

AWS SECURITY, COMPLIANCE & GOVERNANCE · SAMPLE DELIVERABLE

Organization Scan

Week 1 of 6: Access and scan · prepared for StrataForge Robotics, Inc.

Prepared by Charles Sieg · charles.sieg@vantalect.com · vantalect.com

July 17, 2026 · AWS Security, Compliance & Governance Audit: The Hardening (six weeks)

SAMPLE DELIVERABLE

This is a sample deliverable. StrataForge Robotics, Inc. is a fictional company: every account, resource, finding, figure, and person in this report is simulated to show the form and depth of a Vantalect engagement. AWS prices are list prices from the AWS Price List API on October 13, 2026.

Contents


Executive summary	2
What we scanned	4
Access	5
What the role can and cannot do	5
How the role is deployed	6
Method	8
What the scan does not see	9
The organization score	9
By kind of account	10
The lowest-scoring accounts	11
Failed controls	12
The Critical controls	12
The High controls with the most failures	13
Detection coverage before hardening	15
What week 2 does	17
What we need from you this week	18
Appendix: the controls in the score	18
Appendix: all accounts	22

Executive summary

StrataForge's AWS organization scores 32.3% on the Vantalect scan score (71 controls). That is the share of control results that pass across all 48 accounts on the 71 controls we evaluated: all 40 automated controls of the CIS benchmark and 31 more from AWS Foundational Security Best Practices. It is our score, computed the way Security Hub CSPM computes one, and it is not the score Security Hub CSPM would show for the same organization, which covers every control a standard includes ([What we scanned](#) lists ours). Roughly two of every three control results fail somewhere in the organization. Marcus Oyelaran: this is deliverable 1 of the Audit, the scan. It is a baseline, taken before anything has been changed, and week 6 runs the same scan the same way to measure what the Hardening did.

Four things to take from it:

- 1. **No account passes more than 44.7% of its controls.** The best account is sf-log-archive; the median is 35.5%. The lowest is kestrel-legacy-mgmt at 12.2%, and 11 accounts score under 25%. The five accounts that came with the Kestrel Vision acquisition average 18.8%, against 33.6% for the other 43.
- 2. **62 of the 71 controls fail somewhere, and 8 of the failing controls are rated Critical.**
Together they account for 898 failed resource evaluations out of 19,141. Most of those 898 are one control, Config.1, which we expect to include false positives; the rest are a handful of resources each, and a handful of public buckets or open ports is where a breach starts.
- 3. **Detection runs in a minority of the places it should.** AWS Config records in 30.1% of account-Regions, GuardDuty is on in 31.6%, and Security Hub CSPM is on in 15.8%. The organization trail covers all 816. None of the other five services runs in any of the 11 Regions where StrataForge has no workloads.
- 4. **These are raw scan results, not findings.** Week 2 verifies every Critical and High result by hand, removes the false positives, rates each real finding for exploitability and blast radius, and delivers the findings register on July 24, 2026. Do not act on a number in this report as if it were a verified finding.



The 16 lowest-scoring accounts, as the percentage of controls that pass.

What we scanned

The scan covers the whole organization: **49,156 resource evaluations across 48 accounts and 17 enabled Regions, which is 816 account-Regions**. StrataForge uses six of those Regions for workloads; every account has all 17 default Regions enabled, and an unused Region is a place an attacker can work without anyone looking. We scanned all of them.

The scope of the week-one scan.

Item	In scope
Accounts	48, in the organization whose management account is sf-management
Enabled Regions	17 per account: us-east-1, us-east-2, us-west-1, us-west-2, ca-central-1, sa-east-1, eu-central-1, eu-west-1, eu-west-2, eu-west-3, eu-north-1, ap-northeast-1, ap-northeast-2, ap-northeast-3, ap-southeast-1, ap-southeast-2, ap-south-1
Account-Regions	816
Standards	CIS AWS Foundations Benchmark v5.0.0; AWS Foundational Security Best Practices v1.0.0, as implemented by AWS Security Hub CSPM
Controls evaluated	71 distinct controls: 40 from CIS and 48 from FSBP, of which 17 are in both. The full list is in Appendix: the controls in the score ; FSBP has many more controls than the 48 we evaluated
Resource evaluations	49,156, of which 19,141 failed
Out of scope	StrataForge's separate AWS GovCloud (US) organization, which serves its defense programs, was out of scope.

Which benchmark version, and why

The CIS AWS Foundations Benchmark v5.0.0 is the newest version AWS Security Hub CSPM supports (it also supports v3.0.0, v1.4.0, and v1.2.0), and AWS recommends it. CIS itself now publishes v7.1.0, so v5.0.0 is the version we scored against, not the latest CIS release. We use it because the control IDs, severities, and pass or fail logic are the ones your Security Hub CSPM, and any tool built on it, will report. AWS Foundational Security Best Practices v1.0.0 is AWS's own list and is the better guide to services CIS does not cover, such as Amazon Inspector, EKS, and RDS. Both are Security Hub CSPM standards. The newer AWS Security Hub (generally available December 2, 2025) is a different product that correlates findings from several services; it does not change these controls.

