Caitlyn Security White Paper

Caitlyn's Infrastructure and Security

CAITLYN

CUSTOM D

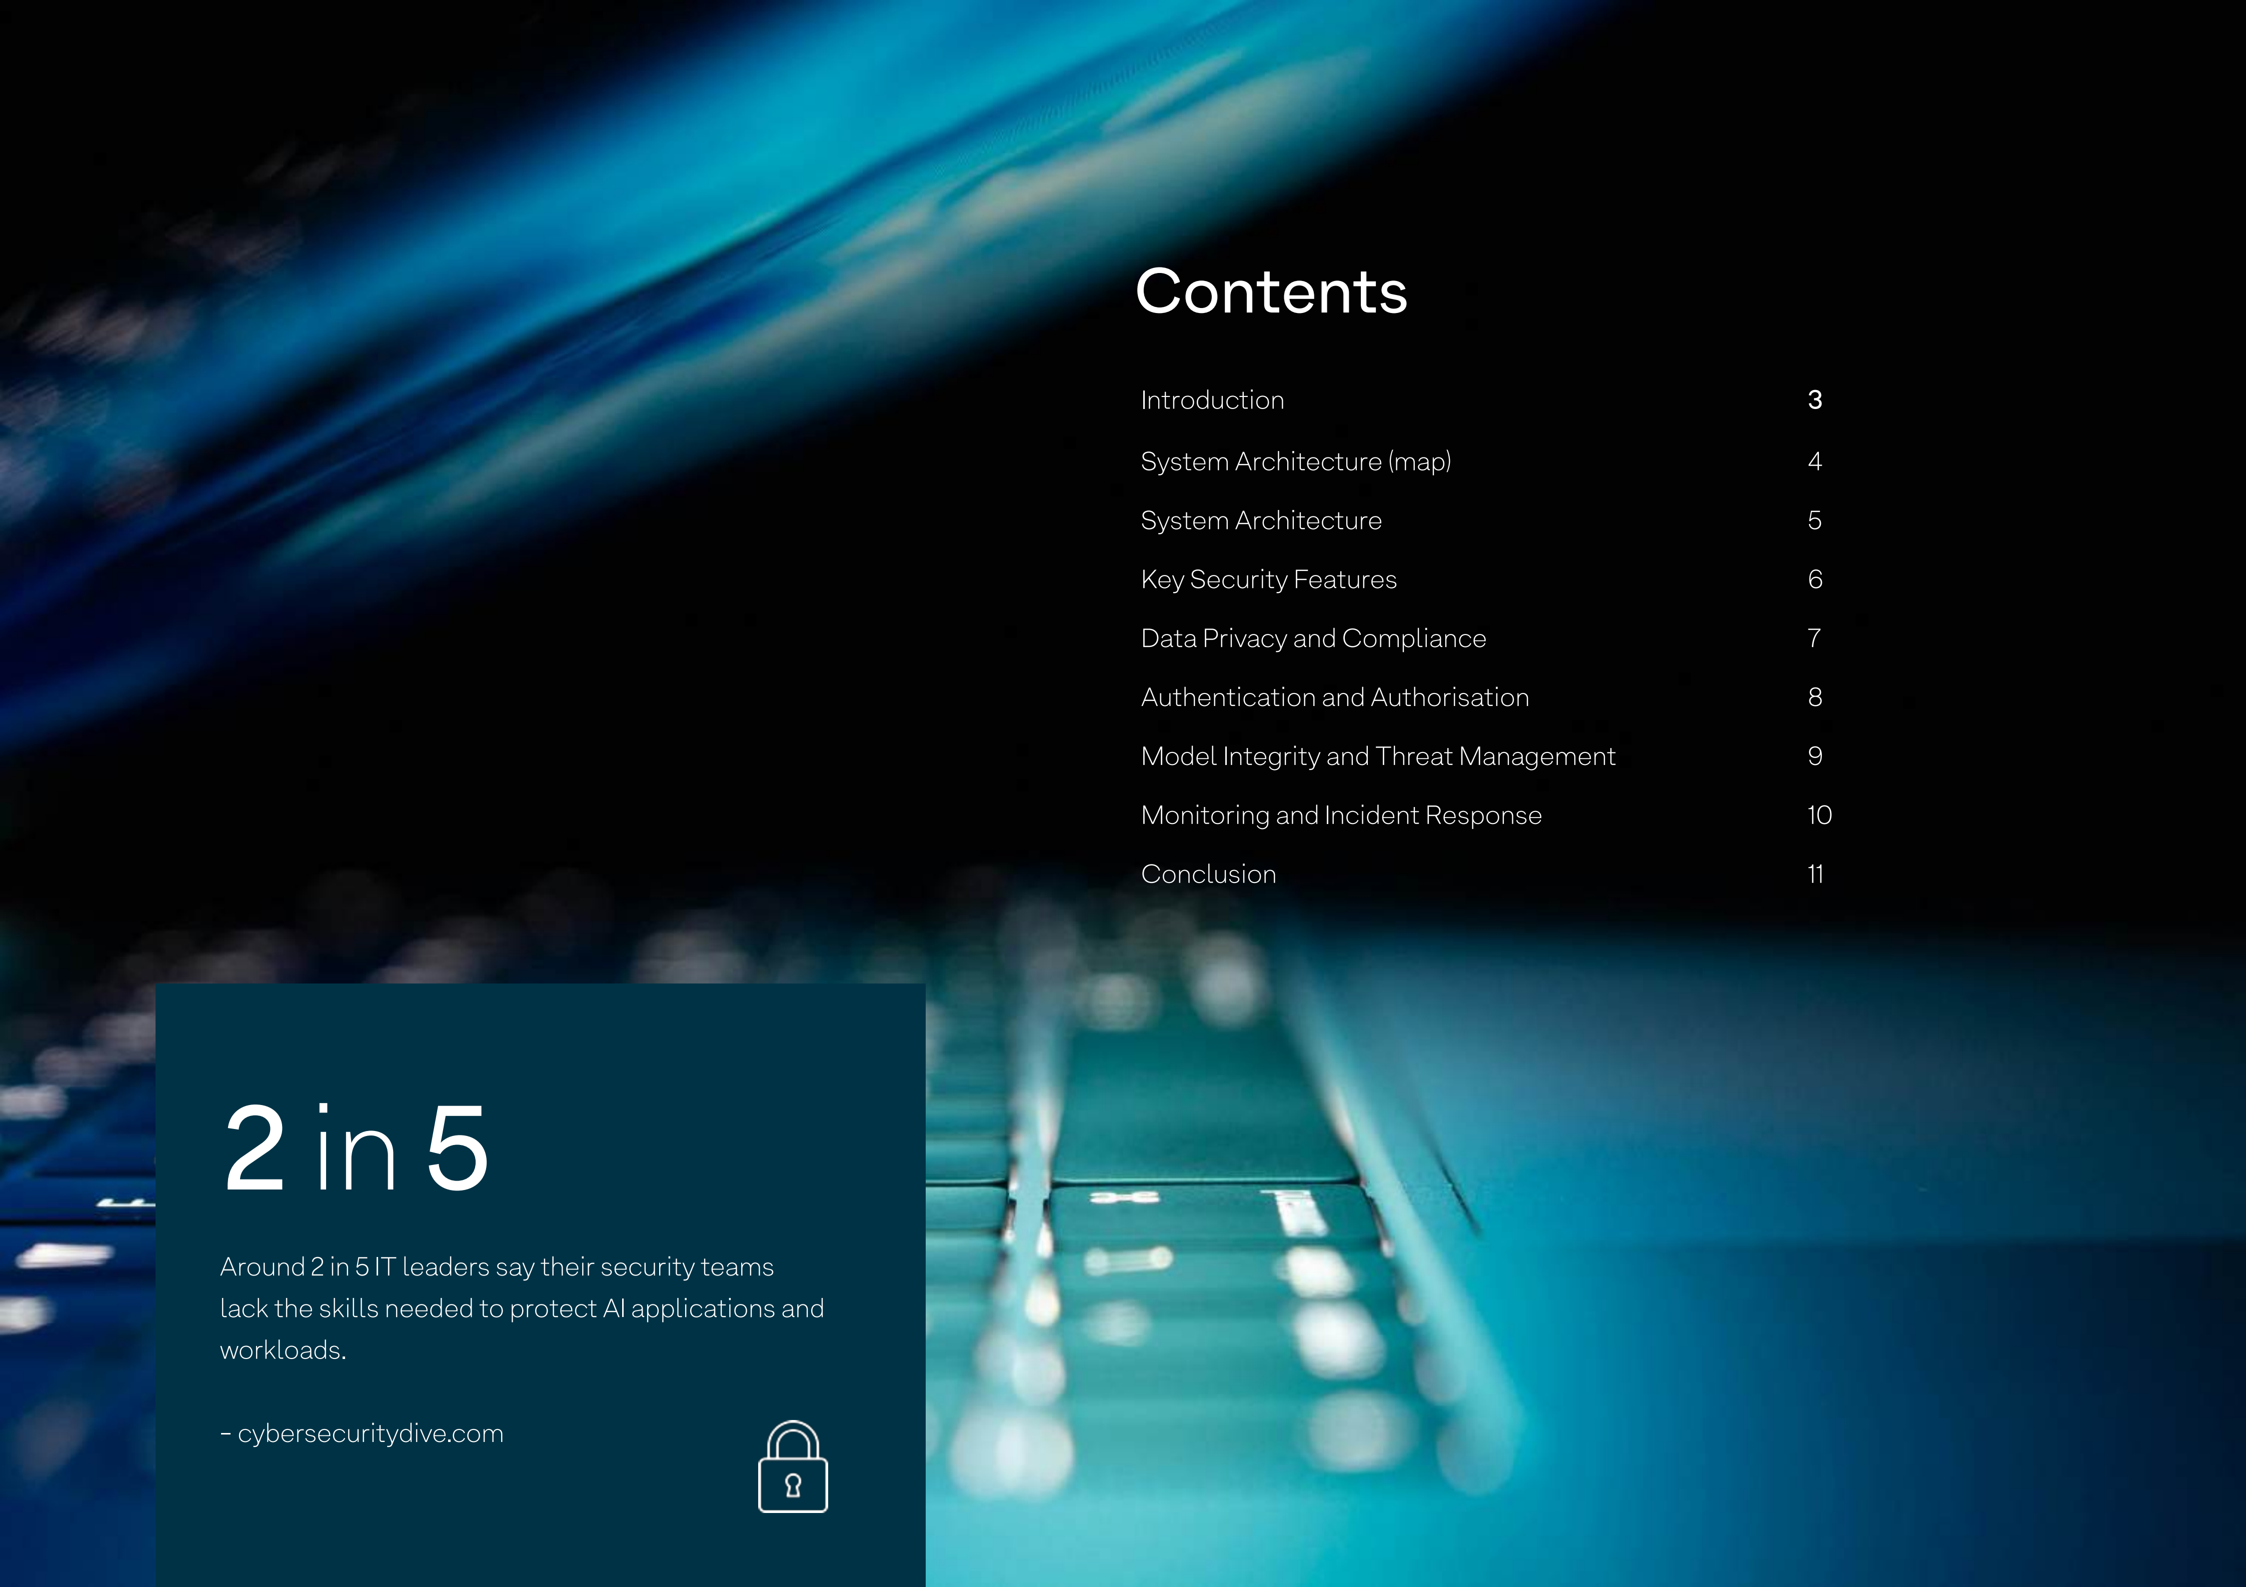

Contents

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

Introduction

Custom D's commitment to security is embedded in Caitlyn's design, from hardened containers and CI/CD-integrated security scans to a comprehensive IAM strategy.

By utilising AWS's advanced security tools and adhering to best practices, Caitlyn maintains high standards for data security and regulatory compliance. As security needs evolve, Custom D will continue enhancing Caitlyn's defences, ensuring it remains a secure, reliable solution for the agricultural industry.

aws

aws certified

Security

SPECIALTY

aws

PARTNER

Advanced Tier Services

System Architecture

Caitlyn's infrastructure is built on AWS services, combining secure storage, data processing, and model inference components, designed with scalability and resilience in mind.

