

AWS SECURITY, COMPLIANCE & GOVERNANCE · SAMPLE DELIVERABLE

Week-Six Re-Scan

Week 6 of 6: Re-scan · prepared for StrataForge Robotics, Inc.

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

August 21, 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

Summary	3
How the re-scan was run	4
Score before and after	6
By standard	7
By account	8
Findings, before and after	11
Every critical, with its closure evidence	13
SF-SEC-001: A 1,127-day-old IAM user key can sign and publish firmware to every robot	14
SF-SEC-002: One IoT policy lets any of 61,400 robots' certificates publish to any robot's command topic	14
SF-SEC-003: 4 GitHub Actions deploy roles can be assumed from any GitHub repository	14
SF-SEC-004: Jupyter is open to the internet on 51 research and sandbox instances, some with no token	14
SF-SEC-005: 5 EBS snapshots in kestrel-prod are public, and snapshot block public access is off in all 816 account-Regions	14
SF-SEC-006: 4 Kestrel dataset buckets are public: 468.8 million objects of labeled images from customer sites	15
SF-SEC-007: A Kestrel MySQL snapshot with annotation records is restorable by any AWS account	15
SF-SEC-008: Kestrel's Jenkins controller is open to the internet on 8080 and 22	15
SF-SEC-009: RDP is open to the internet on 32 PLM CAD workstations that hold product designs	15
SF-SEC-010: Root users of 3 Kestrel accounts have access keys and no MFA	16
The critical accepted in writing	16
The remaining High findings	18
The SOC 2 map, after	20
The gate, and what follows	25
What to do next	26

Summary

The guarantee is met. Every one of the 10 critical findings in the week-two register is closed or accepted in writing: 9 closed and verified by the re-scan, and 1 accepted by Marcus Oyelaran and Priya Raman on August 19, 2026, with controls and an expiry. The gate decided on August 21, 2026 is **PASS**, and the second invoice is issued. [The gate, and what follows](#) shows the computation and what follows from it.

We ran the same scan as week one, the same way, against the same 48 accounts and 17 Regions. **The organization’s Vantalect scan score (71 controls) rose from 32.4% to 91.1%** across both standards. Without the three suppressions the work earned (below), it is 86.3%. The two standards separately: from 38.6% to 89.4% on the CIS AWS Foundations Benchmark v5.0.0 controls, and from 25.1% to 92.9% on the AWS Foundational Security Best Practices controls. Failed controls, counted once per account where a control fails, fell from 1,603 to 211. 2 of the 48 accounts score 100%.

The 54 findings in the register are in four states: **34 closed, 2 accepted in writing, 8 partially remediated, and 10 open**. 5 High findings remain, all partially remediated, and the worst open severity is High. None is critical. Each remaining finding has an owner and a date in the handover report, and none depends on Vantalect. The 48 pull requests that did the work are in the weekly reports for weeks 3 to 5.

The week-one scan and the week-six re-scan, side by side. Open means not closed and not accepted in writing.

Week 6 at a glance	Week one	Week six
Vantalect scan score, both standards	32.4%	91.1%
The same score without the new suppressions	32.4%	86.3%
CIS AWS Foundations Benchmark v5.0.0	38.6%	89.4%
AWS Foundational Security Best Practices v1.0.0	25.1%	92.9%
Failed controls, both standards	1,603	211
Failed resource evaluations	18,894	3,224
Critical findings still open	10	0
High findings still open	20	5
Medium findings still open	15	9

Week 6 at a glance	Week one	Week six
Low findings still open	9	4
Accounts at 100%	0	2

The guarantee

The second half of the fee is invoiced only when every critical finding in the audit is closed, verified by the week-six re-scan; a critical accepted in writing counts as closed. The re-scan found no critical still failing. The gate passed on August 21, 2026. Beyond that gate, no particular result is guaranteed: this is a security engagement, not an audit opinion, an attestation, a certification of compliance, or a penetration test.

How the re-scan was run

The re-scan is the week-one scan repeated, with the same evaluator, the same 71 controls from the same 2 standards, the same accounts, and the same Regions, run in week 6 (the week of August 17, 2026). We changed nothing about how it scores. A scan that is easier the second time is not a result.

Four things need saying, because they decide whether the comparison is honest.


