Security White Paper

Caitlyn's Infrastructure and Security

Caitlyn

2 in 5

Around 2 in 5 IT leaders say their security teams lack the skills needed to protect AI applications and workloads.

- cybersecuritydive.com

Caitlyn

Contents

Introduction	3
System Architecture (map)	4
System Architecture	5–7
Key Security Features	8
Data Privacy and Compliance	9
Authentication and Authorisation	10
Model Integrity and Threat Management	11
Monitoring and Incident Response	12
Conclusion	13

Introduction

Caitlyn is Custom D's advanced generative AI solution, designed to deliver research-driven insights across industries. Built on Amazon Bedrock, Caitlyn features a robust security framework, including hardened containers, zero-trust data storage policies, and advanced compliance measures. Its architecture is purpose-built to handle sensitive data securely while addressing the needs of diverse sectors.

Leveraging a proven track record in FinTech and InsureTech, alongside AWS Specialty Security certifications, Custom D applies its deep security expertise and commitment to continuous improvement. This document highlights Caitlyn's sophisticated architecture and security features, demonstrating its ability to provide reliable, scalable, and secure AI-driven solutions for organisations across industries.

aws certified
Security
SPECIALTY

aws
PARTNER
Advanced Tier
Services

aws

System Architecture

Caitlyn's architecture is designed to run primarily within your own AWS account, ensuring full control, supporting data sovereignty, and meeting compliance mandates.

By leveraging AWS services, Caitlyn provides secure data storage, processing, and inference capabilities. Its architecture prioritises scalability, resilience, and security through multi-layered defences, with only Caitlyn's container image repository and anonymised telemetry operating outside the customer's environment.