Access

We read StrataForge's accounts with one role in each of them: `vantalect-audit-readonly`. It carries the AWS managed policies **SecurityAudit** and **ViewOnlyAccess** and nothing else. It is deployed to all 48 accounts from a CloudFormation StackSet in the management account, and it is the only access Vantalect has.

What the role can and cannot do

The role reads how resources are configured. It cannot create, change, or delete anything, and it cannot read what your systems store.

What the audit role can and cannot do.

The role can	The role cannot
List and describe resources in every service the scan covers: instances, security groups, buckets and their policies, databases, functions, clusters, keys (not their material), roles and policies	Change any resource, policy, or setting, in any account
Read the IAM credential report and account summary: which users have passwords, keys, and MFA, and when each was last used	Read the contents of S3 objects, database rows, queue messages, or log events
Read key and bucket policies, trust policies, and organization structure and policies	Read secret values (Secrets Manager, Systems Manager SecureString parameters) or use a KMS key to decrypt or sign
Read CloudTrail's configuration and its recent event history	Assume any other role, or reach your VPCs or instances

Two of the reads need a deliberate choice on our side. Lambda function configuration includes environment variable values, which can hold secrets, so the collector keeps variable names only and discards the values at the moment it receives them. And the credential report lists people by user name, which we keep in the engagement workspace and delete with the rest of the scan data after the engagement.

How the role is deployed

The role trusts a single Vantalect principal and only when the caller presents an external ID that is unique to StrataForge, so no one else who learns the role's name can use it. Elena Vasquez's team deployed it through the platform change process, as a service-managed StackSet to the organization, plus a direct stack in the management account (StackSets with service-managed permissions do not deploy to the management account).

YAML

```
AWSTemplateFormatVersion: '2010-09-09'
Description: Vantalect read-only audit role (SecurityAudit and ViewOnlyAccess). It cannot change anything.
Parameters:
  VantalectAuditPrincipal:
    Type: String
    Description: ARN of the Vantalect audit role that is allowed to assume this role
  ExternalId:
    Type: String
    NoEcho: true
    MinLength: 16
    Description: External ID agreed for this engagement
```

```

Resources:
  AuditRole:
    Type: AWS::IAM::Role
    Properties:
      RoleName: vantalect-audit-readonly
      MaxSessionDuration: 3600
      AssumeRolePolicyDocument:
        Version: '2012-10-17'
        Statement:
          - Effect: Allow
            Principal:
              AWS: !Ref VantalectAuditPrincipal
            Action: sts:AssumeRole
            Condition:
              StringEquals:
                sts:ExternalId: !Ref ExternalId
      ManagedPolicyArns:
        - arn:aws:iam::aws:policy/SecurityAudit
        - arn:aws:iam::aws:policy/job-function/ViewOnlyAccess
      Tags:
        - Key: Purpose
          Value: vantalect-security-audit
Outputs:
  AuditRoleArn:
    Value: !GetAtt AuditRole.Arn

```

BASH

```

# In the management account: create the StackSet and deploy it to every member account.
aws cloudformation create-stack-set \
  --stack-set-name vantalect-audit-readonly \
  --template-body file://vantalect-audit-readonly.yaml \
  --permission-model SERVICE_MANAGED \
  --auto-deployment Enabled=false \
  --capabilities CAPABILITY_NAMED_IAM \
  --parameters ParameterKey=VantalectAuditPrincipal,ParameterValue=<vantalect-audit-role-arn> \
    ParameterKey=ExternalId,ParameterValue=<external-id>

aws cloudformation create-stack-instances \
  --stack-set-name vantalect-audit-readonly \
  --deployment-targets OrganizationalUnitIds=<root-ou-id> \
  --regions us-east-1 \
  --operation-preferences MaxConcurrentPercentage=100,FailureToleranceCount=0

# The management account is not a StackSet target: create the same role there directly.
aws cloudformation create-stack \
  --stack-name vantalect-audit-readonly \
  --template-body file://vantalect-audit-readonly.yaml \
  --capabilities CAPABILITY_NAMED_IAM \
  --parameters ParameterKey=VantalectAuditPrincipal,ParameterValue=<vantalect-audit-role-arn> \
    ParameterKey=ExternalId,ParameterValue=<external-id>

```

Every use of the role is visible to you. Each call we make runs under a session named for the engagement, so CloudTrail in each account records the AssumeRole event and every API call after it. To see them:

BASH

```
aws cloudtrail lookup-events \
  --lookup-attributes AttributeKey=EventName,AttributeValue=AssumeRole \
  --max-results 50 \
  --region us-east-1 \
  --query 'Events[?contains(CloudTrailEvent, `vantalect-audit-readonly`)].[EventTime,Username]'
```

At the end of week 6 your engineers remove the role; we change nothing in your accounts. It takes three commands, because the StackSet does not reach the management account: delete the stack instances, delete the StackSet, and delete the stack you created directly in the management account. We then ask you to confirm in CloudTrail that nothing assumed the role afterward.

