenant Architecture**: Application inference profiles enable cost isolation and tracking across multiple internal teams or external customers.

Claude Code Terminal-based and IDE-integrated coding agent for software development teams, enabling AI-assisted development workflows directly within existing environments. Claude Code is an Anthropic product that can be used in a configuration pattern to call Claude via API through Amazon Bedrock, using AWS IAM authentication and consumption-based pricing.

- **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 and IDE extensions (VS Code, JetBrains).
- **Security Controls**: Configurable permissions for file system access, command execution, and network operations. Organisation-wide policy enforcement through managed configuration.
- **AWS Authentication**: Uses standard AWS SDK credential chain—IAM roles, CLI configuration, or environment variables—eliminating separate API key management.
- **Enterprise Identity Integration**: Connects with enterprise identity providers (Okta, Azure AD, Auth0, Cognito) for centralised access control and audit trails.
- **Prompt Caching**: Cost and latency optimisation for large codebase processing through Bedrock's prompt caching capability.

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 via AWS, eliminating the need for on-premises infrastructure or software installation.

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 access via Bedrock:

Reserved Tier (aka Reserved Service): Committed capacity on AWS Bedrock with target 99.5% uptime for model responses and prioritised access during peak demand periods. Enables reservation of specific tokens-per-minute capacity with flexible allocation between input and output tokens. Available in 1-month or 3-month commitment terms. Automatic overflow to on-demand Standard tier when reserved capacity is exceeded ensures uninterrupted operations. Priced in tokens-per-minute as per Pricing Document.

Provisioned Throughput: Dedicated model capacity for high-volume, predictable workloads requiring guaranteed throughput. Charged hourly with 1-month or 6-month commitment options. Suitable for enterprise-grade use cases with stable, high-volume inference requirements.

Training and Enablement: Library of AWS and Anthropic 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.

2. Data Protection

Information Assurance

a. Infrastructure Layer

Our service via AWS Bedrock benefits from AWS's leading security and data protection capabilities, including data encryption, network isolation, access controls, audit logging — all governed by AWS certifications and Bedrock's security architecture. Details of Bedrock's Security and Compliance offering is available at https://aws.amazon.com/bedrock/security-compliance/

b. Model Layer

Anthropic maintains a comprehensive information security programme at the model layer 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

Claude on Amazon Bedrock is a stateless inference service – prompts and responses are processed in real-time and are not stored by the service after processing. According to AWS's published documentation, customer data processed via Amazon Bedrock is encrypted at rest and in transit, is not used to train foundation models, and is not shared with model providers including Anthropic.

As a stateless API service, there is no persistent customer data within Bedrock that requires backup. Customers are responsible for implementing backup and recovery procedures for:

- Application code and configurations that integrate with Bedrock
- Any data stored in customer-managed systems before or after Bedrock processing
- AWS account configurations, IAM policies, and service settings

AWS provides enterprise-grade backup capabilities (AWS Backup) for customer-built applications and data stores. Specific backup and recovery requirements should be discussed with AWS directly as part of the customer's AWS engagement.

For details on Amazon Bedrock's data handling practices, refer to: https://aws.amazon.com/bedrock/security-compliance/

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 buyers 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:** In collaboration with our reseller, we can assist you in completing the G-Cloud Call-off Contract Order Form, documenting agreed service levels, pricing, and any specific terms.

Availability of Trial Service

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

- **API Credits:** Complimentary API credits may be provided to qualified public sector organisations for proof-of-concept testing.
- **Pilot Programme:** Trial periods can potentially 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

Claude on Amazon Bedrock benefits from AWS's enterprise-grade service management practices combined with Anthropic's AI-specific operational standards. Service management follows AWS's well-established procedures for incident management, change management, and customer support.