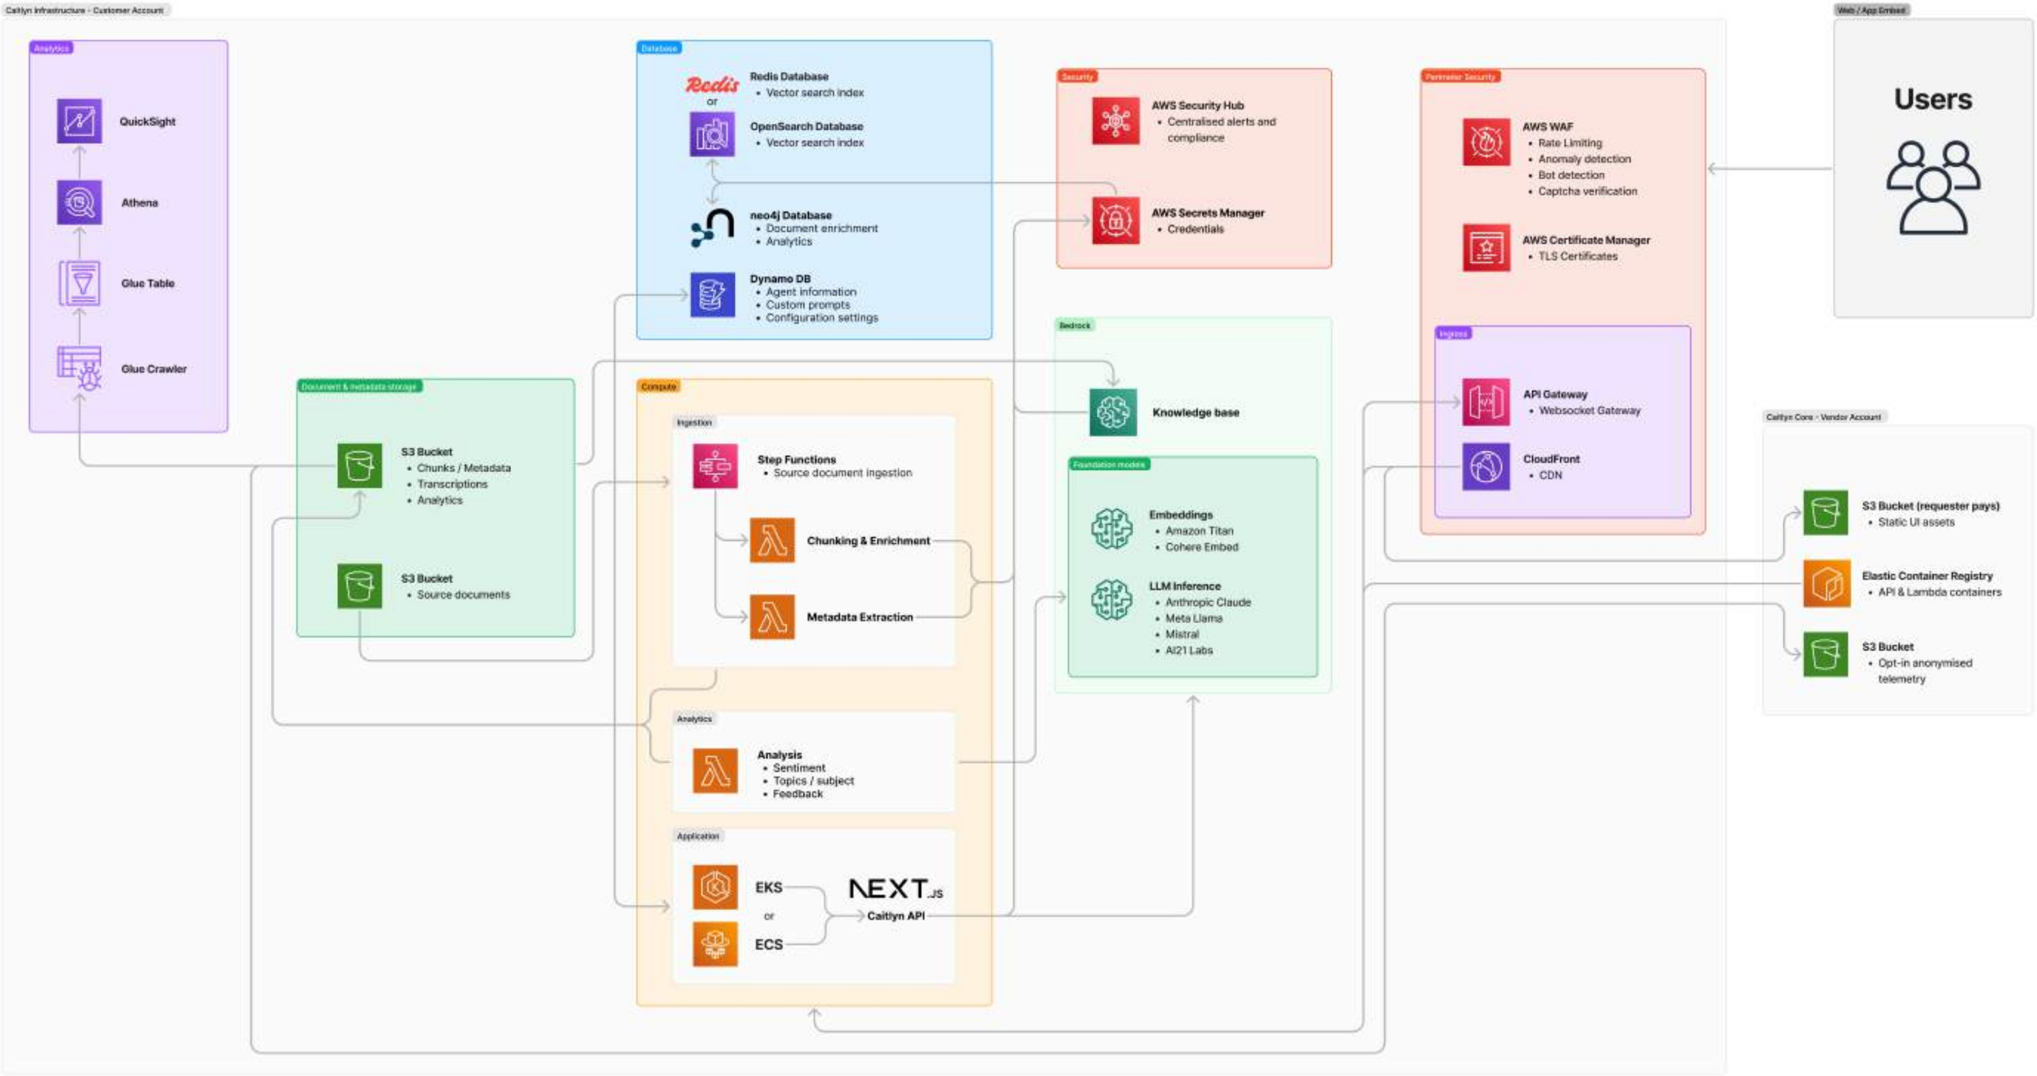

System Architecture

Data Storage

Caitlyn stores data across multiple AWS services:

Amazon S3 Buckets: Used for storing source data, processed chunks, metadata, and transcriptions, each bucket is configured with encryption and restrictive bucket policies.

DynamoDB: Holds agent-specific information, custom prompts, and configuration settings.

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

- **Redis:** Serves as a vector database for indexing and managing search data.
- **OpenSearch:** Serves as a hybrid vector & keyword search index.
- **Neo4J (required):** Serves as a graph database for enriching document context.

Data Processing

Processing workflows leverage AWS Lambda functions and Step Functions to handle:

File Detection and Processing: Step Functions trigger Lambda functions for chunking, enrichment, metadata extraction, and analytics.

Containerised Execution: Caitlyn typically runs on Amazon ECS with Fargate for container orchestration, though EKS is also an option for larger deployments, or those wishing to leverage existing infrastructure. The containers are hardened for security, and container security scans are integrated into the CI/CD pipeline to detect vulnerabilities early.

Data Analytics and Logging

QuickSight, Athena, Glue Table, and Glue Crawler are represented in the Analytics section, supporting Caitlyn's logging and analytics functions for monitoring, compliance, and data insights. These all run within your AWS account.

Telemetry: Delivers anonymised usage data to our central Caitlyn AWS account, with an opt-out option your implementation.

AI Model Inference: Caitlyn uses foundation models for advanced natural language processing:

- **LLM Inference (Claude Sonnet 3) and Embeddings Model (Cohere Embed-English)** are employed for language understanding and response generation.
- **Amazon Bedrock** supplies the foundational models, ensuring reliability and scalability.

