Vantalect

AWS SECURITY, COMPLIANCE & GOVERNANCE · SAMPLE DELIVERABLE

Audit Sponsor Readout

Week 2 of 6: Sponsor readout · prepared for StrataForge Robotics, Inc.

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

July 24, 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

The score 2

What an attacker could reach today 4

Four paths, in plain terms 5

Decisions we need from you 6

The detection cost increase 7

The four weeks of hardening 8

The guarantee gate 10

Scoring methodology revision 10

The score

StrataForge's AWS organization scores 30.1% on the Acumen raw security score (85 selected versioned controls from the CIS and AWS benchmarks), and it has 10 Critical findings, 6 of them usable by anyone on the internet with no credentials. Marcus Oyelaran and Priya Raman: this is the two-week Audit in one document. We scanned all 48 accounts in 17 Regions, verified every Critical and High result by hand, and rated each real finding for how exploitable it is and how far it reaches. The register has 54 findings: 10 Critical, 20 High, 15 Medium, and 9 Low.

Evidence coverage: **2,852/3,044 nonexcluded account/control units evaluated**; 192 unavailable, 0 not evaluated, 1,036 proven empty, and 0 excluded by policy. Lost evidence can raise a partial ratio; compare the counts and the source scope as well as the score.

The score is the least interesting number here. It says the organization is below where an audit will expect it. The findings say why, and the next section says what an attacker could do with them.

The audit in eight numbers.

Measure	Value
Acumen raw security score, 85 selected versioned controls	30.1%
Verified findings	54
Critical findings	10
Critical findings that need no credentials	6
Critical findings that reach robots in the field	2
Engineering effort to close every Critical	137 hours
Engineering effort for the whole register	1,165 hours
Increase in detection cost	$60,745 a month

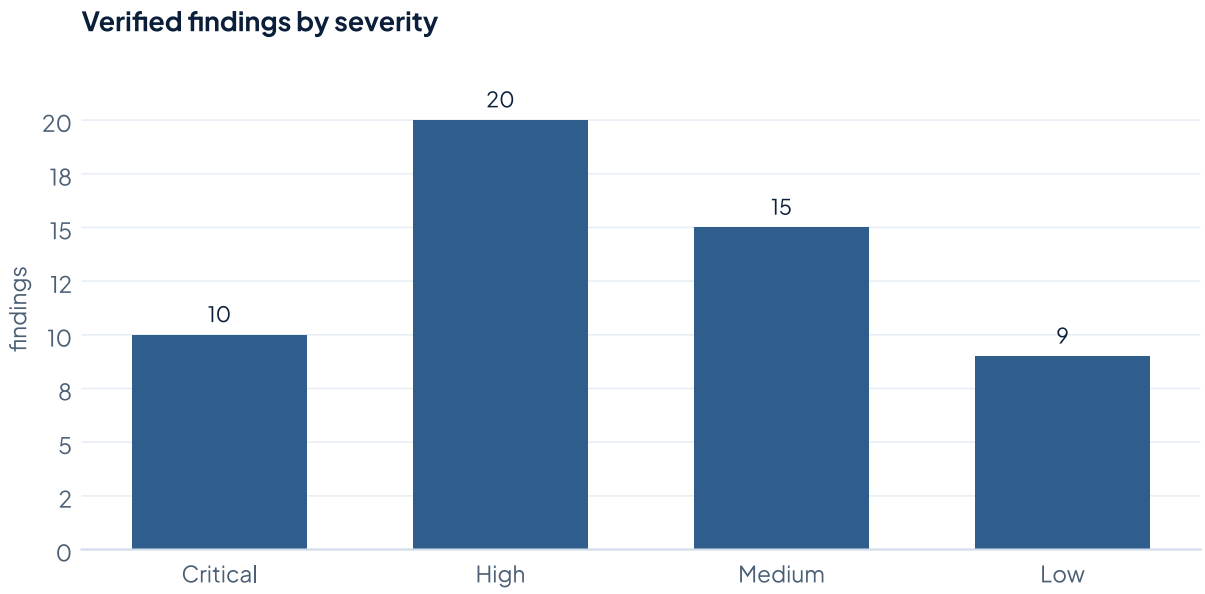

Verified findings by severity.