BASH

```
aws cloudformation delete-stack-instances --stack-set-name vantalect-audit-readonly \
  --deployment-targets OrganizationalUnitIds=<root-ou-id> --regions us-east-1 --no-retain-stacks
aws cloudformation delete-stack-set --stack-set-name vantalect-audit-readonly

# In the management account:
aws cloudformation delete-stack --stack-name vantalect-audit-readonly
```

Method

The scan is the same procedure that week 6 repeats, so the two are comparable by construction.

- 1. List the organization.** The account list, organizational units, and policies come from the management account's AWS Organizations API.
- 2. Assume the audit role in each account and read it.** For each of the 48 accounts and each of the 17 enabled Regions, the collector pages through the describe and list calls for every service a control touches, and reads the IAM credential report once per account. Throttled calls are retried with backoff, and a call that cannot be completed is recorded as a gap rather than a pass.

3. **Evaluate every control.** Each Security Hub CSPM control's logic is applied to the configuration collected, giving one result per account, Region, and control: how many resources were evaluated and how many failed. We evaluate the controls ourselves, rather than reading your Security Hub CSPM findings, because Security Hub CSPM is on in only 129 of the 816 account-Regions; relying on it would have scored the places that already have a scanner and missed the rest.
4. **Score the way Security Hub CSPM does.** A control passes in an account when none of its checks fail in any Region. Controls with no data (the service is not in use) are left out, and the score is passed controls divided by passed plus failed. The organization score pools all 48 accounts' results. We call it the Vantalect scan score because it covers the 71 controls in [Appendix: the controls in the score](#), not every control in the two standards. Control titles and severities are Security Hub CSPM's own.
5. **Keep the raw results.** Every failed check is stored with its account, Region, and resource, so week 2 can verify it and the re-scan can show it gone.

Two limits to hold in mind. The CIS benchmark's CloudWatch metric-filter recommendations are manual checks in Security Hub CSPM, so they are not scored. And the scan sees configuration, not behavior: it cannot tell you who has used a weakness, only that the weakness exists.

What the scan does not see

- Application code and the dependencies it runs on (a penetration test or the Amazon Inspector findings in week 5 cover that).
- Anything outside AWS's APIs: Okta, GitHub, the factory networks, and the on-premises systems that connect to AWS. Week 2 reads the AWS side of each of those connections, such as the roles GitHub Actions assumes, and not the systems themselves.
- The data stored in buckets, databases, and queues.
- StrataForge's separate AWS GovCloud (US) organization.

The organization score

32.3% of control results pass across the organization on the Vantalect scan score: 767 passed and 1,604 failed, with each account's controls counted once. The two standards tell the same story from different directions.

Passed and failed control results, pooled across all 48 accounts.

Standard	Passed	Failed	Score
CIS AWS Foundations Benchmark v5.0.0	541	860	38.6%
AWS Foundational Security Best Practices v1.0.0	395	1,181	25.1%
Both, with shared controls counted once	767	1,604	32.3%

CIS scores higher than FSBP (38.6% against 25.1%) because it is narrower: it concentrates on identity, logging, storage, and networking basics, and StrataForge has done some of those. FSBP also checks the detection and scanning services (GuardDuty, Inspector, Config, Security Hub CSPM), where most Regions have nothing turned on; the four Inspector controls, for example, fail in all 48 accounts.

By kind of account

The unweighted average of account scores by kind of account.

Kind of account	Accounts	Average score
Production	22	34.0%
Non-production	12	29.5%
Shared services	5	34.3%
Security	3	43.2%
Sandbox	4	18.6%
Management	1	36.7%
Suspended	1	23.4%

Production accounts average 34.0%, barely above the organization as a whole, and sandboxes average 18.6%. Even the security accounts, the best kind at 43.2%, pass fewer than half of their controls. That pattern says most of what fails is an organizational default (detection off, no account-level guardrails), not the habit of one team.

The lowest-scoring accounts

The 16 lowest-scoring accounts. All 48 accounts are listed in [Appendix: all accounts](#).