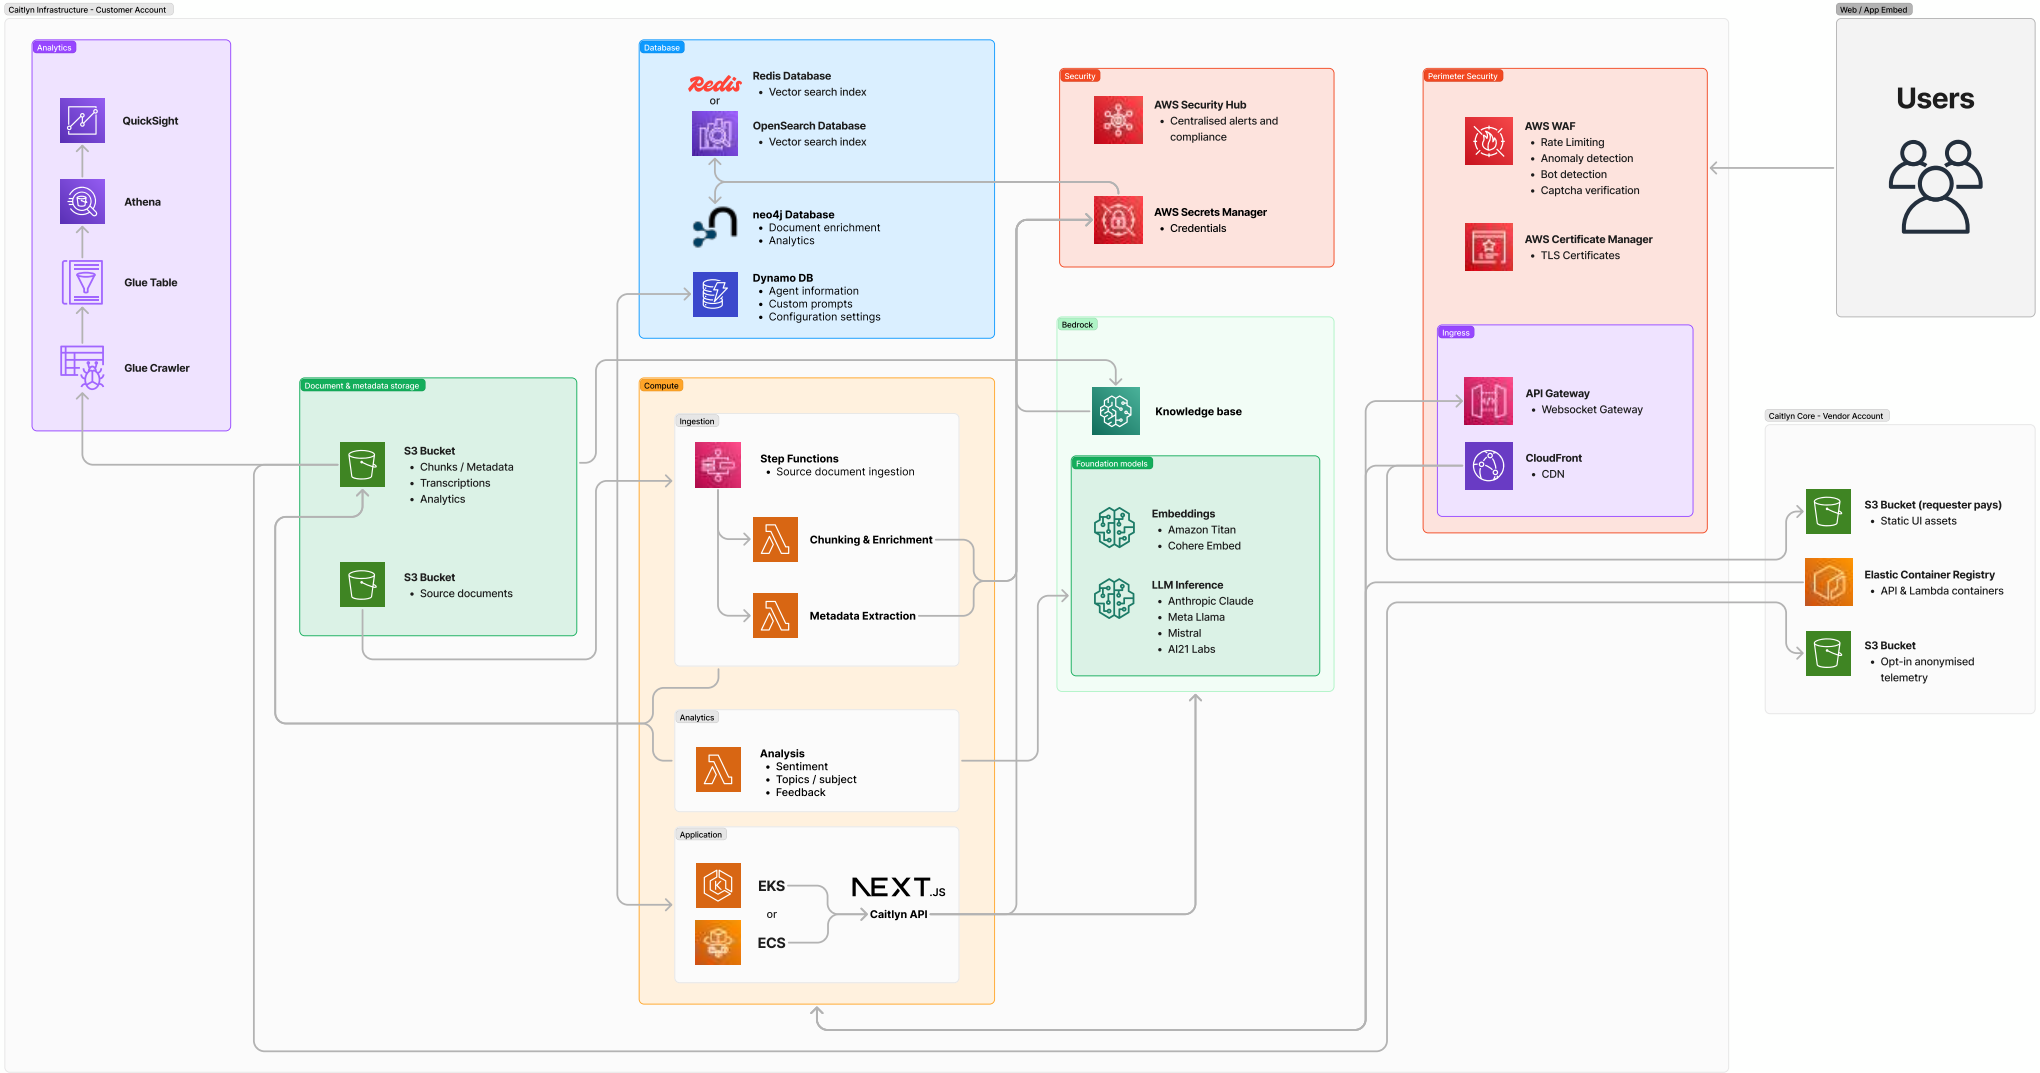

System Architecture

Data Storage

Caitlyn stores data across multiple AWS services:

Amazon S3 Buckets: Caitlyn uses S3 for source data, metadata, and processed content.

- Encryption: Data at rest is encrypted with AES-256 encryption, ensuring protection against unauthorised access.
- Access Controls: IAM roles and bucket policies enforce least-privilege access, limiting exposure to sensitive data.
- Object Lock and Versioning: Enabled to ensure data integrity, supporting Write Once, Read Many (WORM) for critical data. Version retention can be managed by S3 lifecycle rules, with a default “non current” version retention of 6 years. Previous versions are utilised by chat sessions that used that document version for inference, so we can clearly show citations, and updated content.

DynamoDB: Holds configuration settings, agent-specific information, and custom prompts. Protected by role-based authentication and granular IAM policies, DynamoDB ensures access is strictly limited to authorised users and services.

Knowledge base indexing options: Ensuring fast retrieval for AI-driven queries.

- Redis: Operates as a vector database with end-to-end encryption and controlled access to manage search data.
- OpenSearch: Provides hybrid keyword and vector search capabilities, secured by IAM policies and encrypted storage.
- Neo4J (required): Serves as a graph database for context enrichment, with end-to-end encryption to ensure sensitive relationships and insights are securely protected.

Data Processing

Caitlyn employs secure, scalable workflows powered by AWS services to ensure reliable and compliant data handling:

File Detection and Processing: AWS Step Functions securely orchestrate workflows for tasks such as chunking, enrichment, metadata extraction, and analytics. All operations are conducted within a controlled environment, adhering to strict access and encryption standards.

Containerised Execution: Caitlyn runs on Amazon ECS with Fargate, with optional deployment on EKS for larger or custom infrastructure needs. Containers are built from hardened images, container images patched monthly, and scanned using CI/CD-integrated security tools to proactively address vulnerabilities.