For the SOC 2 audit: of 38 criteria mapped, 15 have a gap on the AWS side today, 9 are partly met, 3 are met, 2 are inherited from AWS, and 9 are organizational and outside what AWS shows. The criteria for logical access, protection against outside threats, and confidential information carry the most findings, and the Critical findings are what an auditor would find first. The map is in the audit report; it is a readiness map, not an opinion.

What an attacker could reach today

An attacker on the internet, with no credentials and no inside knowledge, can do 6 of the 10 things below. Another 3 need one stolen credential or one application bug, and 1 needs a foothold inside. The table lists all ten, in the order we recommend closing them.

The Critical findings, with what it takes to use each and what it reaches.

ID	What is wrong	What an attacker does	What they reach
SF-SEC-001	A 1,127-day-old IAM user key can sign and publish firmware to every robot	One leaked key is enough: no MFA, source-IP, or VPC-endpoint condition, and the key has been copied to two laptops.	A signed image in the distribution bucket reaches all 212,000 connected robots.
SF-SEC-002	One IoT policy lets any of 61,400 robots' certificates publish to any robot's command topic	Needs one certificate and key, which gen-2 ForgeArm units keep on the filesystem rather than in a secure element.	Command topics sf/v1/<thing>/cmd reach motion planning on every robot on the policy.
SF-SEC-003	4 GitHub Actions deploy roles can be assumed from any GitHub repository	Any GitHub user can mint a token with audience sts.amazonaws.com from their own workflow.	gha-artifact-publisher pushes the images every workload account deploys.
SF-SEC-004	Jupyter is open to the internet on 51 research and sandbox instances, some with no token	Nine instances were launched with an empty notebook token: anyone reaching 8888 runs code.	Instance roles read the Crucible sensor-log lake and Kestrel model artifacts over IMDSv1.
SF-SEC-005	5 EBS snapshots in kestrel-prod are public, and snapshot block public access is off in all 816 account-Regions	Any AWS account can create a volume from a public snapshot.	Volumes of the Kestrel perception API hosts, including application configuration.
SF-SEC-006	4 Kestrel dataset buckets are public: 468.8 million objects of labeled images from customer sites	Anyone on the internet can list and download them.	361 TB of images from customer sites; about 4% of a 200-image sample show people or badge IDs.
SF-SEC-007	A Kestrel MySQL snapshot with annotation records is restorable by any AWS account	Any AWS account can restore it into its own database.	Annotator identities and customer-site labels for 1.9 million images.

ID	What is wrong	What an attacker does	What they reach
SF-SEC-008	Kestrel's Jenkins controller is open to the internet on 8080 and 22	Jenkins 2.3× with anonymous read of job configuration; the instance allows IMDSv1.	Jenkins holds the Kestrel CI keys and a root access key for kestrel-legacy-mgmt.
SF-SEC-009	RDP is open to the internet on 32 PLM CAD workstations that hold product designs	RDP brute force and credential stuffing against Windows hosts.	Product designs for every robot line.
SF-SEC-010	Root users of 3 Kestrel accounts have access keys and no MFA	A root key used by a Jenkins job last month lives in Jenkins' credential store.	Root of kestrel-prod and kestrel-data, which hold the perception APIs and the labeled datasets.

Four paths, in plain terms

Through Kestrel Vision. Kestrel's Jenkins controller is open to the internet (SF-SEC-008). It stores Kestrel's CI keys and a root key, and any of those CI keys can give itself administrator and, through a trust left over from before the acquisition, run four Kestrel accounts, including the ones that hold the perception APIs and the labeled datasets. That is the hard way in. The easy way is that 4 Kestrel dataset buckets are already open to the whole internet: 361 TB of images from customer sites; about 4% of a 200–image sample show people or badge IDs. A database snapshot of annotator identities can be restored by any AWS account (SF-SEC-007). Images from customer sites are the data your customers would least like to see exposed.