RankAccount	Account ID	OU	Kind	CIS	FSBP	Overall	Failed controls
1kestrel-legacy-mgmt (Kestrel)	...9120	Workloads/Kestrel	Production	13.8%	8.0%	12.2%	36
2sandbox-hackathon	...4078	Sandbox	Sandbox	19.2%	12.5%	14.9%	40
3kestrel-data (Kestrel)	...9098	Workloads/Kestrel	Production	16.7%	11.5%	16.7%	35
4sandbox-platform	...4011	Sandbox	Sandbox	26.9%	10.3%	18.2%	36
5kestrel-dev (Kestrel)	...9076	Workloads/Kestrel	Non-production	23.1%	13.6%	18.9%	30
6sandbox-robotics	...4033	Sandbox	Sandbox	26.9%	12.9%	19.6%	37
7crucible-research	...7088	Workloads/Crucible	Non-production	25.8%	14.3%	20.8%	42
8kestrel-prod (Kestrel)	...9031	Workloads/Kestrel	Production	22.2%	15.6%	20.8%	42
9sandbox-ml	...4056	Sandbox	Sandbox	27.6%	14.3%	21.6%	40
10fleet-dev	...6037	Workloads/Fleet/NonProd	Non-production	26.7%	20.0%	22.8%	44
11suspended-legacy-ci	...4099	Suspended	Suspended	23.1%	21.9%	23.4%	36
12kestrel-staging (Kestrel)	...9054	Workloads/Kestrel	Non-production	26.7%	22.7%	25.6%	29
13sf-shared-services	...1188	Infrastructure	Shared services	34.5%	19.4%	26.9%	38
14crucible-sim-prod	...7019	Workloads/Crucible	Production	35.5%	22.2%	29.6%	38
15crucible-dev	...7101	Workloads/Crucible	Non-production	34.6%	21.9%	29.8%	33
16erp-sap-nonprod	...8039	Workloads/Enterprise	Non-production	35.7%	20.7%	30.4%	32

The lowest three are `kestrel-legacy-mgmt`, `sandbox-hackathon`, and `kestrel-data`: a Kestrel management account left over from before the acquisition, a hackathon sandbox, and the Kestrel data account. They have in common that nobody was running them as part of the landing zone. Of the 16 lowest, 5 are Kestrel accounts and 4 are sandboxes. Altogether, 43 of the 48 accounts score under 40%.

Failed controls

19,141 of the 49,156 resource evaluations failed, and 10,399 of those are Critical or High. A control rated Critical by Security Hub CSPM does not mean the finding is Critical for StrataForge: week 2 adjusts every rating for how exploitable it is here and how far it would reach.

Failed resource evaluations by Security Hub CSPM severity. 62 of the 71 controls evaluated failed at least once.

Severity	Controls that failed	Failed resource evaluations	Share of failures
Critical	8	898	4.7%
High	23	9,501	49.6%
Medium	24	7,385	38.6%
Low	7	1,357	7.1%
All	62	19,141	100.0%

The Critical controls

All 8 Critical controls failed. `Config.1` accounts for 810 of the 898 failed Critical evaluations; the other 7 controls account for 88 between them. The small numbers are the ones to read closely.

Reading the table below as a security engineer would, before verification:

- **Config.1** fails in 810 of 816 account-Regions. Some of that is real (Config is not recording) and some is a known false positive: Control Tower’s own recorder uses its own role, so the service-linked-role half of the test fails even where recording is complete. Week 2 separates the two.
- **IAM.6 and IAM.9** say that 47 accounts’ root users lack hardware MFA and 7 have no MFA at all. **IAM.4** says 3 root users have access keys, which are credentials that cannot be limited.
- **EC2.19** finds 21 security groups in 10 accounts that allow the internet to reach a high-risk port such as 22, 3389, or 8888.
- **EC2.1** finds 5 EBS snapshots that any AWS account can restore, and **EC2.182** (snapshot block public access) is off everywhere, so the same exposure can recur.

- **S3.2** finds 4 buckets in 1 account that allow public read, and **RDS.1** finds 1 database snapshot that anyone can restore. Those two are the results we will verify first, because a public bucket is exposed to anyone who finds it.

If verification confirms an exposure that is public and unauthenticated, we tell Marcus Oyelaran the same day, instead of waiting for the readout, so that he can decide whether to act before week 3.

The Critical controls that failed, most failed resources first.


Control	What it checks	Standards	Failed resources	Accounts	Regions
Config.1	AWS Config should be enabled and use the service-linked role for resource recording	CIS, FSBP	810	48	17
IAM.6	Hardware MFA should be enabled for the root user	CIS, FSBP	47	47	1
EC2.19	Security groups should not allow unrestricted access to ports with high risk	FSBP	21	10	1
IAM.9	MFA should be enabled for the root user	CIS	7	7	1
EC2.1	Amazon EBS snapshots should not be publicly restorable	FSBP	5	1	1
S3.2	S3 general purpose buckets should block public read access	FSBP	4	1	1
IAM.4	IAM root user access key should not exist	CIS, FSBP	3	3	1
RDS.1	RDS snapshot should be private	FSBP	1	1	1

The High controls with the most failures

Altogether 23 High controls failed. These are the ten with the most failed resources.

The ten High controls with the most failed resources.