Access Integration: Caitlyn integrates with your organisation’s SSO or SAML providers, enabling seamless authentication. With out-of-the-box support for RBAC, granular user and group access rules can be applied to files, ensuring only authorised users can access or manage sensitive data.

Source File Retention: Source files are securely retained to support citations and traceability. Storage policies ensure these files are encrypted and access is tightly controlled to maintain compliance and security throughout their lifecycle.

System Architecture

Data Analytics and Logging

Caitlyn provides secure and comprehensive analytics and logging capabilities, ensuring sensitive data is protected throughout these processes while supporting compliance, monitoring, and operational insights:

- Analytics Tools: AWS services such as QuickSight, Athena, Glue Tables, and Glue Crawlers are used to process and analyse data securely. These tools enforce encryption in transit and at rest, ensuring that analytics operations are conducted without exposing sensitive information. Strict IAM policies govern access to these services, limiting visibility to authorised users.
- Logging and Auditing: AWS CloudTrail and S3 Access Logs capture detailed records of system and user activities, creating a robust audit trail. Logs are encrypted and access-controlled, ensuring they are protected from unauthorised access while enabling compliance with frameworks like GDPR and SOC 2.
- Retention Policies: All analytics and logging data are subject to configurable retention policies, allowing organisations to determine how long logs and analytical outputs are preserved. This ensures compliance with regulatory requirements while controlling storage costs. By default, analytics and telemetry data are retained for 6 months.
- Telemetry: Anonymised telemetry data is used to optimise system performance and user experience. To protect privacy, no customer-specific data is included in telemetry, and organisations can opt out of telemetry sharing to align with their internal security requirements.