Through GitHub. A workflow in any GitHub repository can assume 4 StrataForge deploy roles, so anyone with a free GitHub account can run as one (SF-SEC-003). One of them pushes the container images that every workload account then deploys. A tampered image would run in production with StrataForge's own credentials.

Through the robots. One access key can sign firmware and publish it, and a signed image reaches all of the connected robots (SF-SEC-001). One leaked key is enough: no MFA, source-IP, or VPC-endpoint condition, and the key has been copied to two laptops. Separately, one robot certificate can publish to any robot's command topic (SF-SEC-002), and the gen-2 ForgeArm units keep their certificate key on the filesystem. Both are safety findings: they are about what a robot in a customer's building could be told to do.

Through a notebook or a desktop. Jupyter is open to the internet on 51 research and sandbox instances, some with no password, and 1,631 instances across the organization still allow the older metadata service that lets one application bug steal the instance's credentials (SF-SEC-004, SF-SEC-014). Remote desktop is open to the internet on the workstations that hold product designs for every robot line (SF-SEC-009).

What we have not looked for

This is a configuration review. We did not look for evidence that any of these has been used, and the lack of alerts means little: GuardDuty runs in 31.6% of account-Regions and not in any Kestrel account. Whether anyone has already read the public buckets or restored the snapshot is a question for S3 and CloudTrail access history, which Jonah Reyes's team and Nadia Petrova's Security Operations team can answer. Until they do, treat the contents of the public buckets as disclosed.

Decisions we need from you

Seven decisions come out of the Audit. The first three gate week 3.

The decisions, in the order we would take them.

No.	Decision	Who decides, and when
1	**Agree the register,** and agree that its 10 Critical findings (SF-SEC-001 to SF-SEC-010) are the ones the guarantee is measured on	Marcus Oyelaran, at this readout
2	**Commit the owners.** Week 3 carries 309 hours of change that Vantalect drafts as pull requests. The 6 owners (Sam Okafor, Elena Vasquez, Hiro Tanaka, Jonah Reyes, Mei Chen, Grace Liu) each review and merge the changes to their own accounts within two business days, starting the week of July 27, 2026	Priya Raman, this week
3	**Approve the fleet changes.** The IoT policy and per-robot certificate fixes (SF-SEC-002 and the High finding SF-SEC-030) touch robots in the field, so they roll out by thing group on a bench first, and the robots at the Bremen and Ota sites wait for their maintenance windows	Marcus Oyelaran and Sam Okafor, this week
4	**Approve removing root credentials from every member account** and keeping one hardware-MFA root in the management account. A root-only task then runs from a short-lived root session started in the management account	Marcus Oyelaran, before week 3 ends

No.	Decision	Who decides, and when
5	**Take Kestrel into the landing zone:** enroll its five accounts in Control Tower, remove the trust left over from before the acquisition, and retire the Jenkins controller or put it behind single sign-on	Priya Raman, with Jonah Reyes, by week 5
6	**Approve the detection cost increase** of $60,745 a month on the bill's basis, list prices less the 7% discount (next section)	Priya Raman and Ravi Menon, before week 5
7	**Decide who can accept a risk in writing.** An accepted Critical counts as closed for the guarantee, and it stays on the register and the SOC 2 map. We recommend the CISO signs, and that any accepted Critical has an expiry date	Marcus Oyelaran, before the week-six re-scan

Two deserve a sentence. **Decision 3 is the one with a safety dimension.** A policy that is tighter than the robots expect can disconnect them, which is why the rollout is staged and the bench test comes first. The finding is verified: Reproduced with the fleet team on a bench robot on 07/22/2026: its certificate published to another robot's command topic. **Decision 2 is the one that decides whether week 3 happens.** Week 3 is the heaviest week of the engagement, and the Criticals are only 137 of its 309 hours; the rest is the guardrails at the top of the organization and the High findings that share the week.

The detection cost increase

Turning detection on everywhere raises StrataForge's monthly security-service bill by $60,745, from $58,900 to $119,645. That is 103% more, or $728,940 a year, for GuardDuty and its protection plans, Security Hub CSPM, AWS Config, and Amazon Inspector in all 17 Regions of all 48 accounts. The figures are on the basis the monthly bill uses: AWS list prices retrieved October 13, 2026, each Region at its own rate, less the 7% discount in StrataForge's private pricing agreement.