Control	What it checks	Standards	Failed resources	Accounts	Regions
EC2 . 8	EC2 instances should use Instance Metadata Service Version 2 (IMDSv2)	CIS, FSBP	1,631	40	5
EC2 . 182	Block public access settings should be enabled for Amazon EBS snapshots	FSBP	816	48	17
Inspector . 3	Amazon Inspector Lambda code scanning should be enabled	FSBP	816	48	17
Inspector . 4	Amazon Inspector Lambda standard scanning should be enabled	FSBP	816	48	17
Inspector . 1	Amazon Inspector EC2 scanning should be enabled	FSBP	808	48	17
Inspector . 2	Amazon Inspector ECR scanning should be enabled	FSBP	808	48	17
EC2 . 2	VPC default security groups should not allow inbound or outbound traffic	CIS, FSBP	787	48	17
IAM . 28	IAM Access Analyzer external access analyzer should be enabled	CIS	768	48	16
GuardDuty . 1	GuardDuty should be enabled	FSBP	558	48	17
GuardDuty . 6	GuardDuty Lambda Protection should be enabled	FSBP	258	43	6



Failed resources by control, most severe controls first.

Two patterns matter more than any single row. The Inspector controls (Inspector.1 to Inspector.4) and EC2.182 fail in nearly every account-Region, which is one decision not made rather than thousands of separate defects: Amazon Inspector is off and EBS snapshot block public access is not turned on anywhere. And EC2.8 (IMDSv2) fails on 1,631 instances in 40 accounts: it is the weakness that turns a server-side request forgery bug in an application into stolen instance credentials, so its count tells us how much a single application flaw could reach.

The 24 Medium and 7 Low controls that failed carry 8,742 failed evaluations between them. We list them in the raw results we keep for week 2 and do not print them here.

Detection coverage before hardening

The most important thing in this report may be what is not turned on. For each account and Region we checked whether six services that detect or record trouble are running: the CloudTrail organization trail, Amazon GuardDuty, AWS Security Hub CSPM, AWS Config, IAM Access Analyzer, and Amazon Inspector. That is 4,896 yes-or-no cells for the organization; this section summarizes them.

Detection and recording coverage by service before the Hardening. The six workload Regions hold 288 account-Regions and the other 11 hold 528.

Service	Account-Regions covered	Share of 816	In the six workload Regions	In the 11 other Regions	Accounts with none
CloudTrail (organization trail)	816	100.0%	288	528	0
GuardDuty	258	31.6%	258	0	5
Security Hub CSPM	129	15.8%	129	0	5
AWS Config	246	30.1%	246	0	7
IAM Access Analyzer	48	5.9%	48	0	0
Amazon Inspector	8	1.0%	8	0	40

Three observations:

- 1. The organization trail is the one control that is everywhere** (100.0%), so the record of what happened exists. It is a 2019 trail, and week 2 looks at whether it can be trusted: log file validation, encryption, and who can stop it.
- 2. Everything else runs in the six workload Regions only, and unevenly.** GuardDuty covers 258 of the 288 account-Regions in those Regions and 0 of the 528 elsewhere. Security Hub CSPM and IAM Access Analyzer cover far less: Access Analyzer reads only one Region.
- 3. The Kestrel accounts have no GuardDuty at all.** All 5 acquired accounts have it off in every Region; a compromise there would not raise an alert.

By kind of account, as the share of that group's account-Regions covered:

Detection coverage by kind of account, as a share of account-Regions.

Kind of account	Accounts	CloudTrail	GuardDuty	Security Hub CSPM	AWS Config	Access Analyzer	Inspector
Production	22	100.0%	30.5%	15.2%	30.5%	5.9%	2.1%
Non-production	12	100.0%	29.4%	14.7%	29.4%	5.9%	0.0%
Shared services	5	100.0%	35.3%	17.6%	35.3%	5.9%	0.0%
Security	3	100.0%	35.3%	17.6%	35.3%	5.9%	0.0%
Sandbox	4	100.0%	35.3%	17.6%	26.5%	5.9%	0.0%
Management	1	100.0%	35.3%	17.6%	35.3%	5.9%	0.0%
Suspended	1	100.0%	35.3%	17.6%	0.0%	5.9%	0.0%

The full matrix, 48 accounts by 6 services by 17 Regions, is kept in the engagement workspace. Week 5 turns every cell on, in every account and Region, and routes the alerts to your on-call.

What week 2 does

Week 2 turns this scan into something you can act on. Three things change between this report and the next.

- **Verification.** Every Critical and High result is checked by hand, with the account owner where it is needed: is the thing really exposed, can it be reached, and is it there on purpose? We read configuration and, where the test is safe and approved in writing, reproduce the weakness; we do not exploit anything we have not been cleared to.
- **False positives.** A result that verification shows to be wrong or intentional stays in the register, with the reason and the setting that makes the scanner stop reporting it. We expect Config.1 on Control Tower's recorders to be the largest, and a few intentionally open ports (for example, a relay that must accept connections from the internet) to appear.
- **Exploitability and blast radius.** Each real finding starts from its benchmark severity and is raised or lowered for how likely it is to be used against StrataForge and how far the damage would reach, by five written rules. A Medium benchmark result that lets one stolen key sign firmware for the whole robot fleet is not Medium.

