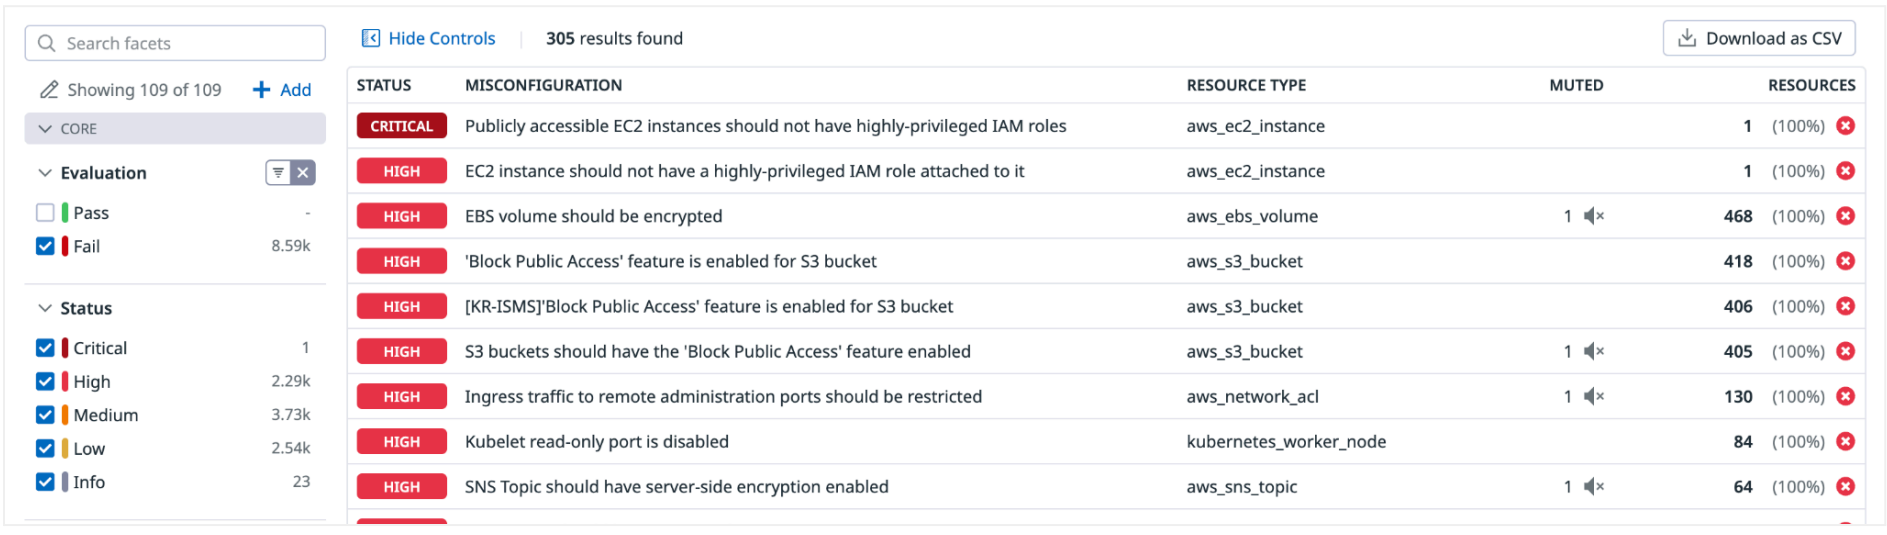
[Design Article]: New Cloud Security Management Compliance

QUICK LINKS **Full Project Definition Doc** | **Figma File** | **Full Prototype**

> ***Disclaimer***
>
> This 2024 CSPM redesign proposal follows the latest improvements in Datadog Cloud Security. With Cloud Security Management's inception, CSPM becomes its compliance subproduct. Product management recommendations and engineering questions have been gathered here. A first round of user interviews has been made with our *"internal customers"* and we are conducting user testing sessions with real customers.
>
> ⏳ **Open for feedback until Friday, June 29th, 2023**.