Content Delivery and API Security

Amazon CloudFront: Provides a secure content delivery network (CDN) for Caitlyn's API, enhancing performance and protecting endpoints from distributed threats.

AWS WAF: Configured to include rate limiting, anomaly detection, and a full-screen CAPTCHA to mitigate bot activity and unauthorized access attempts.

Key Security Features

Caitlyn's architecture leverages AWS's security services to implement a comprehensive, multi-layered defence strategy:

CloudFront and WAF Protection:

- API Security: CloudFront ensures secure content delivery with HTTPS/TLS encryption. AWS WAF adds protection through rate limiting, anomaly detection, and a CAPTCHA feature to prevent unauthorised access.

- DDoS Protection (Optional with AWS Shield): While AWS Shield is available as an option, it further strengthens Caitlyn's defences against distributed denial-of-service (DDoS) attacks.

Container Security:

- Hardened Containers: Caitlyn's ECS containers are hardened to mitigate vulnerabilities.

- Container Security Scans: Integrated into CI/CD pipelines, these scans detect and remediate potential threats early, ensuring that deployed containers meet stringent security standards.

Application Security:

- SAST Scans, and known vulnerability audit tools are run in CI/CD during merge review, and pre-deploy pipelines.

Centralised Security Monitoring:

- AWS Security Hub aggregates security alerts, supporting compliance tracking and quick detection of vulnerabilities across Caitlyn's infrastructure.

Secrets Management

- AWS Secrets Manager securely stores and manages sensitive information like API keys and database credentials, ensuring that only authorised services and users can access critical information.

S3 Secure Configurations

- 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.

