1. Executive Summary

Our SaaS products are AWS-native deployed entirely within the customer's AWS account via a single CloudFormation stack. No customer data leaves the account boundary. This document describes the security architecture, data handling practices, and compliance alignment for enterprise evaluation. This architecture eliminates third-party risk by design — no vendor-hosted infrastructure processes or stores customer data. The entire software stack operates as first-party resources within the customer's AWS estate, inheriting the security posture of the customer's existing controls.

2. Deployment Model

- All resources (Lambda, S3, OpenSearch, Firehose, SQS, DynamoDB, VPC) are created in the customer's account and region.
- The product is delivered as a CloudFormation template with pre-built Lambda deployment packages hosted in a public-read S3 bucket. After deployment, no ongoing dependency on Perfware infrastructure except for periodic AWS entitlement checks.
- Customer retains full ownership and control of all resources, encryption keys, and data.
- Infrastructure-as-Code (CloudFormation) ensures every deployment is auditable, repeatable, and version-controlled. No manual configuration or drift.
- Deployment packages are integrity-verified via S3 ETags and deployed atomically — partial deployments are automatically rolled back by CloudFormation.

3. Data Flow & Boundaries

- Log data flows: CloudWatch Logs -> Subscription Filter -> Firehose -> S3 (encrypted) -> Lambda processing -> OpenSearch / Datalake.
- All data remains within the customer's VPC and AWS account at all times.
- Cross-account log forwarding (optional) uses IAM AssumeRole with customer-controlled trust policies.
- No internet egress required for core processing (VPC interface endpoints used).
- Amazon Bedrock AI calls (where applicable) are made via VPC interface endpoints. Prompts and responses never traverse the public internet. AWS's data use policy ensures customer inputs are not used for model training.
- Metric and log data processed by the system is never aggregated, sampled, or transmitted externally for analytics or telemetry purposes.

4. Encryption

- **At rest:** S3 buckets use SSE-S3 or customer-managed KMS CMK (configurable). OpenSearch uses node-to-node encryption and encryption at rest. SQS queues support KMS CMK encryption. DynamoDB uses AWS-owned keys.
- **In transit:** All internal communication uses TLS 1.2+. VPC endpoints ensure traffic stays on the

AWS backbone. ALB uses HTTPS with customer-provided ACM certificates.

- **Key management:** KMS keys are created within the customer's account with automatic annual rotation enabled. Customer retains full key policy control.
- SNS alert notifications use AWS-managed encryption (SSE). Sensitive metric values are never included in plaintext email subjects.
- DynamoDB tables support point-in-time recovery (PITR) on higher tiers, providing encrypted backup without operational overhead.

5. Authentication & Access Control

- Subscription Editor access is protected by Amazon Cognito User Pools with ALB authentication integration.
- Role-based access control: admin (full access) and viewer (read-only) roles via Cognito groups.
- MFA support: configurable as OFF, OPTIONAL, or ON via CloudFormation parameter.
- Single Sign-On: supports SAML 2.0 federation via Cognito identity providers.
- All API operations require authenticated sessions; session cookies are HttpOnly and Secure.
- OIDC federation support: Okta, Azure AD, Google Workspace, and any SAML 2.0 or OpenID Connect provider via Cognito identity federation.
- Session management: 60-second keepalive with automatic session expiry detection. No long-lived tokens stored client-side.
- API Gateway / ALB-level authentication ensures no unauthenticated request reaches application code.

6. Network Security

- Lambda functions run in isolated VPC subnets with no internet access.
- Interface VPC endpoints (SQS, CloudWatch Logs, SNS, Monitoring, Lambda, Glue, STS) provide private connectivity to AWS services.
- ALB security groups restrict HTTPS access to customer-specified CIDR ranges.
- OpenSearch domain is deployed within the VPC with security group restrictions.
- No public endpoints except the ALB (which requires Cognito authentication).
- Gateway VPC endpoints for S3 and DynamoDB ensure high-throughput data operations never leave the AWS network.
- Security group rules follow least-privilege: Lambda-to-OpenSearch on port 443 only, no broad CIDR allowances.
- Public subnet exposure is limited to the ALB listener (HTTPS/443) — no SSH, no management ports exposed.