- **The baseline is the week-one scan after the agreed false positives were suppressed.** At the week-two readout we and the sponsor agreed on 6 false positives: conditions that fail a check for a reason that is not a risk, such as the Control Tower recorder failing the Config.1 service-linked-role test. 2 of the 6 fail no Security Hub CSPM check at all, so they move no score. The unsuppressed week-one score of 32.3% and the baseline of 32.4% differ by only 0.1 point, and CIS does not move at all, which is why we say so here rather than hide it.
- **The agreed fixes are applied the way they were agreed.** Two are control parameters (Config.1 and EC2.18) and two are documented suppressions (IAM.22 and KMS.4). The re-scan also suppresses three controls that the work itself made not applicable: IAM.6 and IAM.9 where root credentials no longer exist (a root user with no password and no MFA cannot fail an MFA check), and S3.5 where the TLS resource control policy now denies plain-HTTP requests. S3.5 is not suppressed in sf-management, which resource control policies do not cover, so it still fails there. These new suppressions move the organization score by 4.8 points, so we publish both figures: 91.1% with them and 86.3% without. Suppressed evaluations rose from 247 to 315, and every suppression has its reason on the finding.

- **A finding counts as closed only when its pull requests have merged and none of its original failed checks still fails.** Not when we say so, and not when the owner says so. The per-critical evidence is in [Every critical, with its closure evidence](#).
- **The score is a count of controls, not of risk, and of 71 controls, not of every control Security Hub CSPM has.** A control passes or fails for an account as a whole, and one failing resource fails it. An account at 100% passes every one of the 71 controls we evaluate in it; it is not an account with no risk, and it may still fail a control outside the list. Security Hub CSPM's own score for the organization would differ. That is why the register, not the score, carries the guarantee.

How the Vantalect scan score works

For each account and standard the score is the controls that pass divided by the controls that pass or fail; a control with no data is left out. The organization score pools all 48 accounts: every account's passed and failed counts are summed and then divided. A "failed control" in this report is one control failing in one account.

Score before and after



The 14 lowest-scoring accounts at week one, with the score they started from and the gain by week six.

By standard

The organization score by standard, before and after, and at week six without the new suppressions, with the number of controls failing in at least one account's evaluation.

Standard	Score, week one	Score, week six	Week six without new suppressions	Failed controls, week one	Failed controls, week six
CIS AWS Foundations Benchmark v5.0.0	38.6%	89.4%	81.0%	860	144
AWS Foundational Security Best Practices v1.0.0	25.1%	92.9%	88.6%	1,180	114
Both standards	32.4%	91.1%	86.3%	1,603	211

The suppressions the work earned, which are in the week-six score and not in the week-one baseline.

Control	Accounts	ChecksReason suppressed
IAM.6	47	47No root credentials exist (centralized root access management); root tasks run through sts:AssumeRoot
IAM.9	47	47No root credentials exist (centralized root access management); root tasks run through sts:AssumeRoot
S3.5	21	217Compensating control: resource control policy sf-rcp-enforce- tls denies requests without TLS

What the 211 remaining failed controls are: they trace to the 18 findings still open, and the largest groups are the instances that still allow IMDSv1 (SF-SEC-014), the IAM users whose keys are older than 90 days (SF-SEC-012), the EKS clusters whose endpoints are public (SF-SEC-016), the unencrypted volumes and databases (SF-SEC-032, SF-SEC-035), and the default network ACLs (SF-SEC-043). The remaining-highs table in [The remaining High findings](#) gives each its count of still-failing checks.

By account


The score of every account under each standard and both together, before and after, with failed controls across both.

Account	OU	CIS before	CIS after	FSBP before	FSBP after	Overall before	Overall after	Failed before	Failed after
fleet-dev	Workloads/Fleet/NonProd	26.7%	83.3%	20.0%	82.5%	22.8%	82.5%	44	10
kestrel-prod	Workloads/Kestrel	22.2%	81.2%	15.6%	87.2%	20.8%	84.5%	42	9
fleet-staging	Workloads/Fleet/NonProd	40.0%	86.7%	30.8%	84.6%	35.7%	85.7%	36	8
fleet-api-prod	Workloads/Fleet/Prod	48.4%	87.1%	32.5%	85.0%	39.7%	86.2%	35	8
sf-cicd	Infrastructure	38.7%	81.5%	24.3%	91.4%	32.7%	86.3%	37	7
crucible-research	Workloads/Crucible	25.8%	83.9%	14.3%	91.4%	20.8%	86.8%	42	7
crucible-sim-prod	Workloads/Crucible	35.5%	83.9%	22.2%	91.7%	29.6%	87.0%	38	7
customer-portal-prod	Workloads/Fleet/Prod	44.1%	88.2%	27.0%	89.2%	35.7%	87.5%	36	7
mfg-mes-prod	Workloads/Enterprise	42.3%	80.8%	27.3%	90.9%	35.4%	87.5%	31	6
fleet-eu-prod	Workloads/Fleet/Prod	42.4%	84.8%	30.0%	90.0%	37.9%	87.9%	36	7
fleet-ota-prod	Workloads/Fleet/Prod	48.5%	84.8%	36.1%	91.7%	40.7%	88.9%	32	6
erp-sap-nonprod	Workloads/Enterprise	35.7%	85.7%	20.7%	93.1%	30.4%	89.1%	32	5
kestrel-staging	Workloads/Kestrel	26.7%	86.7%	22.7%	90.3%	25.6%	89.6%	29	5
customer-portal-dev	Workloads/Fleet/NonProd	33.3%	86.7%	25.0%	93.8%	30.6%	89.8%	34	5
fleet-eu-staging	Workloads/Fleet/NonProd	40.0%	86.7%	28.1%	90.6%	34.7%	89.8%	32	5
corp-apps	Workloads/Enterprise	41.9%	90.0%	29.4%	91.2%	36.5%	90.2%	33	5
crucible-data	Workloads/Crucible	41.4%	89.7%	28.6%	91.4%	35.3%	90.2%	33	5
fleet-perf	Workloads/Fleet/NonProd	43.3%	90.0%	32.4%	91.2%	37.3%	90.2%	32	5