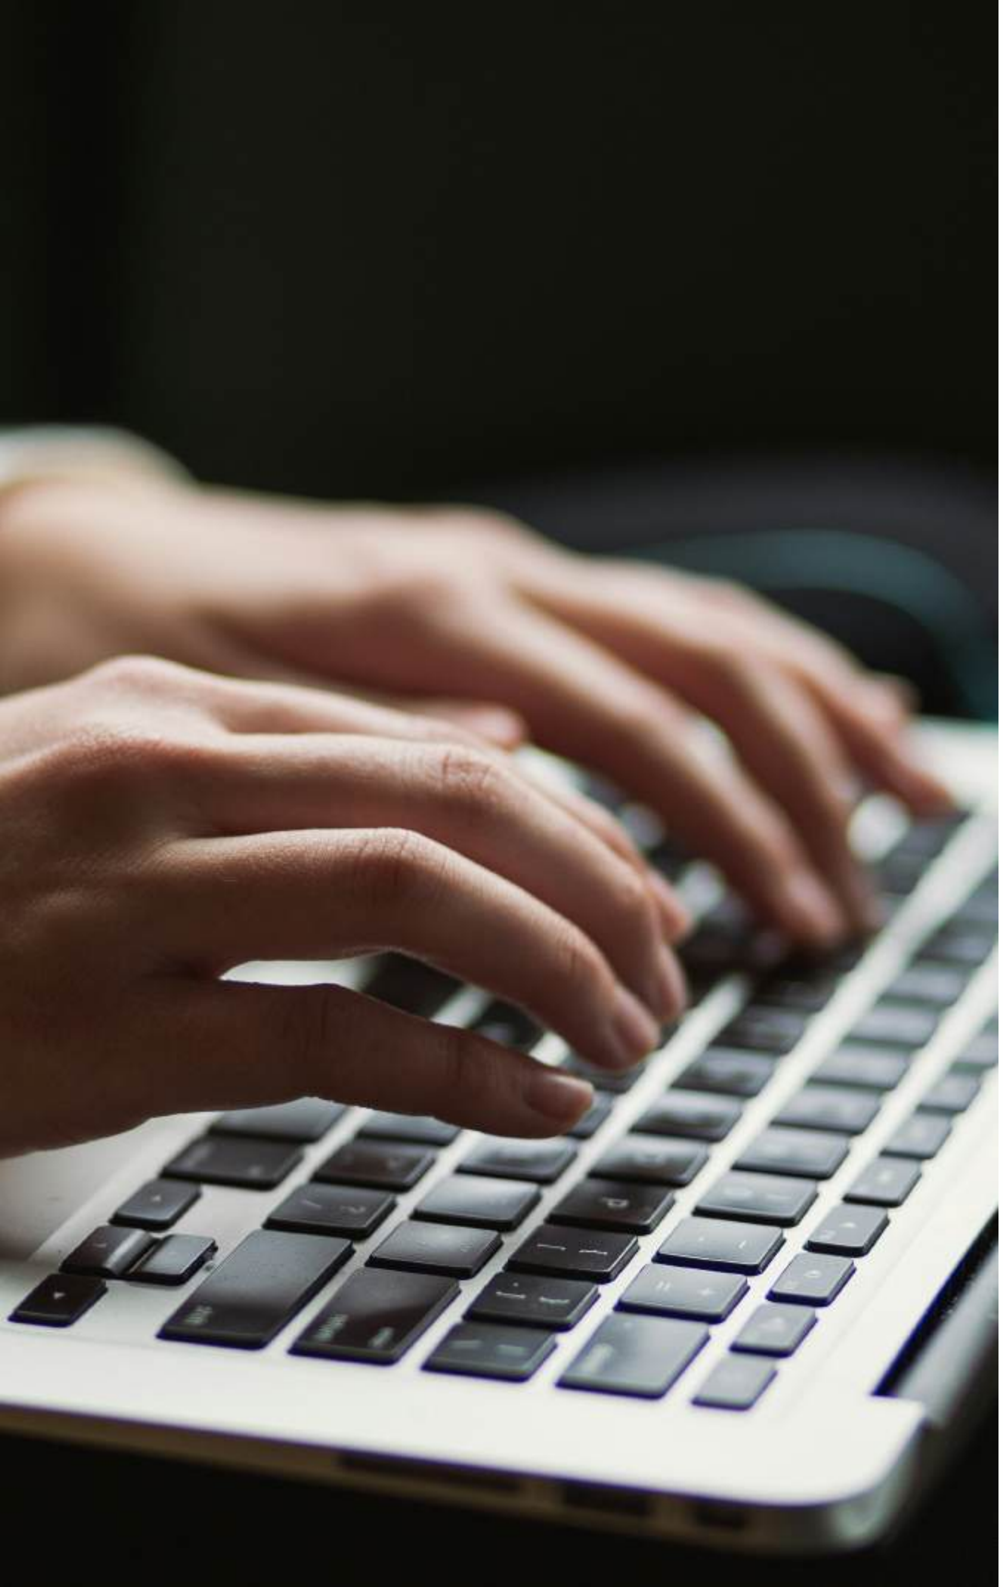

Data Privacy & Compliance

Caitlyn is designed to comply with data privacy regulations, including GDPR, to ensure transparency and user control over data:

Data Minimisation

Caitlyn collects only essential data, limiting stored information to reduce risk.

Anonymisation and User Control

Data can be anonymised for analytics purposes, and users can request data deletion or modification, ensuring that Caitlyn aligns with privacy regulations.

Compliance Monitoring

AWS Security Hub provides continuous compliance assessments to ensure Caitlyn's practices align with regulatory requirements and industry standards.

Authentication and Authorisation

Identity and access management (IAM) is foundational to Caitlyn's security framework, ensuring that access is controlled and limited:

Principle of Least Privilege

IAM policies follow the least-privilege principle, granting minimal access necessary for each role.

Multi-Factor Authentication (MFA)

MFA provides an additional layer of security for sensitive operations.

Access Reviews

Permissions are regularly audited to maintain alignment with Caitlyn's security policies, reducing the risk of unauthorised access.

Authentication Providers

The Caitlyn application supports role-based authentication with granular permissions, which can be mapped to common external authentication providers, such as Google Workspace, and Microsoft Entra ID.