On July 24, 2026 you receive deliverables 2 to 7 together: the findings register (with the prescriptive fix for every Critical and High finding), the identity and access review, the SOC 2 compliance map, the governance review, the remediation roadmap, and the sponsor readout.

What we need from you this week

- **Tell us now what is public on purpose.** A bucket, a security group, a snapshot, or a DNS record that is open deliberately saves us a false positive and saves your engineers a question. Elena Vasquez's team can send the list to the shared workspace.
- **Keep the audit role in place** through week 6, and tell us if any account owner is asked to remove it.
- **Hold the office hours.** The account owners for the accounts at the top of [The 16 lowest-scoring accounts](#). All 48 accounts are listed in [appendix-accounts](#) will get the first questions, mostly yes-or-no ones.

Appendix: the controls in the score

The 71 controls behind the Vantalect scan score, with the standards that include each. Severities and titles are Security Hub CSPM's own.

Every control the week-one scan and the week-six re-scan evaluate.

Control	What it checks	Severity	Standards
Account.1	Security contact information should be provided for an AWS account	Medium	CIS
CloudTrail.1	CloudTrail should be enabled and configured with at least one multi-Region trail that includes read and write management events	High	CIS, FSBP
CloudTrail.2	CloudTrail should have encryption at-rest enabled	Medium	CIS, FSBP
CloudTrail.4	CloudTrail log file validation should be enabled	Low	CIS, FSBP
CloudTrail.7	Ensure S3 bucket access logging is enabled on the CloudTrail S3 bucket	Low	CIS
Config.1	AWS Config should be enabled and use the service-linked role for resource recording	Critical	CIS, FSBP
EC2.2	VPC default security groups should not allow inbound or outbound traffic	High	CIS, FSBP
EC2.6	VPC flow logging should be enabled in all VPCs	Medium	CIS, FSBP
EC2.7	EBS default encryption should be enabled	Medium	CIS, FSBP

Control	What it checks	Severity	Standards
EC2.8	EC2 instances should use Instance Metadata Service Version 2 (IMDSv2)	High	CIS, FSBP
EC2.21	Network ACLs should not allow ingress from 0.0.0.0/0 to port 22 or port 3389	Medium	CIS
EC2.53	EC2 security groups should not allow ingress from 0.0.0.0/0 to remote server administration ports	High	CIS
EC2.54	EC2 security groups should not allow ingress from ::/0 to remote server administration ports	High	CIS
EFS.1	Elastic File System should be configured to encrypt file data at-rest using AWS KMS	Medium	CIS
EFS.8	EFS file systems should be encrypted at rest	Medium	CIS
IAM.2	IAM users should not have IAM policies attached	Low	CIS
IAM.3	IAM users' access keys should be rotated every 90 days or less	Medium	CIS, FSBP
IAM.4	IAM root user access key should not exist	Critical	CIS, FSBP
IAM.5	MFA should be enabled for all IAM users that have a console password	Medium	CIS, FSBP
IAM.6	Hardware MFA should be enabled for the root user	Critical	CIS, FSBP
IAM.9	MFA should be enabled for the root user	Critical	CIS
IAM.15	Ensure IAM password policy requires minimum password length of 14 or greater	Medium	CIS
IAM.16	Ensure IAM password policy prevents password reuse	Low	CIS
IAM.18	Ensure a support role has been created to manage incidents with AWS Support	Low	CIS
IAM.22	IAM user credentials unused for 45 days should be removed	Medium	CIS
IAM.26	Expired SSL/TLS certificates managed in IAM should be removed	Medium	CIS
IAM.27	IAM identities should not have the AWSCloudShellFullAccess policy attached	Medium	CIS
IAM.28	IAM Access Analyzer external access analyzer should be enabled	High	CIS

Control	What it checks	Severity	Standards
KMS.4	AWS KMS key rotation should be enabled	Medium	CIS
RDS.2	RDS DB Instances should prohibit public access, as determined by the PubliclyAccessible configuration	Critical	CIS, FSBP
RDS.3	RDS DB instances should have encryption at-rest enabled	Medium	CIS, FSBP
RDS.5	RDS DB instances should be configured with multiple Availability Zones	Medium	CIS
RDS.13	RDS automatic minor version upgrades should be enabled	High	CIS
RDS.15	RDS DB clusters should be configured for multiple Availability Zones	Medium	CIS
S3.1	S3 general purpose buckets should have block public access settings enabled	Medium	CIS, FSBP
S3.5	S3 general purpose buckets should require requests to use TLS	Medium	CIS, FSBP
S3.8	S3 general purpose buckets should block public access	High	CIS, FSBP
S3.20	S3 general purpose buckets should have MFA delete enabled	Low	CIS
S3.22	S3 general purpose buckets should log object-level write events	Medium	CIS
S3.23	S3 general purpose buckets should log object-level read events	Medium	CIS
EC2.18	Security groups should only allow unrestricted incoming traffic for authorized ports	High	FSBP
EC2.19	Security groups should not allow unrestricted access to ports with high risk	Critical	FSBP
EC2.1	Amazon EBS snapshots should not be publicly restorable	Critical	FSBP
EC2.182	Block public access settings should be enabled for Amazon EBS snapshots	High	FSBP
S3.2	S3 general purpose buckets should block public read access	Critical	FSBP
S3.9	S3 general purpose buckets should have server access logging enabled	Medium	FSBP
RDS.1	RDS snapshot should be private	Critical	FSBP