Detection cost before and after, on the bill's basis.

	Per month	Per year
Detection today (GuardDuty, Security Hub CSPM, AWS Config)	$58,900	$706,800
Detection after week 5	$119,645	$1,435,740
Increase	**$60,745**	**$728,940**

Four things about the number. First, "today" is not an estimate: it is what the July 2026 bill carried for the three services that were running, which the cost engagement's baseline report shows as the same figure, and it includes nothing for the five Kestrel accounts, which have none of them. Second, "after" is the same three services in every account and Region ($85,689) plus $33,935 for the GuardDuty protection plans and Amazon Inspector, which are new. Those are priced from the instances, functions, clusters, and databases we counted, at each Region's price; four lines rest on usage the organization has not yet generated (container image volume, Kubernetes audit events, Lambda network activity, and malware scans), and the week-5 report states each assumption. Third, IAM Access Analyzer's external-access findings are free and the alert routing costs pennies, so neither is in the table; each protection plan has a free trial (GuardDuty's is 30 days for a newly enabled account), and we will compare the first real bill against this estimate in week 6. Fourth, the organization trail's data events for S3 and Lambda cost $56,764 a month already and are unchanged by this work; we mention them so that the total is not misread. The increase is the price of being able to see an attack; it is decision 6.

The four weeks of hardening

If you approve the roadmap, the Hardening carries it out through StrataForge's own change process, in the four repositories, with the engineers who own each account. Every change is a pull request read as code in a twice-weekly change review before it merges, and nothing is applied without its owner's yes.

The Hardening, week by week.

Week	Starts	What goes in	Findings	Hours
3	Jul 27	Every Critical closed as code, proven in non-production before production. Service control policies, a resource control policy that requires TLS, and an EC2 declarative policy go in at the top, so what closed stays closed	24	309
4	Aug 3	Identity: people on single sign-on, long-lived keys retired, workload roles cut to least privilege, root and break-glass locked down and alarmed	9	174
5	Aug 10	Detection and response: GuardDuty, Security Hub CSPM, IAM Access Analyzer, AWS Config, and Amazon Inspector in every account and Region, a new organization trail in an archive no application account can alter, and runbooks	11	116
6	Aug 17	The same scan as week 1, run the same way; the before and after; a 2-hour handover workshop		
After	Aug 21	The backlog: findings and parts of findings that wait on a date or a dependency outside the engagement	10	566