These measures ensure that analytics and logging processes not only deliver valuable insights but also uphold the highest standards of data protection, retention, and compliance.

AI Model Inference

Caitlyn's AI inference processes are designed with robust security and privacy measures to ensure data is handled responsibly:

- Secure Model Hosting: Caitlyn leverages Amazon Bedrock to access foundation models, ensuring inference operations occur within a secure environment. Bedrock prevents third-party models from accessing customer data, ensuring it is excluded from model training or refinement, safeguarding sensitive information.
- Compliance Assurance: Caitlyn's inference processes are designed to align with global privacy and security standards, such as GDPR and SOC 2, by minimising data usage, enforcing strict access controls, and maintaining detailed audit logs to support transparency, accountability, and regulatory compliance.

This demonstrates our commitment to aligning Caitlyn's inference workflows with best practices and regulatory standards.

System Architecture

Content Delivery and API Security

Caitlyn's content delivery and API security measures ensure reliable performance and robust protection against threats:

- Amazon CloudFront: Provides a secure content delivery network (CDN) for end-user access to Caitlyn, ensuring fast and reliable access. CloudFront protects endpoints using HTTPS/TLS encryption and includes default DDoS protection through AWS Shield. Customers can optionally enable AWS Shield Advanced for enhanced threat mitigation.
- AWS WAF: Implements rate limiting, anomaly detection, and CAPTCHA verification to safeguard APIs against bots and unauthorised access. Managed and custom rule sets provide layered security tailored to specific requirements.
  - AWS Managed Rules
    - AWS WAF Bot Control rule group
    - Amazon IP reputation list
    - Anonymous IP list
    - Core rule set
    - Linux operating system
  - Custom Rules
    - Inference abuse protection
- RBAC Integration: Caitlyn supports integration with organisational SSO and SAML providers, enabling granular API access control through role-based access control (RBAC) configurations.

These measures provide high-performance content delivery while maintaining stringent API security, ensuring data integrity and protection from unauthorised access.

Key Security Features

Caitlyn's security framework implements a multi-layered approach to protect data, ensure compliance, and mitigate risks:

Zero-Trust Architecture
Caitlyn employs zero-trust principles across its infrastructure, including restrictive IAM policies, encrypted data storage, and access controls to minimise exposure and enforce least-privilege access.

Container Security
Hardened containers are regularly patched and scanned using CI/CD-integrated security tools, ensuring vulnerabilities are addressed before deployment.

DDoS Protection
All endpoints are protected by AWS Shield, with the option to enable Shield Advanced for enhanced protection. An organisations existing Shield Advanced subscription can cover Caitlyn's endpoints with no additional cost.

Endpoint Protection
AWS WAF provides robust endpoint protection, including rate limiting, anomaly detection, and CAPTCHA verification to safeguard external interfaces from unauthorised access, bot activity, and other threats.

RBAC and SSO Integration
Caitlyn integrates with leading SSO providers using both OAuth and SAML protocols, enabling seamless, secure access for users. Support is available for Microsoft Entra ID, Amazon Cognito, and Google Workspace, with the flexibility to accommodate custom authentication providers as well.

Logging and Auditing
Detailed logs of system and user activity are maintained using AWS CloudTrail, CloudWatch Logs and S3 Access Logs, providing a robust audit trail for compliance and anomaly detection.

Compliance-Ready Design
Caitlyn aligns with global regulatory standards such as GDPR and SOC 2, ensuring secure and compliant operations with transparent audibility.

These features provide a secure, scalable, and resilient foundation for Caitlyn's operations, ensuring robust protection across all layers of its architecture.