7. Data Retention & Lifecycle

- S3 lifecycle rules automatically expire log objects based on configurable retention periods.
- OpenSearch ISM policies automatically delete indexes based on configurable retention (separate for app and audit).
- Datalake S3 objects have independent lifecycle rules per index type.
- All retention values are customer-configurable via CloudFormation parameters.
- S3 versioning with noncurrent version expiration ensures configuration history is preserved for audit while automatically cleaning old versions after 90 days.
- Incomplete multipart uploads are automatically aborted after 7 days to prevent orphaned storage charges.
- Anomaly and state records in DynamoDB use TTL-based expiration — no manual purge required.

8. Audit & Logging

- All Lambda invocations are logged to CloudWatch Logs with structured output.
- S3 access logging captures all bucket operations to a dedicated access log bucket.
- CloudTrail (customer-managed) captures all API activity across the stack resources.
- Subscription Editor maintains S3 version history of all configuration changes with user attribution.
- CloudWatch alarms monitor processing errors, queue depths, and infrastructure health.
- Embedded Metric Format (EMF) provides real-time operational telemetry (invocation counts, error rates, processing durations) without additional monitoring infrastructure.
- Configuration changes are tracked via S3 object versioning with user attribution from Cognito identity claims.
- Dead letter queues (DLQs) capture failed processing events for forensic analysis without data loss.

9. Pattern Detection & Data Handling

- Pattern detection (PII, credentials, SQL injection) runs entirely within Lambda at processing time.
- **Tag mode:** Adds metadata flags without modifying log content.
- **Redact mode:** Replaces sensitive content before indexing (original is never stored in OpenSearch/datalake).
- **Filter mode:** Drops matched events entirely – never written to any persistent store.
- Custom pattern rules are defined by the customer and executed locally.
- Detection rules execute in ephemeral Lambda environments with no persistent local storage — processed data exists in memory only during the invocation lifecycle.

10. Shared Responsibility Model

Perfware is responsible for:

- Secure software development practices for the product code.
- Ensuring the CloudFormation template follows AWS security best practices.
- Providing timely security updates via new product versions.
- Maintaining the integrity of the deployment packages in the public S3 bucket.

The customer is responsible for:

- AWS account security (IAM, CloudTrail, GuardDuty, SCPs).
- Network configuration (VPC, security groups, NACLs).
- Cognito user management and MFA enforcement.
- KMS key policies and rotation schedules.
- Data retention configuration appropriate to their compliance requirements.
- Applying product updates (new CloudFormation template versions).

11. Compliance Framework Alignment

The following summarizes how perfware.cloud aligns with common compliance frameworks:

HIPAA: Encryption at rest and in transit, access controls, audit logging, no PHI exfiltration. Customer signs BAA with AWS. Product does not store, process, or transmit PHI outside the account.

FedRAMP: Runs within customer's FedRAMP-authorized AWS environment. No external dependencies. Supports required encryption standards, access controls, and continuous monitoring.

SOC 2: Supports all Trust Service Criteria – Security (encryption, access control, network isolation), Availability (auto-scaling, alarms, DLQs), Confidentiality (pattern redaction, CMK encryption), Processing Integrity (idempotent processing, DynamoDB locks).

PCI DSS: Does not store cardholder data. Supports network segmentation, encryption, access logging, and monitoring controls.

GDPR: All data processing occurs within customer-specified AWS region. Pattern redaction supports data minimization. Customer controls data retention and deletion.

12. Vulnerability Management

- Lambda runtime: Python (AWS-managed, patched automatically).
- Dependencies: Minimal (boto3 provided by AWS runtime). Custom library zip built from pinned versions.
- OpenSearch: AWS-managed service with automatic security patching.
- Nginx (dashboards proxy): Amazon Linux 2023 AMI with automatic updates via SSM.
- No custom binaries or third-party agents deployed.

- No third-party runtime dependencies beyond the AWS SDK (boto3), which is provided and patched by the AWS Lambda runtime.
- Static analysis and infrastructure validation are performed prior to every release using CloudFormation template validation and Python AST checks.

13. Incident Response

- Security issues in the product code: Report to security@perfware.cloud. Patches provided within 48 hours for critical issues.
- Infrastructure incidents: Customer-managed via their existing AWS incident response procedures. Product provides CloudWatch alarms and SNS notifications for operational awareness.
- Diagnostics collection scripts generate support bundles containing only resource metadata and operational metrics — no customer log content or sensitive data is included.
- Circuit breaker patterns prevent cascade failures: if downstream services become unavailable, processing degrades gracefully rather than failing open.

14. Contact

Security inquiries: security@perfware.cloud

General support: support@perfware.cloud