Account	OU	CIS before	CIS after	FSBP before	FSBP after	Overall before	Overall after	Failed before	Failed after
sandbox-m1	Sandbox	27.6%	89.7%	14.3%	91.4%	21.6%	90.2%	40	5
fleet-data-prod	Workloads/Fleet/Prod	41.4%	89.7%	30.6%	91.7%	36.5%	90.4%	33	5
fleet-ingest-prod	Workloads/Fleet/Prod	46.2%	92.3%	32.4%	89.2%	38.5%	90.4%	32	5
erp-sap-prod	Workloads/Enterprise	48.4%	90.3%	30.6%	91.7%	38.9%	90.7%	33	5
digital-twin-prod	Workloads/Fleet/Prod	45.5%	87.9%	31.6%	94.7%	39.3%	91.1%	34	5
fleet-jp-prod	Workloads/Fleet/Prod	42.4%	87.9%	33.3%	94.9%	38.6%	91.2%	35	5
kestrel-dev	Workloads/Kestrel	23.1%	92.3%	13.6%	90.3%	18.9%	91.3%	30	4
crucible-dev	Workloads/Crucible	34.6%	92.3%	21.9%	90.6%	29.8%	91.5%	33	4
digital-twin-dev	Workloads/Fleet/NonProd	34.5%	89.7%	24.2%	93.9%	30.6%	91.8%	34	4
fleet-ota-staging	Workloads/Fleet/NonProd	44.8%	89.7%	27.3%	93.9%	36.7%	91.8%	31	4
sf-shared-services	Infrastructure	34.5%	89.7%	19.4%	94.4%	26.9%	92.3%	38	4
sandbox-platform	Sandbox	26.9%	90.9%	10.3%	92.6%	18.2%	92.5%	36	3
suspended-legacy-ci	Suspended	23.1%	90.9%	21.9%	93.3%	23.4%	93.0%	36	3
fleet-remote-ops-prod	Workloads/Fleet/Prod	42.3%	92.3%	24.1%	93.1%	34.1%	93.2%	29	3
sf-network-hub	Infrastructure	50.0%	92.3%	24.1%	93.1%	36.4%	93.2%	28	3
sandbox-robotics	Sandbox	26.9%	92.3%	12.9%	93.5%	19.6%	93.5%	37	3
kestrel-legacy-mgmt	Workloads/Kestrel	13.8%	92.3%	8.0%	93.8%	12.2%	93.6%	36	3
sandbox-hackathon	Sandbox	19.2%	96.2%	12.5%	93.8%	14.9%	93.6%	40	3
kestrel-data	Workloads/Kestrel	16.7%	92.6%	11.5%	97.0%	16.7%	93.8%	35	3
plm-prod	Workloads/Enterprise	46.9%	90.6%	25.0%	96.9%	37.3%	94.1%	32	3