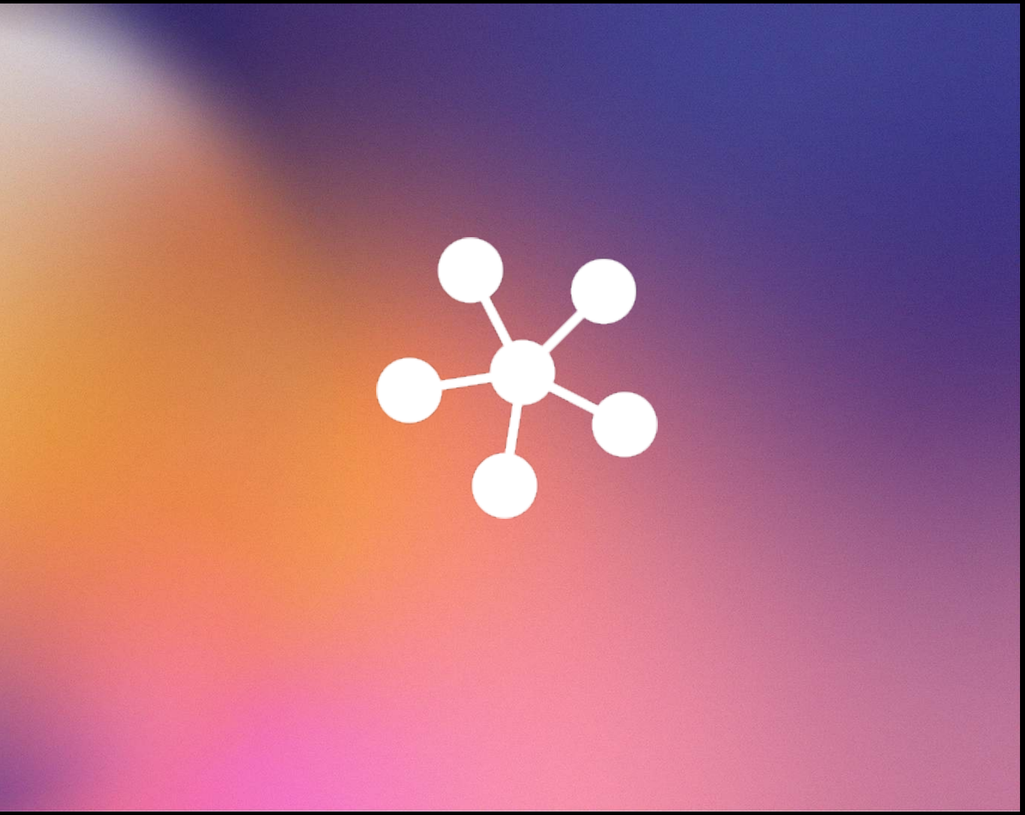

Model Integrity & Threat Management

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

Data Validation and Controlled Inputs

Models are trained on validated data sources, reducing the likelihood of model manipulation or exploitation.

Malicious Prompt Detection

Unit tests are conducted to detect and mitigate risks from malicious prompt inputs, helping protect the model from adversarial use.

Regular Model Updates

Models are updated to ensure relevancy, accuracy, and security, with tight control over training data to maintain model integrity.

Monitoring & Incident Response

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

Continuous Monitoring

AWS Security Hub centralises monitoring and alerting across all AWS services, enabling swift response to anomalies.

Application Monitoring

GlitchTip or Sentry is used for real-time application performance and error monitoring, allowing quick identification and remediation of issues that may impact security. This is typically run as a centralised Caitlyn service, but can also be deployed and run within your own AWS account if required.

Infrastructure Monitoring

CloudWatch Alarms are configured for ECS and EKS clusters, ensuring prompt alerting for performance and security-related issues within the compute infrastructure.

Logging and Auditing

AWS CloudTrail and S3 Access Logs track all user interactions, providing an audit trail for investigating potential issues.

Incident Response Plan

Custom D has a structured incident response plan in place, covering threat detection, containment, and root cause analysis. Notifications are issued to affected stakeholders as required, and all incidents are reviewed post-resolution to refine security practices.

Conclusion

Custom D's commitment to security is embedded in Caitlyn's design, from hardened containers and CI/CD-integrated security scans to a comprehensive IAM strategy.

By utilising AWS's advanced security tools and adhering to best practices, Caitlyn maintains high standards for data security and regulatory compliance. As security needs evolve, Custom D will continue enhancing Caitlyn's defences, ensuring it remains a secure, reliable solution for the agricultural industry.

CAITLYN

CUSTOM D

getcaitlyn.ai

Copyright © Caitlyn | Developed by the team at Custom D

customd.com