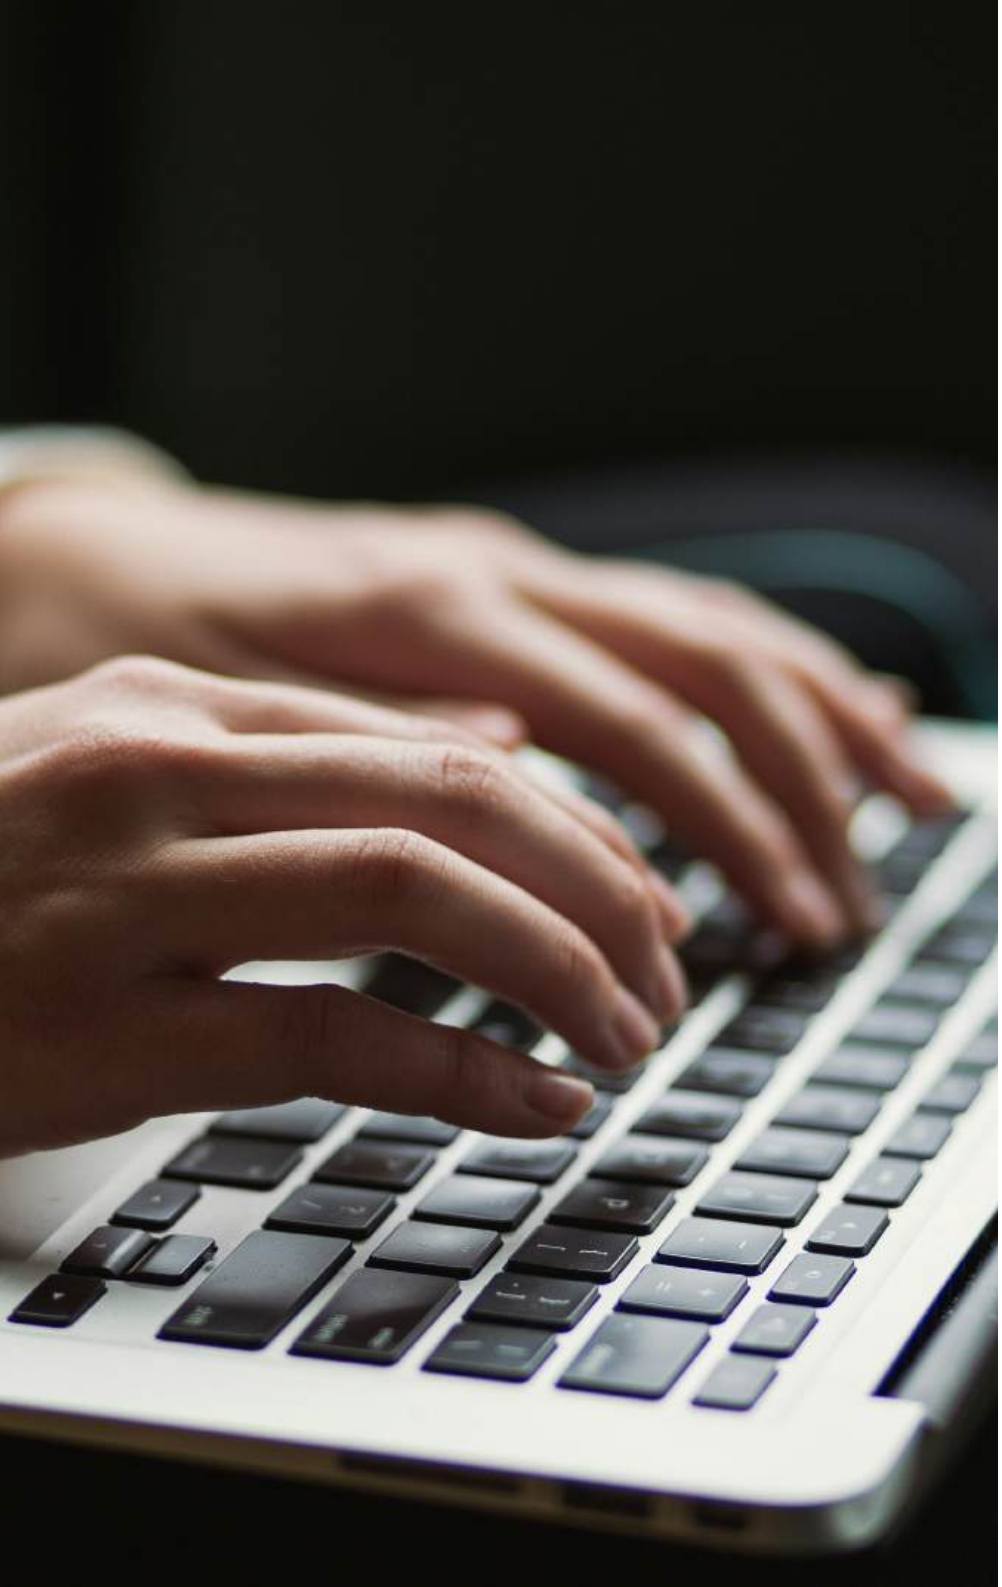

Data Privacy & Compliance

Caitlyn meets rigorous data privacy and compliance standards, ensuring sensitive information is handled responsibly and transparently:

Data Minimisation:
Caitlyn processes only the data necessary for its operations, reducing risk exposure by adhering to principles of data minimisation. Inputs and outputs are scoped to prevent unnecessary data retention.

Compliance Alignment:
Caitlyn's infrastructure and workflows are designed to align with global regulatory standards, such as GDPR and SOC 2. This includes measures like encrypted storage, strict access controls, and robust logging for audibility.

Encryption Standards:
All data is encrypted both at rest (AES-256) and in transit (TLS), ensuring protection against unauthorised access during processing and storage. These encryption practices extend to analytics, inference inputs/outputs, and logs.

Retention Policies:
Data retention is configurable to meet organisational and regulatory requirements. By default, logs and analytical outputs are retained for six months, with options to adjust retention periods to balance compliance and cost-efficiency.

Data Control and Transparency:
Caitlyn operates primarily within the customer's AWS account, granting full control over infrastructure and compliance practices. Anonymised telemetry, used for performance enhancements, is optional and never includes identifiable customer data.

Authentication and Authorisation

Caitlyn employs robust identity and access management (IAM) measures to secure user interactions and protect sensitive data. These measures ensure that Caitlyn's access controls are both secure and flexible, enabling organisations to manage access efficiently while maintaining high levels of security.

SSO and SAML Integration

Caitlyn integrates seamlessly with leading SSO providers using OAuth and SAML protocols, supporting out-of-the-box compatibility with Microsoft Entra ID, Amazon Cognito, and Google Workspace. Custom authentication providers can also be configured to meet unique organisational requirements.

Role-Based Access Control (RBAC)

Caitlyn's architecture supports granular RBAC, allowing organisations to define fine-grained user and group permissions. This ensures that access to sensitive data and features is limited to authorised users only.

This includes access to knowledge base management and tuning functionality, as well as user access to knowledge base content

Multi-Factor Authentication (MFA)

MFA adds an additional layer of security to prevent unauthorised access, requiring users to verify their identity through a secondary factor.

Access Auditing

Permissions are regularly reviewed to ensure alignment with organisational policies. Logs of authentication events are maintained to provide transparency and support compliance requirements.