Account	OU	CIS before	CIS after	FSBP before	FSBP after	Overall before	Overall after	Failed before	Failed after
crucible-training-prod	Workloads/Crucible	45.2%	93.5%	26.5%	94.1%	36.5%	94.2%	33	3
crucible-models	Workloads/Crucible	50.0%	90.6%	31.4%	97.1%	41.5%	94.3%	31	3
corp-data-warehouse	Workloads/Enterprise	42.9%	92.9%	29.0%	100.0%	36.2%	95.7%	30	2
sf-log-archive	Security	57.1%	92.9%	32.3%	100.0%	44.7%	95.7%	26	2
fleet-analytics-prod	Workloads/Fleet/Prod	35.7%	92.9%	25.0%	100.0%	33.3%	95.8%	32	2
sf-management	Root	46.7%	93.3%	27.3%	100.0%	36.7%	95.9%	31	2
sf-security-tooling	Security	51.7%	93.1%	33.3%	100.0%	42.9%	95.9%	28	2
sf-forensics	Security	53.8%	95.5%	28.6%	100.0%	41.9%	97.4%	25	1
sf-backup	Infrastructure	50.0%	100.0%	25.9%	100.0%	36.6%	100.0%	26	0
sf-identity	Infrastructure	54.2%	100.0%	29.6%	100.0%	39.0%	100.0%	25	0

Findings, before and after



The register's findings by severity, with how many are still open at week six.

The 54 findings by severity and by where each ended.

Severity	In the register	Closed	Accepted in writing	Still open
Critical	10	9	1	0
High	20	15	0	5
Medium	15	6	0	9
Low	9	4	1	4
Total	54	34	2	18

The four states a finding can be in at week six.

Status	FindingsWhat it means
Closed	34The pull requests merged and none of the original failed checks still fails
Accepted in writing	2A named executive accepted the residual risk, for a stated time, with compensating controls
Partially remediated	8Work merged; a remainder waits on a window, an upgrade, or a customer, and has an owner and a date
Open	10Not started in the engagement; scheduled in the roadmap with an owner and a date

All 54 findings in the register are accounted for. 0 findings disagree between the register and the re-scan, which is the check we ran last: a finding marked closed whose checks still fail, or marked open whose checks all pass, would appear in that count.

Every critical, with its closure evidence

The 10 critical findings from the week-two register, each with the pull requests that fixed it, the Security Hub CSPM controls we re-ran, and the evidence that it is closed. The re-scan counts a control as re-run when it failed in week one on the resources the finding named. "Now passing" are the ones that failed and no longer do; "removed" are resources that no longer exist.

The closure of each critical finding, as the re-scan measured it. A finding with no automated control is verified by the command in its evidence.

Finding	Status	Closed	Pull requests	Controls re-run	Checks	Now passing	Still failing	Suppressed
SF-SEC-001	closed	Jul 28	fleet-infra#4323, fleet-infra#4328	IAM.2, IAM.3	2	2	0	0
SF-SEC-002	accepted	accepted	fleet-infra#4329, fleet-infra#4332, fleet-infra#4336	no automated control	0	0	0	0
SF-SEC-003	closed	Jul 27	fleet-infra#4340	no automated control	0	0	0	0
SF-SEC-004	closed	Jul 28	crucible-infra#1987	EC2.18, EC2.19, EC2.53, EC2.54	19	19	0	0
SF-SEC-005	closed	Jul 27	platform-landing-zone#613	EC2.1, EC2.182	821	821	0	0
SF-SEC-006	closed	Jul 28	crucible-infra#1992	S3.1, S3.2, S3.8	8	8	0	0
SF-SEC-007	closed	Jul 27	crucible-infra#1982	RDS.1	1	1	0	0
SF-SEC-008	closed	Jul 28	crucible-infra#1996	EC2.18, EC2.19, EC2.53	3	3	0	0
SF-SEC-009	closed	Jul 28	enterprise-infra#905	EC2.18, EC2.19, EC2.53	9	9	0	0
SF-SEC-010	closed	Jul 28	platform-landing-zone#619	IAM.4, IAM.6, IAM.9	9	3	0	6

SF-SEC-001: A 1,127-day-old IAM user key can sign and publish firmware to every robot

Scope: 1 IAM user.

Closed July 28, 2026. Pull requests: fleet-infra#4323, fleet-infra#4328. The re-scan re-ran 2 checks across IAM.2, IAM.3: 2 now pass. **Evidence.** The access key of IAM user ota-signing-ci is deactivated (the key and user are deleted after the next clean release); the key policy allows kms:Sign only to ota-release-signer through the production GitHub environment with two approvals; a Sign by any other principal pages fleet on-call.

SF-SEC-002: One IoT policy lets any of 61,400 robots' certificates publish to any robot's command topic

Scope: 61,400 robot certificates.

Accepted in writing August 19, 2026; the full text is in [The critical accepted in writing](#). Pull requests: fleet-infra#4329, fleet-infra#4332, fleet-infra#4336. There is no Security Hub CSPM control for this finding; its evidence is Device Defender's audit and the thing-group counts in the weekly robot report.