Security teams catch just 35% of cloud misconfigurations. The rest? Buried in noise until breaches or audits surface them. To close these gaps, Datadog's Cloud Security Posture Management helps security professionals assess their high-level compliance posture, investigate cloud misconfigurations, and compile security reports to pass mandatory compliance audits. CSPM performs continuous scanning and tracks every resource for configuration checks across cloud accounts, hosts, and containers. This proposal details how we will turn detection gaps into proactive insights through human-centric workflows.



I/ Problem Framing

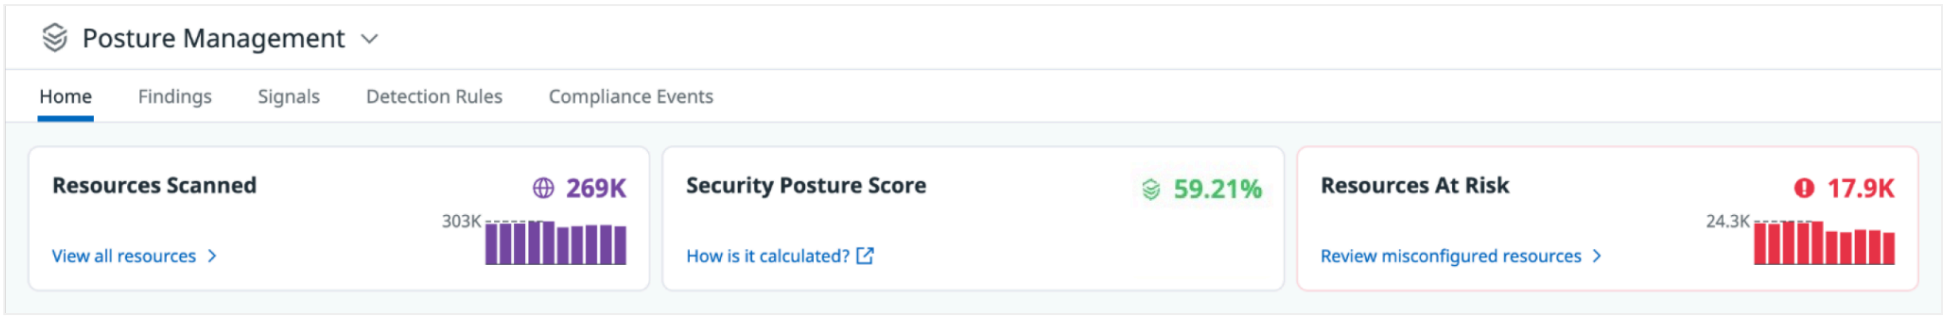
Context

Today, security teams struggle to identify and remediate daily upcoming cloud misconfigurations while meeting complex government and industry regulations frameworks. Cloud Security Posture Management (CSPM) is a Datadog CSM sub-product that scans organizations' cloud infrastructure for misconfigurations and compliance risks. Datadog's Security products are very technical: they are built by engineers for engineers. Their complexity is often hard to simplify for users who need to answer daily stressful alert events. Adding to the complexity, users often can't directly fix security issues and must work with DevOps teams who frequently lack the proper context or prioritization, to tackle them quickly.

