



# A.I. Data Privacy, Security and Compliance

last updated: August 7th 2025

## Executive Summary

Alive5 delivers enterprise-grade AI solutions that prioritize your data security above all else. Built on Amazon Bedrock infrastructure with Anthropic's Constitutional AI, our platform ensures your proprietary data never leaves your control, responses come exclusively from your content, and built-in safeguards maintain professional standards. This combination provides the perfect balance of cutting-edge AI capabilities with the security, compliance, and trustworthiness that business-critical applications demand.

## Enterprise AI with Amazon & Anthropic

Alive5 has strategically partnered with Amazon Web Services (AWS) Bedrock to deliver an enterprise-grade AI infrastructure that meets the most stringent security, compliance, and data protection requirements. Amazon Bedrock represents the gold standard in enterprise AI infrastructure, providing foundational models through a fully managed service designed specifically for business-critical applications.

As a trusted AWS partner, Alive5 leverages Bedrock's enterprise-focused infrastructure to ensure our clients receive AI solutions that prioritize data sovereignty, regulatory compliance, and operational security above all else. This partnership enables us to offer enterprise-grade AI capabilities while maintaining the highest standards of data protection and privacy that business leaders demand.



## Why Amazon Bedrock for AI Infrastructure

Amazon Bedrock is AWS's flagship generative AI service, specifically engineered to address the unique challenges enterprises face when implementing AI solutions. Unlike consumer-focused AI platforms, Bedrock provides enterprise customers with:

- Complete control over their data with technical guarantees preventing unauthorized access
- Robust security frameworks with multiple layers of protection
- Compliance certifications that meet requirements of highly regulated industries including healthcare, financial services, and government sectors
- Zero data retention policies ensuring customer inputs are never stored or used for training

## Data Privacy, Security and Compliance

Amazon Bedrock and Anthropic have jointly established comprehensive enterprise safeguards that address the most stringent business requirements:

### Infrastructure and Security

**AWS Infrastructure:** Leverages AWS's robust and secure cloud infrastructure with multiple layers of security and enterprise-grade compliance certifications

**Data Isolation:** Complete customer data separation with technical guarantees that data cannot be accessed by other customers or used for model training

**Zero Data Retention:** Foundation models on Bedrock do not retain, store, or learn from customer inputs or outputs after request completion

**Encryption:** AWS uses advanced encryption for data at rest and in transit, protecting customer information from unauthorized access

**Access Controls:** Robust Identity and Access Management (IAM) tools provide granular control over who can access AI models and associated data

**SOC2-Type 2 Compliance:** Comprehensive security controls audited by independent third parties

## **Compliance and Governance**

**Enterprise Compliance Certifications:** AWS maintains comprehensive compliance certifications including GDPR, HIPAA, SOC 2, FedRAMP, and industry-specific standards

**Virtual Private Cloud (VPC) Support:** Customers can deploy Bedrock within their own VPC, adding an extra layer of network isolation and security

**Comprehensive Auditing:** AWS CloudTrail provides detailed logging of all API calls, enabling thorough security analysis and compliance audits

**Model Governance:** Advanced tools for managing model versions and controlling which models can be used in production environments

**Data Residency Options:** Regional data storage options help meet specific data residency and sovereignty requirements

## **Privacy by Design**

**Privacy-First Architecture:** Anthropic's AI models are designed with privacy considerations from the ground up

**Transparency:** Both AWS and Anthropic provide comprehensive documentation on data handling practices and privacy measures

**Customer Control:** Complete customer control over data processing, storage, and access policies

**No Cross-Customer Data Sharing:** Technical architecture prevents any possibility of data sharing between customers

# Why Anthropic Claude

Our top priority has been identifying an AI model that excels in safety, security, and data privacy built for professional workloads. After extensive evaluation, we found that Anthropic's Claude Sonnet outperformed other available models in key enterprise requirements. Anthropic’s Constitutional AI approach provides the perfect balance of safety and compliance we require, while Amazon Bedrock's enterprise-grade infrastructure delivers the security foundation our users demand. This combination ensures unparalleled performance in terms of accuracy, speed, and trustworthiness for business-critical applications.