SF-SEC-003: 4 GitHub Actions deploy roles can be assumed from any GitHub repository

Scope: 4 IAM roles.

Closed July 27, 2026. Pull requests: fleet-infra#4340. **Evidence.** All four trust policies require sub = repo:strataforge/<repository>:environment:production; an assume attempt from Vantalect's test repository now fails with AccessDenied.

SF-SEC-004: Jupyter is open to the internet on 51 research and sandbox instances, some with no token

Scope: 51 instances.

Closed July 28, 2026. Pull requests: crucible-infra#1987. The re-scan re-ran 19 checks across EC2.18, EC2.19, EC2.53, EC2.54: 19 now pass. **Evidence.** No security group admits 8888 or 22 from the internet; notebooks open through Session Manager port forwarding.

SF-SEC-005: 5 EBS snapshots in kestrel-prod are public, and snapshot block public access is off in all 816 account-Regions

Scope: 5 EBS snapshots.

Closed July 27, 2026. Pull requests: platform-landing-zone#613. The re-scan re-ran 821 checks across EC2.1, EC2.182: 821 now pass. **Evidence.** No public snapshots; the EC2 declarative policy sets snapshot block public access to block all sharing in every account and Region.

SF-SEC-006: 4 Kestrel dataset buckets are public: 468.8 million objects of labeled images from customer sites

Scope: 4 S3 buckets.

Closed July 28, 2026. Pull requests: crucible-infra#1992. The re-scan re-ran 8 checks across S3.1, S3.2, S3.8: 8 now pass. **Evidence.** Buckets imported into CDK, private, with bucket and account Block Public Access on; S3.2 and S3.8 pass.

SF-SEC-007: A Kestrel MySQL snapshot with annotation records is restorable by any AWS account

Scope: 1 DB snapshot.

Closed July 27, 2026. Pull requests: crucible-infra#1982. The re-scan re-ran 1 check across RDS.1: 1 now passes. **Evidence.** describe-db-snapshot-attributes shows no restore attribute value; RDS.1 passes.

SF-SEC-008: Kestrel's Jenkins controller is open to the internet on 8080 and 22

Scope: 1 instance.

Closed July 28, 2026. Pull requests: crucible-infra#1996. The re-scan re-ran 3 checks across EC2.18, EC2.19, EC2.53: 3 now pass. **Evidence.** Jenkins answers only through an internal ALB with Okta sign-in; every stored AWS key was rotated or retired.

SF-SEC-009: RDP is open to the internet on 32 PLM CAD workstations that hold product designs

Scope: 3 security groups (the finding counts security groups; its title counts the workstations attached to them).

Closed July 28, 2026. Pull requests: enterprise-infra#905. The re-scan re-ran 9 checks across EC2.18, EC2.19, EC2.53: 9 now pass. **Evidence.** No security group admits 22 or 3389 from 0.0.0.0/0 or ::/0; access through Session Manager and Fleet Manager.

SF-SEC-010: Root users of 3 Kestrel accounts have access keys and no MFA

Scope: 3 root users.

Closed July 28, 2026. Pull requests: platform-landing-zone#619. The re-scan re-ran 9 checks across IAM.4, IAM.6, IAM.9: 3 now pass, and 6 are suppressed with a written reason (the root user no longer has credentials, so the MFA checks cannot apply). **Evidence.** Centralized root access management enabled; member root credentials deleted with sts:AssumeRoot and the IAMDeleteRootUserCredentials task policy; get-account-summary shows AccountAccessKeysPresent 0.

The critical accepted in writing

Risk acceptance for SF-SEC-002

Accepted by: Marcus Oyelaran, Chief Information Security Officer, and Priya Raman, Chief Technology Officer, on August 19, 2026.

Expires: December 15, 2026. After that date the finding is open again unless the signers renew it in writing.

Scope: 3,412 ForgeArm gen-2 robots on firmware 3.x, at customer sites, that cannot use the v2 topic layout until firmware 4.2 is installed in each customer's maintenance window.

Statement: StrataForge accepts, until December 15, 2026, the residual risk that a certificate taken from one of the 3,412 ForgeArm gen-2 robots still on firmware 3.x can publish and subscribe to that one robot's own telemetry, shadow, and jobs topics in the legacy sf/v1 layout. It cannot reach another robot's topics, and no robot certificate can reach a command topic. Firmware 4.2 will be installed in each customer's scheduled maintenance window, after which the robot moves to sf-robot-scoped-v2 and the legacy policy is deleted.

Compensating controls.