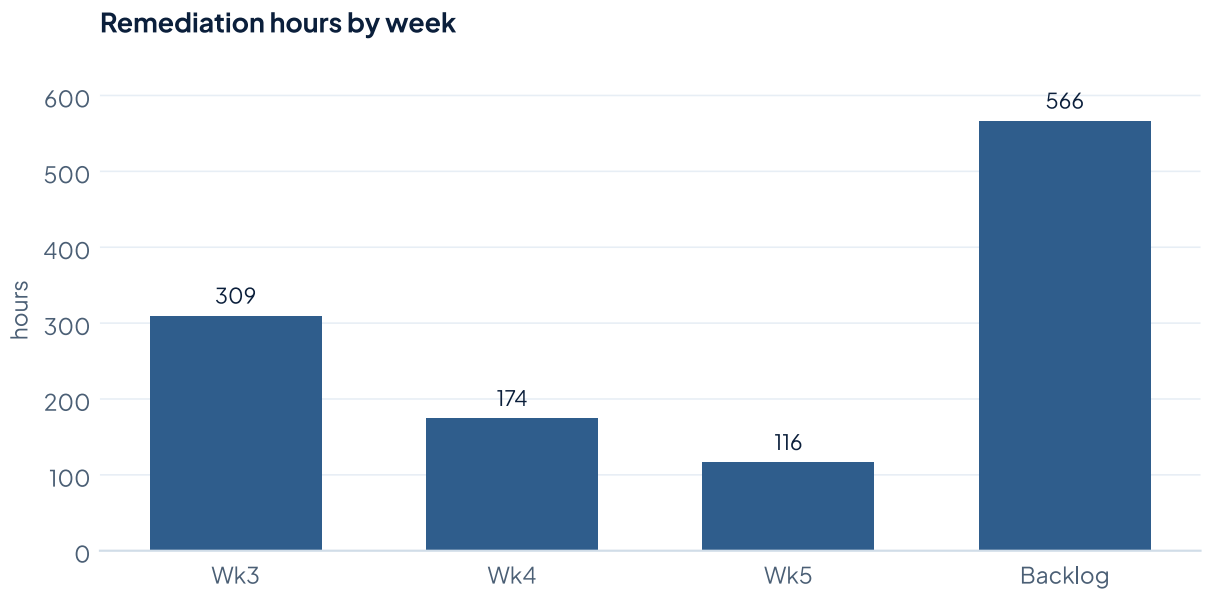

Remediation hours by week. Backlog hours include the parts of earlier findings that wait.

You will see Elena Vasquez, your platform lead, and us in a weekly 30-minute sponsor check-in and in the change reviews. About 566 hours of work are left for StrataForge after week 6; the audit report gives each item an owner and a date, and the handover workshop is where they are taken on.

The guarantee gate

> **The guarantee**
>
> The second half of the Hardening's $35,000 fee, $17,500, is invoiced only when every Critical finding in this audit is closed, verified by the week-six re-scan. If the gate is not met, the second invoice is not issued and you keep everything delivered. A Critical finding you accept in writing as a risk counts as closed, and so does a fix you approved and then chose not to apply.

The gate is measured on the 10 Critical findings in this register, SF-SEC-001 to SF-SEC-010, as agreed with Marcus Oyelaran today. New findings the re-scan turns up are reported with a proposed fix, and are not part of the gate. Beyond the gate, we cannot promise that an environment will not be breached: the Hardening reduces risk, and it is not an audit opinion, an attestation, a certification, or a penetration test.

What we ask of you at this readout: agree the register, name the owners for week 3, and decide the other five items above. Everything else is already in motion, and the first change review is the week of July 27, 2026.

Scoring methodology revision

This revision uses **Acumen's `acumen-security-score/2` raw method**, ratified October 26, 2026. It is an independent simulation score, not a claim to reproduce the AWS console. The selected universe is 37 CIS AWS Foundations 5.0.0 requirements and 48 FSBP 1.0.0 controls, an overall union of 85 versioned controls. Each control counts once per account after all required Regions and members are rolled up. One failure dominates unavailable evidence; otherwise any unavailable required member prevents a pass. Proven empty inventory is separate from missing evidence. The score is passed divided by passed plus failed, with no score when that denominator is empty.

Accepted risk, remediation decisions, and suppression leave raw technical verdicts unchanged. The native control tables remain drilldowns and do not define this denominator. Each ratio retains its numerator, denominator, and nine-place half-even fraction in the facts; report percentages round once. Six required members have no shipped implementation in this source revision: authorized-port security groups, EBS snapshot public-access blocking, ECR image scanning, EKS version support, administrator policies, and RDS snapshot sharing outside the organization. They remain unavailable rather than being dropped. The legacy Elasticsearch inventory (native ES.2) is not modeled and also remains unavailable; an unmodeled service is never declared empty.

Evidence coverage: **2,852/3,044 nonexcluded account/control units evaluated**; 192 unavailable, 0 not evaluated, 1,036 proven empty, and 0 excluded by policy. Lost evidence can raise a partial ratio; compare the counts and the source scope as well as the score.

Raw account/control counts and evidence coverage for this report's assessment. Proven empty and policy-excluded units are outside evidence coverage; unavailable and unrun units remain inside it. An account percentage followed by a fraction states that alias's evidence coverage for the account.

Alias	Passed	Failed	Unavailable	Not evaluated	Proven empty	Policy excluded	Evidence coverage
cis	521	848	0	0	407	0	1,369/1,369
fsbp	337	1,146	192	0	629	0	1,483/1,675
overall	858	1,994	192	0	1,036	0	2,852/3,044

© 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.