Control	What it checks	Severity	Standards
IAM.1	IAM policies should not allow full "*" administrative privileges	High	FSBP
IAM.21	IAM customer managed policies that you create should not allow wildcard actions for services	Low	FSBP
Lambda.2	Lambda functions should use supported runtimes	Medium	FSBP
EKS.1	EKS cluster endpoints should not be publicly accessible	High	FSBP
EKS.2	EKS clusters should run on a supported Kubernetes version	High	FSBP
ECR.1	ECR private repositories should have image scanning configured	High	FSBP
GuardDuty.1	GuardDuty should be enabled	High	FSBP
Opensearch.2	OpenSearch domains should not be publicly accessible	Critical	FSBP
ES.2	Elasticsearch domains should not be publicly accessible	Critical	FSBP
ELB.1	Application Load Balancer should be configured to redirect all HTTP requests to HTTPS	Medium	FSBP
SSM.1	Amazon EC2 instances should be managed by AWS Systems Manager	Medium	FSBP
GuardDuty.5	GuardDuty EKS Audit Log Monitoring should be enabled	High	FSBP
GuardDuty.6	GuardDuty Lambda Protection should be enabled	High	FSBP
GuardDuty.7	GuardDuty EKS Runtime Monitoring should be enabled	High	FSBP
GuardDuty.8	GuardDuty Malware Protection for EC2 should be enabled	High	FSBP
GuardDuty.9	GuardDuty RDS Protection should be enabled	High	FSBP
GuardDuty.10	GuardDuty S3 Protection should be enabled	High	FSBP
GuardDuty.11	GuardDuty Runtime Monitoring should be enabled	High	FSBP
GuardDuty.12	GuardDuty ECS Runtime Monitoring should be enabled	Medium	FSBP

Control	What it checks	Severity	Standards
GuardDuty.13	GuardDuty EC2 Runtime Monitoring should be enabled	Medium	FSBP
Inspector.1	Amazon Inspector EC2 scanning should be enabled	High	FSBP
Inspector.2	Amazon Inspector ECR scanning should be enabled	High	FSBP
Inspector.3	Amazon Inspector Lambda code scanning should be enabled	High	FSBP
Inspector.4	Amazon Inspector Lambda standard scanning should be enabled	High	FSBP

Appendix: all accounts

Vantalect scan score by account, lowest first. Account IDs are masked to their last four digits.

RankAccount	Account ID	OU	Kind	Business unit	CIS	FSBP	Overall	Failed controls
1kestrel-legacy-mgmt	...9120	Workloads/Kestrel	Production	Kestrel Vision	13.8%	8.0%	12.2%	36
2sandbox-hackathon	...4078	Sandbox	Sandbox	Platform	19.2%	12.5%	14.9%	40
3kestrel-data	...9098	Workloads/Kestrel	Production	Kestrel Vision	16.7%	11.5%	16.7%	35
4sandbox-platform	...4011	Sandbox	Sandbox	Platform	26.9%	10.3%	18.2%	36
5kestrel-dev	...9076	Workloads/Kestrel	Non-production	Kestrel Vision	23.1%	13.6%	18.9%	30
6sandbox-robotics	...4033	Sandbox	Sandbox	Fleet Cloud	26.9%	12.9%	19.6%	37
7crucible-research	...7088	Workloads/Crucible	Non-production	Crucible	25.8%	14.3%	20.8%	42
8kestrel-prod	...9031	Workloads/Kestrel	Production	Kestrel Vision	22.2%	15.6%	20.8%	42
9sandbox-ml	...4056	Sandbox	Sandbox	Crucible	27.6%	14.3%	21.6%	40
10fleet-dev	...6037	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	26.7%	20.0%	22.8%	44