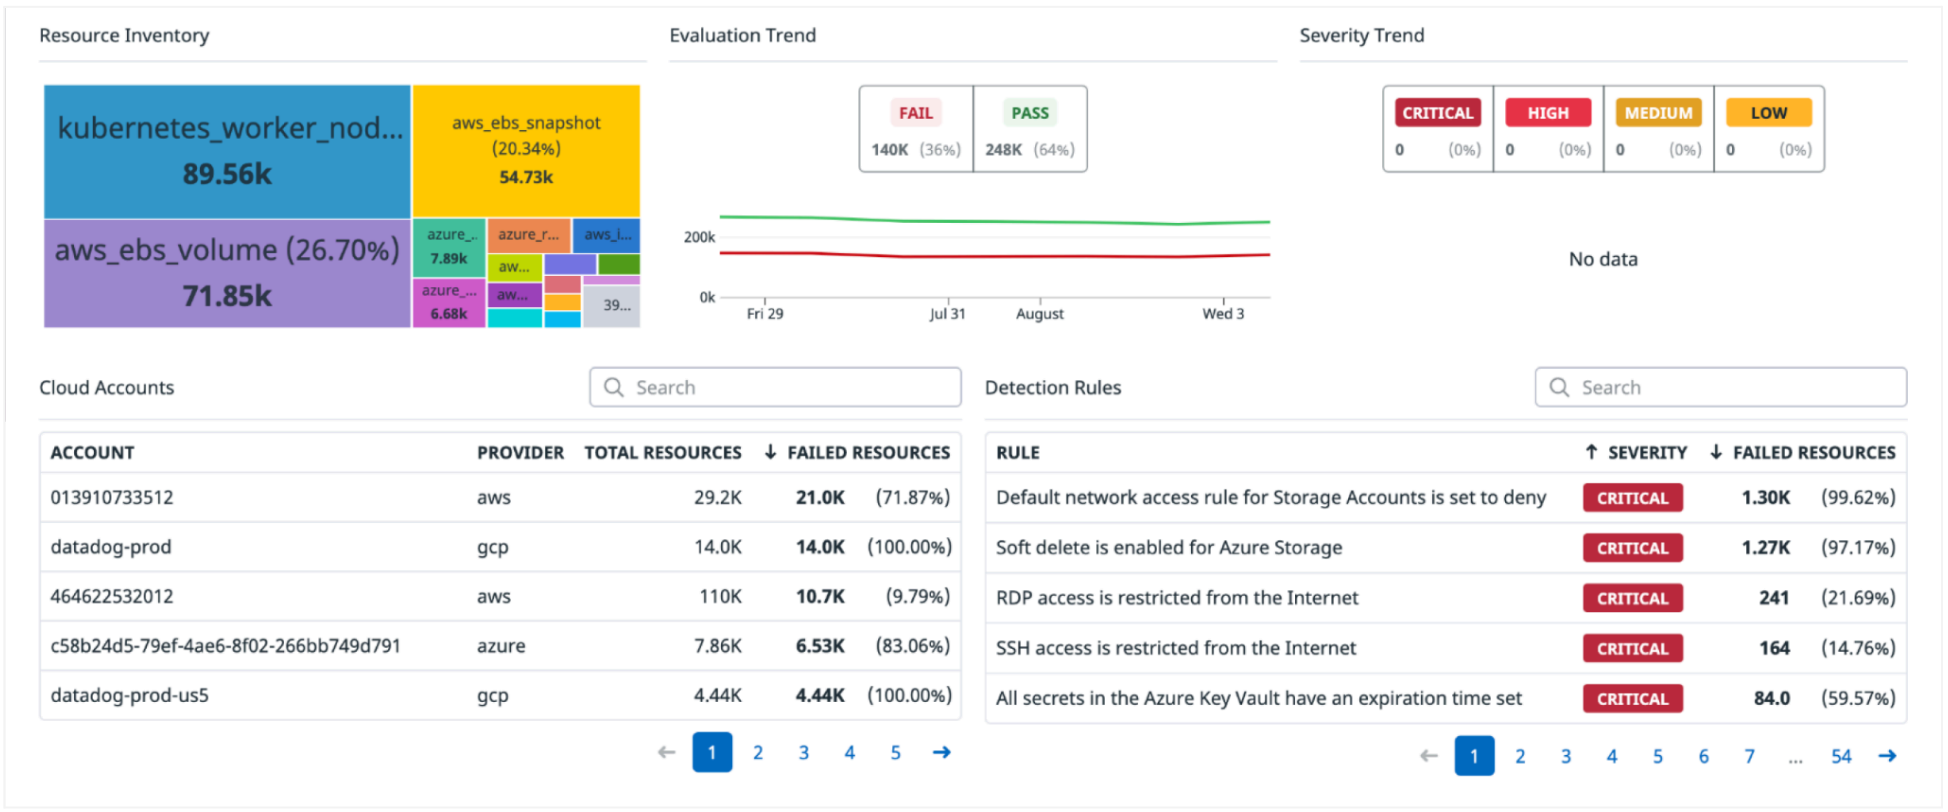
Problems

We've noticed a trend where new users seldom return to Datadog's CPSM pages after their initial setup. We believe this high churn is partly due to the failure to showcase concrete value during - and immediately following - the onboarding process. As a result, new security users and expert users began to doubt the usefulness of CSPM as they couldn't quickly understand what to fix, where to start their investigation, and what to prioritize.