- **Industry Standards Followed:** Amazon Bedrock operates within AWS's comprehensive compliance framework, which includes ISO 27001, SOC 2 Type II, and numerous other certifications. Anthropic maintains its own compliance portfolio including ISO 27001, SOC 2 Type II, ISO 42001 (AI Management System), CSA STAR, and Cyber Essentials Certification. More information about Anthropic's compliance artefacts and certifications can be found in our Trust Centre (https://trust.anthropic.com/). AWS compliance information is available at https://aws.amazon.com/compliance/. Combined, this provides customers with dual assurance across both the infrastructure layer (AWS) and the AI model layer (Anthropic).
- **Incident Management Practices:** AWS operates a formal incident management system with structured response procedures. Real-time service status and incident notifications are available via AWS Health Dashboard (https://health.aws.amazon.com/) and the AWS Service Health Dashboard. Customers can configure AWS Health events to trigger notifications via Amazon EventBridge, SNS, or integrate with existing ITSM tools. Post-incident reports are published for significant service events.
- **Change Management Practices:** Model updates and new features on Amazon Bedrock follow AWS's change management processes. Model versions are documented at https://docs.aws.amazon.com/bedrock/latest/userguide/models-supported.html. Customers can use specific model version ARNs to maintain consistency and control when to adopt new versions. Customers can monitor their API usage via AWS CloudWatch and Cost Explorer to track model utilisation. Anthropic has committed to preserving the weights of all publicly released models for the lifetime of the company to ensure optionality for future availability.
- **Support:** AWS provides tiered support plans with varying response times and access to technical resources:
    - **Basic** (free): Account and billing questions only
    - **Business Support+** (from $29/month): AI-powered assistance with 30-minute response for critical cases, seamless escalation to AWS experts
    - **Enterprise** (from $5,000/month): Designated Technical Account Manager (TAM), 15-minute response for production-critical cases
    - **Unified Operations** (from $50,000/month): Dedicated team including TAM, Domain Engineer, and Billing Specialist with 5-minute response times
- Support is available via AWS Support Centre. For Anthropic-specific model questions regarding Claude's capabilities, safety, or acceptable use policies, customers can contact Anthropic via support@anthropic.com.

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

Service Constraints

- **Maintenance Windows:** Anthropic's Claude services are designed to operate continuously but AWS may, at times, conduct scheduled maintenance. Details of AWS Bedrock service availability is publicly available at https://status.aws.amazon.com/
- **Configuration:** Claude models are accessed through Amazon Bedrock as a managed service. Configuration and administration are handled through AWS-native tooling:
    - **Access and governance:** Model access is managed via AWS IAM policies and Service Control Policies (SCPs). AWS Organizations provides centralised account and access management. Documentation is available at docs.aws.amazon.com/bedrock
    - **Model selection and API configuration:** Customers select models from the Bedrock model catalogue and configure inference parameters (temperature, max tokens, etc.) through the AWS SDK or Console.
    - **Service tiers and capacity:** Bedrock offers Standard, Priority, and Reserved service tiers. The Reserved tier provides dedicated compute capacity with a target of 99.5% model response uptime, with overflow to the Standard tier.
    - **Model lifecycle:** Once launched on Bedrock, models transition through Active → Legacy (usually 6 months before end of life, EOL) → Extended Access (usually 3 months) → EOL. AWS notifies active users at each stage. Documentation is available at: docs.aws.amazon.com/bedrock/latest/userguide/model-lifecycle.html
    - **Regional availability:** Model availability varies by AWS region. Customers should verify availability in their chosen region(s) for data residency and operational requirements.

Service Levels

- **Availability:** Amazon Bedrock carries a 99.9% monthly uptime SLA, measured at the regional level. Scheduled maintenance is not included in considering this SLA.
- **Support:** Support is provided through AWS Support plans. Business Support includes response times of under 1 hour for production system outages. Enterprise Support extends this to under 15 minutes for business-critical outages, with 24/7 access to AWS technical experts.
- **Incident monitoring:** Service status and incident notifications are available via the AWS Service Health Dashboard at status.aws.amazon.com.

Outage and Maintenance Management

- **Status Page:** Claude via AWS Bedrock is a managed service with real-time service status available at status.aws.amazon.com.

- **Incident communication:** AWS communicates service incidents and degradation via the Service Health Dashboard and AWS Health Notifications. Customers can configure alarms to monitor Bedrock-specific metrics and trigger notifications.
- **Post-incident reporting:** AWS publishes post-event summaries for significant infrastructure-level incidents affecting core AWS services.
- **Cross-region capacity:** Amazon Bedrock supports cross-region inference via inference profiles, which automatically route requests to regions with available capacity when the source region is under load or quota-constrained. Geographic inference profiles (US, EU) keep traffic within defined boundaries for data residency compliance. Global inference profiles route across all commercial regions for maximum throughput but should only be used where data residency requirements permit.

Rate limiting and quotas

Amazon Bedrock enforces per-account quotas on model inference, including requests per minute (RPM), tokens per minute (TPM), and tokens per day (TPD). These quotas are region-specific and model-specific, and default values vary. On-demand inference uses a shared capacity pool; during periods of high demand, requests may be throttled even when within account quotas. Customers can request quota increases via AWS Support, though increases are subject to capacity availability and are not guaranteed. For workloads requiring dedicated, predictable throughput, Provisioned Throughput provides reserved compute capacity with defined commitment tiers.

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.

Outcomes/Deliverables

Customers can expect the following illustrative outcomes from Claude deployment via AWS:

- API access to state-of-the-art AI models, suited for implementation as a virtual collaborator in, for example, text generation, multimodal and text analysis, and reasoning tasks.
- AI-enhanced 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.

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 API/Claude Developer Platform, potentially including via Bedrock, 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, 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.

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:

- 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.