- sf-robot-legacy-v1-restricted lets each robot use only its own sf/v1/<thing> topics, its shadow, and its jobs (the jobs deliver the firmware update), and denies every cloud command topic to robot certificates.

- Device Defender Detect profile on the legacy thing group alarms on authorization failures and message-volume anomalies to PagerDuty.
- Weekly report to the CISO of robots remaining on firmware 3.x, by customer site.
- Robots on firmware 3.x accept commands only from the site's local controller, which authenticates over a separate channel.

The finding is One IoT policy lets any of 61,400 robots' certificates publish to any robot's command topic. The acceptance is the signed copy of the draft that week 3 carried: the same text, the same four controls, and the same expiry. What changed is that 3,412 robots are still on firmware 3.x, and 1,240 pilot robots that trusted the shared staging test key were fixed separately in week 3.

Under the guarantee, a critical accepted in writing counts as closed. We think that is right for this finding. It is a known, bounded risk with controls that work and an end date, held by the two people who can carry it, and the alternative was to update 3,412 robots at customer sites in a week, which would have been the larger risk. It is not a reason to accept any other critical.

A second finding was accepted in writing, a Low: SF-SEC-049, MFA delete, accepted by Marcus Oyelaran on August 14, 2026 until August 14, 2027, with Object Lock and versioning as the controls.

The remaining High findings

5 High findings are partially remediated. Each moved from a worse state this engagement; each has a remainder that cannot finish inside six weeks for a reason outside the code. They are not critical, and none is a surprise to its owner.

The 5 High findings still open at week six, with the owner, the date, and the count of Security Hub CSPM checks that still fail because of each.

ID	Finding	What remains	Owner	Due	Checks still failing
SF-SEC-015	144 Lambda functions have secret-named environment variables	Move the fleet-dev functions' secrets to Secrets Manager	Sam Okafor	September 30, 2026	0
SF-SEC-016	All 16 EKS cluster API endpoints accept connections from 0.0.0.0/0	Make the non-production endpoints private once CI runners reach them through the Transit Gateway	Sam Okafor	September 30, 2026	4
SF-SEC-030	23 robot certificates are shared by 1,840 robots at 3 customer sites	Re-provision the robots at the Bremen and Ota sites during their maintenance windows	Sam Okafor	September 30, 2026	0
SF-SEC-014	1,631 instances still allow IMDSv1	Switch the instances that still make IMDSv1 calls after their agents and SDKs are upgraded	Elena Vasquez	October 30, 2026	358
SF-SEC-012	1,009 active access keys are older than 90 days; 400 are over three years old	Replace the factory gateway users' keys with IAM Roles Anywhere during factory maintenance windows	Grace Liu	November 20, 2026	112

Why each is not finished, plainly:

- **SF-SEC-015.** Lambda functions keep DATABASE_URL, PARTNER_API_KEY, GITHUB_TOKEN, DB_PASSWORD, or SAP_RFC_PASSWORD in plain environment variables, readable by anyone with lambda:GetFunctionConfiguration. Store them in Secrets Manager, grant read to the function role, and rotate every value that was exposed. Remaining: Move the fleet-dev functions' secrets to Secrets Manager. Sam Okafor, September 30, 2026.
- **SF-SEC-016.** EKS clusters expose a public API endpoint open to 0.0.0.0/0 (EKS.1). Turn on private endpoint access, limit the public endpoint to the corporate egress range, and move to private-only once CI runs inside the VPC. Remaining: Make the non-production endpoints private once CI runners reach them through the Transit Gateway. Sam Okafor, September 30, 2026.
- **SF-SEC-030.** Some certificates are shared by many robots (DEVICE_CERTIFICATE_SHARED_CHECK and CONFLICTING_CLIENT_IDS_CHECK), so one robot's connection kicks another off and revoking a stolen certificate would take all of them offline. Give each robot its own certificate through fleet provisioning by claim, with a pre-provisioning hook that lets only listed serial numbers through, and revoke the shared ones. Remaining: Re-provision the robots at the Bremen and Ota sites during their maintenance windows. Sam Okafor, September 30, 2026.
- **SF-SEC-014.** Many instances still allow IMDSv1 (EC2.8), which lets a server-side request forgery bug read instance credentials. An EC2 declarative policy makes IMDSv2 the default for new launches, and existing instances are changed one at a time once the MetadataNoToken metric shows nothing uses v1. Remaining: Switch the instances that still make IMDSv1 calls after their agents and SDKs are upgraded. Elena Vasquez, October 30, 2026.
- **SF-SEC-012.** Many IAM users carry access keys older than 90 days or unused for 45 days or more, and some hold attached policies. Triage from the credential report, deactivate before deleting, replace CI keys with OIDC roles, and block new users and keys. Remaining: Replace the factory gateway users' keys with IAM Roles Anywhere during factory maintenance windows. Grace Liu, November 20, 2026.