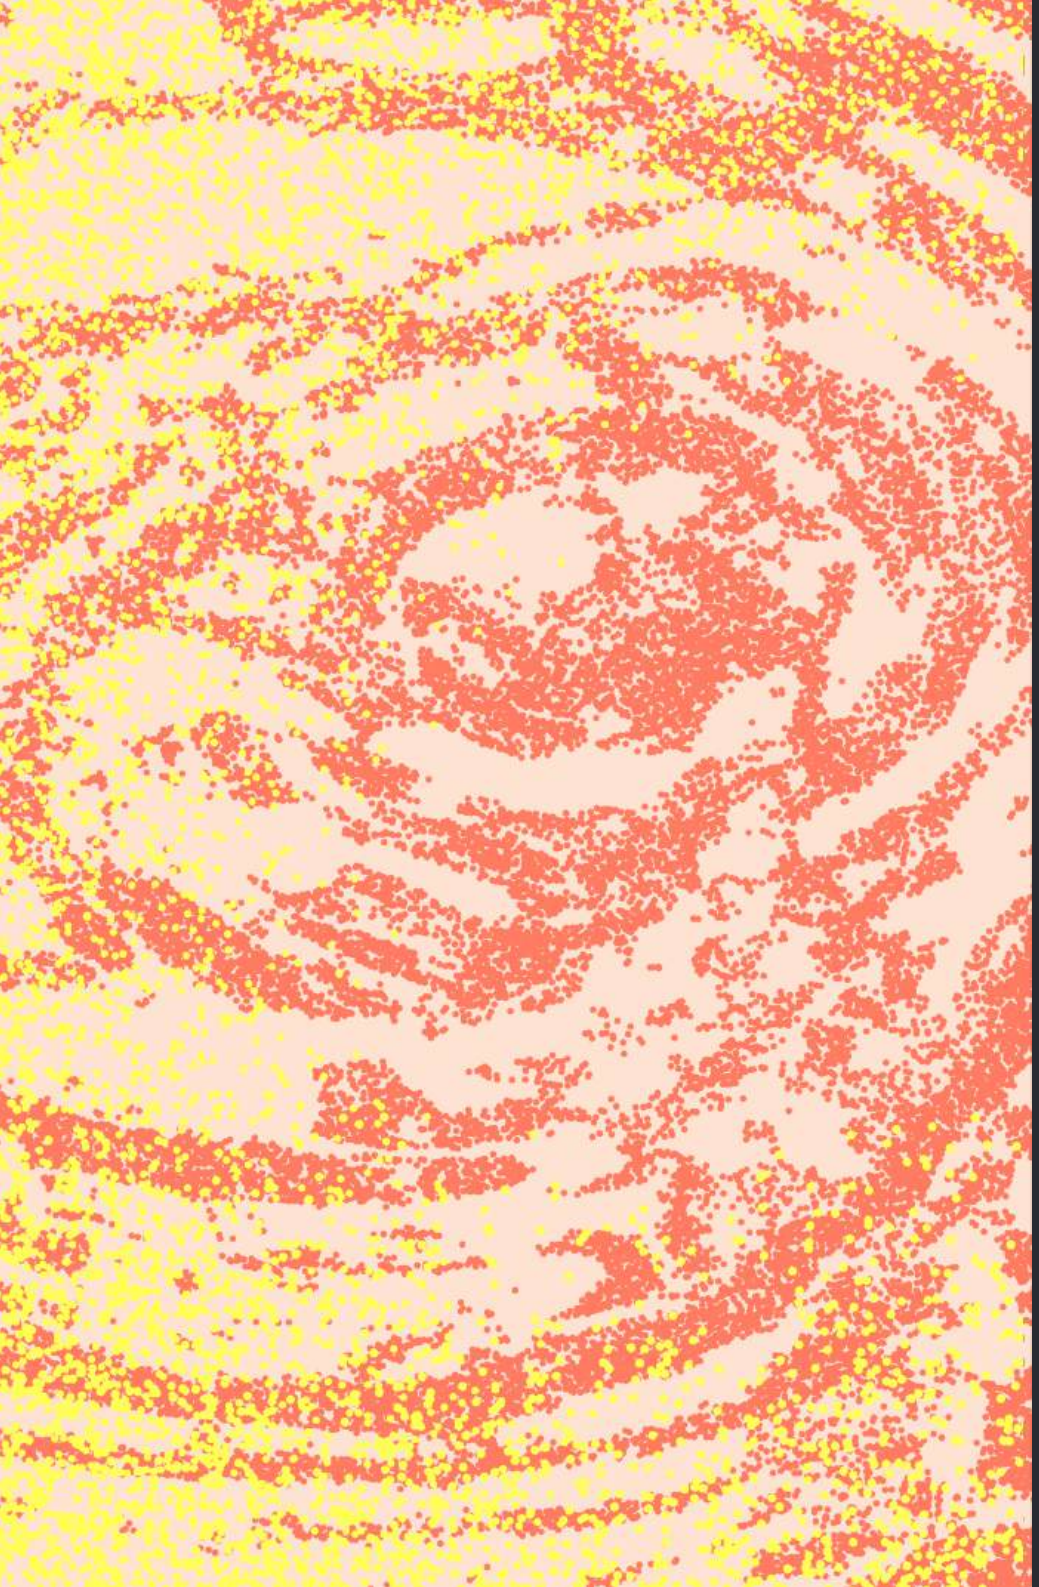

Model Integrity & Threat Management

Caitlyn's AI models rely on controlled data inputs and are protected from adversarial threats:

Data Validation and Controlled Inputs:
Input data is validated to prevent injection of malicious or malformed content. Controlled input channels reduce the likelihood of adversarial manipulation.

Secure Model Hosting
Caitlyn leverages Amazon Bedrock to host and manage foundation models in a secure environment. Bedrock ensures that customer data is isolated and excluded from model training or refinement processes.

Malicious Prompt Detection
Caitlyn employs unit tests and validation processes to detect and mitigate risks from malicious prompt inputs, protecting the model from adversarial use.

Regular Updates
Models are updated periodically to enhance accuracy, relevance, and security. Updates incorporate insights from Custom D's extensive experience across diverse deployments.

Compliance Alignment
Model workflows align with regulatory standards like GDPR and SOC 2, with detailed audit logs maintained to support accountability and compliance reporting.

Monitoring & Incident Response

Caitlyn's monitoring and incident response strategies are designed for rapid detection, containment, and resolution of security events:

Continuous Monitoring
AWS Security Hub centralises monitoring across all services, aggregating alerts for quick detection of vulnerabilities and anomalies. CloudWatch provides additional real-time performance and infrastructure monitoring.

Incident Detection and Resolution
Custom D's incident response plan includes threat detection, containment, root cause analysis, and resolution. Notifications are issued to stakeholders as needed, and all incidents are reviewed post-resolution to refine security practices.

Infrastructure Monitoring
AWS CloudTrail and S3 Access Logs capture detailed records of all activities, creating a complete audit trail to support investigations and compliance requirements.

Proactive Threat Management
Integration with AWS WAF and Shield Advanced provides enhanced protections against DDoS attacks and other external threats, ensuring minimal service disruption.

Compliance-Driven Logging and Alerting
Caitlyn supports compliance frameworks like GDPR and SOC 2 by ensuring that logs are securely stored and alerts are generated for any suspicious activities.

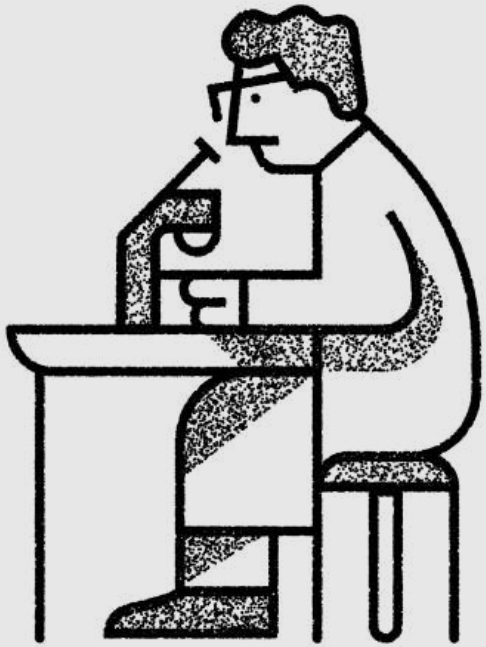

Conclusion

With decades of experience in custom development and infrastructure, particularly in FinTech and InsureTech, Custom D has built a reputation for delivering secure, scalable, and innovative solutions. This expertise, honed over more than 20 years, forms the foundation of Caitlyn's design.

Caitlyn's architecture reflects Custom D's deep understanding of security and operational requirements, providing organisations with a trusted AI platform that prioritises compliance, resilience, and performance. Backed by a proven track record and real-world deployments, Caitlyn delivers the reliability and security that businesses demand in today's fast-evolving digital landscape.