Compliance Overview page



Visibility is the starting point of any security protection plan. You can't protect what you cannot see. In our case, users land on a first section with very preliminary data and limited actionability. Security teams deserve more than a fragmented data and vague overview: here it buries critical risks in a cluttered interface, leaving teams guessing what to prioritize. Without clear insights or actionable steps, remediation becomes a game of chance.



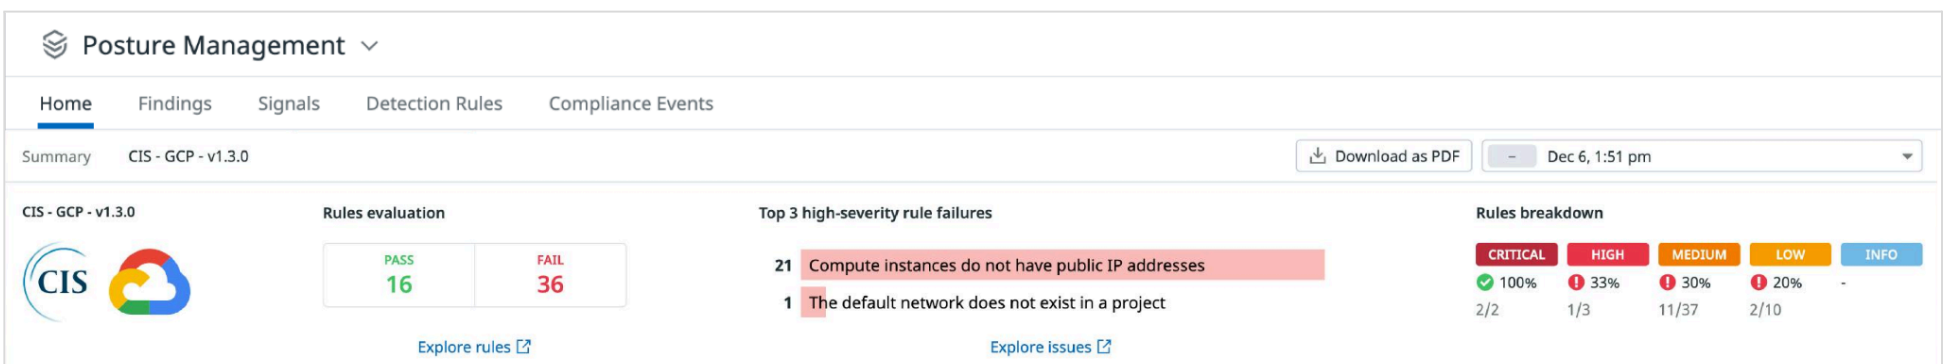
"*Less is More*": users want to see information, but only what's relevant to them. And they don't want to be bombarded with it all at once. Users already battle alert fatigue and fragmented tools: this section worsens the problem by overloading them into a sea of raw data without clear next steps. The combo Treemap / Timeseries / Tables is also breaking workflows as users have to jump between pages to connect this data to the infrastructure security context (Findings Explorer, Detection Rules, Resource Catalog, etc) which does not help with triage and prioritization.



The page ends with a long suite of compliance frameworks, each leading to a dedicated compliance framework report page. The current rules evaluation and top insights are not helping users to understand what changed over time or easing to track progress toward complete remediation. Based on our competitive analysis, we initiated debates with engineering on how we could improve the security posture evaluation for every supported compliance framework: a security posture score in percentage seems more significant to showcase compliance levels, as it would use a weighted ratio of misconfiguration's severity and the number of pass/fail compliance rules for each severity.