The Medium and Low findings that remain are in the roadmap in the handover report, with the same owner-and-date discipline.

The SOC 2 map, after

SOC 2 (2017 Trust Services Criteria, with the revised points of focus of 2022). The map is the compliance mapping from the audit, updated for the re-scan: for each criterion, whether what AWS shows meets it, which findings bear on it, and which are still open. It is a map of AWS-side evidence, not an attestation, not an audit opinion, and not a certification. Your auditor decides what satisfies a criterion; this tells them what to ask for.

The 38 criteria by status, before and after. "Not assessed: organizational" are the criteria that depend on policy, people, and process, outside what AWS shows.

Status	Week two	Week six
Met on the AWS side	3	10
Partially met	9	12
Gap	15	5
Inherited from AWS	2	2
Not assessed: organizational	9	9

The SOC 2 criteria that AWS evidence bears on, at week six. The organizational criteria are omitted; they are the auditor's and StrataForge's, not AWS's.

Criterion	Title	Status	Open findings	Worst open	What AWS shows
CC1.3	Structures, reporting lines, and authorities	Partially met	SF-SEC-047	Low	AWS Organizations OU structure with an accountable owner for every account (Team and Owner tags).
CC2.1	Relevant, quality information	Partially met	SF-SEC-047	Low	AWS Config aggregator inventory and a tagging coverage report that names an owner for each resource.
CC2.2	Internal communication of control information	Met on the AWS side	none	none	Runbooks and PagerDuty routing for Security Hub CSPM and GuardDuty findings, with on-call acknowledgments.
CC2.3	Communication with external parties	Met on the AWS side	none	none	An alternate SECURITY contact on every AWS account.
CC3.2	Identification and analysis of risk	Met on the AWS side	none	none	This findings register, its ratings, and the risk-acceptance log.
CC3.4	Changes that affect internal control	Met on the AWS side	none	none	Integration records for the Kestrel acquisition: Control Tower enrollment and removal of pre-acquisition trusts.
CC4.1	Ongoing and separate evaluations	Met on the AWS side	none	none	Security Hub CSPM scores per account and standard, AWS Config compliance, and CloudTrail digest validation.
CC4.2	Communication of deficiencies	Met on the AWS side	none	none	The remediation roadmap with owners and dates, and the week-six re-scan.
CC5.2	General controls over technology	Met on the AWS side	none	none	SCP, RCP, and declarative policy documents with their attachment targets; Control Tower enrollment status.

Criterion	Title	Status	Open findings	Worst open	What AWS shows
CC6.1	Logical access security, including encryption	Gap	SF-SEC-012, SF-SEC-015, SF-SEC-030, SF-SEC-032, SF-SEC-035, SF-SEC-036, SF-SEC-037	High	IAM Identity Center permission sets and assignments, the IAM credential report, KMS key policies, and encryption defaults.
CC6.2	Provisioning and removal of access	Gap	SF-SEC-012	High	Okta-to-Identity Center SCIM provisioning records, the credential report, and key retirement records.
CC6.3	Least privilege and access changes	Partially met	SF-SEC-046, SF-SEC-048	Low	Access review exports, IAM Access Analyzer unused-access findings, and permissions boundaries.
CC6.4	Physical access	Inherited from AWS	none	none	AWS's SOC 2 report covers the data centers (carved out).
CC6.5	Disposal of protected assets	Inherited from AWS	none	none	AWS media sanitization is covered by AWS's SOC 2 report.
CC6.6	Protection against threats from outside the system boundary	Gap	SF-SEC-014, SF-SEC-016, SF-SEC-030, SF-SEC-034, SF-SEC-043	High	Security group rules open to 0.0.0.0/0 (only the documented ones), EKS endpoint settings, WAF associations, IoT policies.
CC6.7	Restricted transmission and movement of data	Partially met	SF-SEC-034, SF-SEC-036	Medium	The RCP that enforces TLS, load balancer listener policies, and ElastiCache in-transit encryption.
CC6.8	Prevention and detection of malicious software	Partially met	SF-SEC-033	Medium	GuardDuty Malware Protection and Runtime Monitoring coverage, and Amazon Inspector coverage.
CC7.1	Detection of configuration changes and vulnerabilities	Partially met	SF-SEC-031, SF-SEC-033	Medium	Inspector findings against remediation SLAs, AWS Config rule compliance, and Security Hub CSPM control status.