## Model Performance Overview

Claude Sonnet demonstrates superior performance across key enterprise capabilities:

Capability	Strength
Natural Language Understanding	Accurately interprets complex customer queries, even with nuanced or technical phrasing
Knowledge Breadth	Comprehensive understanding across diverse business domains and technical subjects
Contextual Awareness	Maintains conversation context for coherent, relevant multi-turn customer interactions
Professional Communication	Provides efficient, direct answers while maintaining appropriate business tone
Multi-language Support	Communicates fluently in multiple languages for global customer bases
Safety and Ethics	Constitutional AI framework ensures professional standards and appropriate responses

Claude Sonnet 3.7 is the default model powering Alive5's FAQ Bot, ensuring optimal performance and security for our users. Other models available in Alive5 include Haiku, Nova, GPT-4 and GPT-4o.

## Anthropic's Constitutional AI Framework

Anthropic's approach to AI development is rooted in "Constitutional AI," a groundbreaking framework that ensures AI systems operate within carefully defined ethical boundaries. This innovative methodology instills fundamental values and limitations directly into the AI's training process, creating models that are both highly capable and inherently aligned with professional standards.

### **Key Constitutional AI Benefits for Business:**

- **Built-in Safety:** Responses automatically filtered for inappropriate or unprofessional content
- **Ethical Alignment:** AI behavior aligned with human values and business interests from the ground up
- **Consistent Reliability:** Reduces risk of harmful or off-brand responses
- **Professional Standards:** Maintains appropriate tone and content for business communications

### **Anthropic's Enterprise-Focused Approach:**

- **B Corporation Certification:** Commitment to high standards of social and environmental performance
- **Safety-First Development:** Prioritizes AI safety research and responsible scaling
- **Long-term Perspective:** Considers long-term implications of AI development
- **Transparent Practices:** Open dialogue about AI ethics and safety within business communities

## **Implementation Benefits**

By choosing Alive5's AI solution powered by Amazon Bedrock and Anthropic, enterprises gain:

### **Data Protection**

- Your proprietary data never leaves your control
- Zero risk of data being used to train models for competitors
- Complete isolation from other customers' data and activities

### **Response Control and Safety**

- Chatbot responses limited exclusively to your approved content

- Built-in safeguards prevent inappropriate or unprofessional responses
- Constitutional AI framework ensures ethical and appropriate behavior

## Enterprise-Grade Infrastructure

- AWS's industry-leading security and compliance framework
- Multiple layers of encryption and access controls
- Comprehensive audit trails for regulatory compliance

## Trusted Partnership

- Amazon Bedrock's proven track record with Fortune 500 companies
- Anthropic's focus on safe, beneficial AI development
- Alive5's commitment to enterprise-focused AI solutions

## Summary

The AI landscape is evolving rapidly, but enterprise requirements for security, compliance, and data protection remain paramount. Alive5's partnership with Amazon Bedrock and Anthropic ensures that your organization can leverage cutting-edge AI capabilities while maintaining the highest standards of data security and regulatory compliance.

Ready to implement enterprise-grade AI with complete confidence? Reach out to us at [sales@alive5.com](mailto:sales@alive5.com) to get started.

## Links and References

- [Alive5 AI Solutions](#)
- [Amazon Bedrock Enterprise Security](#)
- [Anthropic Trust Center](#)
- [AWS Generative AI Security](#)
- [AWS AI/ML Compliance Framework](#)

- AWS Services in Scope for Compliance

*This document contains Alive5's proprietary information about partnerships and implementation details. External performance benchmarks and compliance certifications should be verified through official AWS and Anthropic documentation.*