Rank	Account	Account ID	OU	Kind	Business unit	CIS	FSBP	Overall	Failed controls
11	suspended-legacy-ci	...4099	Suspended	Suspended	Platform	23.1%	21.9%	23.4%	36
12	kestrel-staging	...9054	Workloads/Kestrel	Non-production	Kestrel Vision	26.7%	22.7%	25.6%	29
13	sf-shared-services	...1188	Infrastructure	Shared services	Platform	34.5%	19.4%	26.9%	38
14	crucible-sim-prod	...7019	Workloads/Crucible	Production	Crucible	35.5%	22.2%	29.6%	38
15	crucible-dev	...7101	Workloads/Crucible	Non-production	Crucible	34.6%	21.9%	29.8%	33
16	erp-sap-nonprod	...8039	Workloads/Enterprise	Non-production	Enterprise IT	35.7%	20.7%	30.4%	32
17	customer-portal-dev	...6103	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	33.3%	25.0%	30.6%	34
18	digital-twin-dev	...6126	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	34.5%	24.2%	30.6%	34
19	fleet-remote-ops-prod	...5139	Workloads/Fleet/Prod	Production	Fleet Cloud	42.3%	20.7%	31.8%	30
20	sf-cicd	...1356	Infrastructure	Shared services	Platform	38.7%	24.3%	32.7%	37
21	fleet-analytics-prod	...5093	Workloads/Fleet/Prod	Production	Fleet Cloud	35.7%	25.0%	33.3%	32
22	fleet-eu-staging	...6081	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	40.0%	28.1%	34.7%	32
23	crucible-data	...7065	Workloads/Crucible	Production	Crucible	41.4%	28.6%	35.3%	33
24	mfg-mes-prod	...8084	Workloads/Enterprise	Production	Enterprise IT	42.3%	27.3%	35.4%	31
25	customer-portal-prod	...5215	Workloads/Fleet/Prod	Production	Fleet Cloud	44.1%	27.0%	35.7%	36
26	fleet-staging	...6014	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	40.0%	30.8%	35.7%	36
27	corp-data-warehouse	...8107	Workloads/Enterprise	Production	Enterprise IT	42.9%	29.0%	36.2%	30
28	sf-network-hub	...1145	Infrastructure	Shared services	Platform	50.0%	24.1%	36.4%	28
29	fleet-data-prod	...5071	Workloads/Fleet/Prod	Production	Fleet Cloud	41.4%	30.6%	36.5%	33
30	crucible-training-prod	...7042	Workloads/Crucible	Production	Crucible	45.2%	26.5%	36.5%	33

RankAccount	Account ID	OU	Kind	Business unit	CIS	FSBP	Overall	Failed controls
31corp-apps	...8129	Workloads/Enterprise	Production	Enterprise IT	41.9%	29.4%	36.5%	33
32sf-backup	...1407	Infrastructure	Shared services	Platform	50.0%	25.9%	36.6%	26
33sf-management	...0007	Root	Management	Platform	46.7%	27.3%	36.7%	31
34fleet-ota-staging	...6148	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	44.8%	27.3%	36.7%	31
35fleet-perf	...6059	Workloads/Fleet/NonProd	Non-production	Fleet Cloud	43.3%	32.4%	37.3%	32
36plm-prod	...8061	Workloads/Enterprise	Production	Enterprise IT	46.9%	25.0%	37.3%	32
37fleet-eu-prod	...5160	Workloads/Fleet/Prod	Production	Fleet Cloud	42.4%	30.0%	37.9%	36
38fleet-ingest-prod	...5012	Workloads/Fleet/Prod	Production	Fleet Cloud	46.2%	32.4%	38.5%	32
39fleet-jp-prod	...5182	Workloads/Fleet/Prod	Production	Fleet Cloud	42.4%	33.3%	38.6%	35
40erp-sap-prod	...8016	Workloads/Enterprise	Production	Enterprise IT	48.4%	30.6%	38.9%	33
41sf-identity	...1202	Infrastructure	Shared services	Platform	54.2%	29.6%	39.0%	25
42digital-twin-prod	...5238	Workloads/Fleet/Prod	Production	Fleet Cloud	45.5%	31.6%	39.3%	34
43fleet-api-prod	...5044	Workloads/Fleet/Prod	Production	Fleet Cloud	48.4%	32.5%	39.7%	35
44fleet-ota-prod	...5117	Workloads/Fleet/Prod	Production	Fleet Cloud	48.5%	36.1%	40.7%	32
45crucible-models	...7124	Workloads/Crucible	Production	Crucible	50.0%	31.4%	41.5%	31
46sf-forensics	...3618	Security	Security	Security	53.8%	28.6%	41.9%	25
47sf-security-tooling	...3374	Security	Security	Security	51.7%	33.3%	42.9%	28
48sf-log-archive	...2291	Security	Security	Security	57.1%	32.3%	44.7%	26

© 2026 Vantalect, LLC. Sample deliverable; StrataForge Robotics is fictional.

AWS, Amazon Web Services, AWS Graviton, and the AWS service names are trademarks of Amazon.com, Inc. or its affiliates. Kubernetes is a registered trademark of The Linux Foundation. GitHub is a trademark of GitHub, Inc. NVIDIA and Isaac Sim are trademarks of NVIDIA Corporation. SAP and SAP S/4HANA are trademarks of SAP SE. Okta is a trademark of Okta, Inc. PagerDuty is a trademark of PagerDuty, Inc. SOC 2 is a registered trademark of the American Institute of Certified Public Accountants. CIS Benchmarks is a trademark of the Center for Internet Security, Inc. Vantalect is not affiliated with or endorsed by any of them. StrataForge Robotics is fictional.