Criterion	Title	Status	Open findings	Worst open	What AWS shows
CC7.2	Monitoring for anomalies	Partially met	SF-SEC-043, SF-SEC-054	Medium	GuardDuty in every account and Region, the organization trail with validation, and alarms on root and break-glass use.
CC7.3	Evaluation of security events	Met on the AWS side	none	none	Triage records for paged findings.
CC7.4	Incident response	Met on the AWS side	none	none	Runbooks, and records of their use and of the quarterly exercise.
CC7.5	Recovery from incidents	Partially met	SF-SEC-045	Medium	Restore tests from AWS Backup and the logically air-gapped vault.
CC8.1	Change management	Partially met	SF-SEC-031, SF-SEC-054	Medium	Reviewed pull requests for infrastructure changes, two-person approval of firmware releases, CI role trust conditions.
CC9.1	Mitigation of business-disruption risk	Partially met	SF-SEC-045	Medium	AWS Backup plans, cross-Region copies, and backup retention settings.
CC9.2	Vendor and business-partner risk	Gap	SF-SEC-012	High	An inventory of vendor access: cross-account roles with ExternalId and no vendor IAM users with long-lived keys.
A1.2	Environmental protections, backups, and recovery infrastructure	Partially met	SF-SEC-045	Medium	Multi-AZ settings, backup retention, and the logically air-gapped vault for production data.
A1.3	Recovery plan testing	Partially met	SF-SEC-045	Medium	Restore-test records for production databases.
C1.1	Identification and protection of confidential information	Gap	SF-SEC-015, SF-SEC-032, SF-SEC-035, SF-SEC-036, SF-SEC-037	High	Data classification tags, encryption settings, and public-access guardrails for buckets and snapshots.

Criterion	Title	Status	Open findings	Worst open	What AWS shows
C1.2	Disposal of confidential information	Met on the AWS side	none	none	S3 lifecycle rules and snapshot retention settings.

5 criteria are still a gap, and each is a gap because of open findings with owners and dates, not because of something unknown. Expect an auditor to ask about exactly these.

- **CC6.1, Logical access security, including encryption.** Open: SF-SEC-012, SF-SEC-015, SF-SEC-030, SF-SEC-032, SF-SEC-035, SF-SEC-036, SF-SEC-037.
- **CC6.2, Provisioning and removal of access.** Open: SF-SEC-012.
- **CC6.6, Protection against threats from outside the system boundary.** Open: SF-SEC-014, SF-SEC-016, SF-SEC-030, SF-SEC-034, SF-SEC-043.
- **CC9.2, Vendor and business-partner risk.** Open: SF-SEC-012.
- **C1.1, Identification and protection of confidential information.** Open: SF-SEC-015, SF-SEC-032, SF-SEC-035, SF-SEC-036, SF-SEC-037.

The gate, and what follows

The guarantee reads, in the statement of work: the second half of the Hardening fee is invoiced only when every critical finding in the week-two register is closed, verified by the week-six re-scan, and a critical accepted in writing counts as closed. This is the computation.

The gate computation, decided August 21, 2026.

Step	Result
Critical findings in the week-two register	10
Closed, verified by the re-scan	9
Accepted in writing as a risk	1
Still open	0
Closed or accepted, divided by critical	10 of 10
Gate result	PASS

Second invoice issued. The second invoice, for \$17,500 (half of the \$35,000 fixed fee, and the part the guarantee protected), is issued on August 21, 2026, due within 15 days under the statement of work. The first half was invoiced at signature. Nothing is owed beyond the fee for the six weeks, and the work product is yours: the code in your four repositories, the policies, the runbooks, the register, the recordings, and this report.

Two sentences the sponsor may need. First, passing the gate is a statement about the critical findings in a register we both agreed in week two; it is not a statement that StrataForge cannot be breached, and it is not a SOC 2 report. Second, the score will drift down if nothing is done: a new account, a new workload, or a new engineer adds failing controls faster than anyone expects. The handover report says how StrataForge keeps it where it is.

What to do next

1. Read the handover report: the runbooks in full, the remaining roadmap with owners and dates, and the guardrails StrataForge now owns.
2. Put the two dates that matter in the calendar. December 15, 2026: the acceptance for SF-SEC-002 expires. September 30, 2026: the earliest of the High findings is due.
3. Run the re-scan yourselves in the first week of each month for the next quarter, with the commands in the handover report, and compare against this report. If the score falls by more than a point, find out why that week.

© 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.