Compliance Framework Report page

When clicking on a compliance framework row in the compliance overview page, we access its dedicated compliance framework report page which provides security insights including ongoing compliance rule failures and their related misconfigured resources.



The navigation still reflects some of CSPM's features that will be discontinued in the coming months. Still, above this, when clicking on a compliance framework row in the CSPM overview page, the navigation doesn't reflect it at all and says to users that they are still in the CSPM overview page. This confusion must be fixed to show the pathway to this page. The top of the page is cluttered with data, making it difficult for users to spot key critical insights. If they want to learn more, their only option is to be redirected to the misconfiguration explorer page and its nested side panel ("*explore rules*" and "*explore issues*" buttons). Security data, such as passed/failed rule evaluation, top failures, or high-severity rules, don't provide the optimal triage value: we lack clear actionable contextual information or guidance to address the everlasting overflow of misconfigurations.



The page remains static, offering limited interactivity and poor filtering or sorting options. The dense presentation of compliance rules, controls, and percentages can overwhelm users: the lack of filtering options makes it harder for users to spotlight the rules with the most fail. Moreover, this compliance framework report page is meant to help security engineers prepare for audits but doesn't provide any compliance trend over time.

Discovery

Main Use case

- *As a CSM / CSPM user* I want to assess my compliance posture at a high level, across industry standards *so I can* pinpoint compliance weak spots and prepare ahead of an audit.

Secondary Use case

- *As a CSM / CSPM user* I want to overview relevant security insights about each framework *so I can* identify failing controls and prioritize my next remediation efforts.

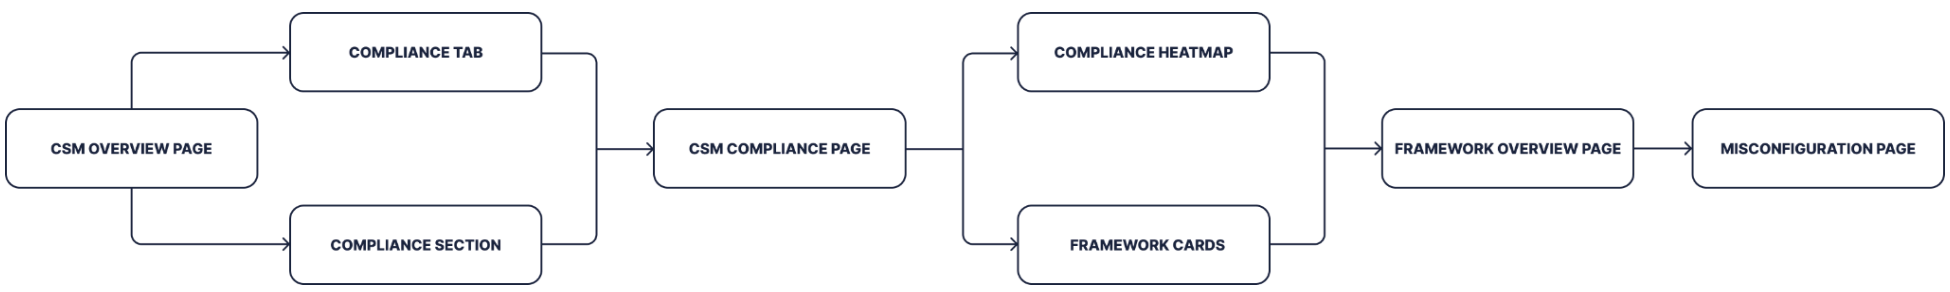
CSPM User Journey

Sec & DevSecOps Teams Journey Map / DATADOG CSPM

Project Scope

Phase	1 Identify Misconfigurations	2 Assess Security Risks	3 Prioritize Issues	4 Remediate Issues	5 Monitor Compliance	6 Compile Audit Evidence
User tasks	• Datadog CSPM fires the alert • User receives Slack notifications • User checks misconfigurations info and confirms ownership	• Access the Datadog CSPM overview page • View the notified risks in the CSM / CSPM explorer page	• Access the targeted framework detail page • Analyze the severity of each misconfiguration • Gather information • Share on Slack if deeper analysis required	• Assign remediation tasks to relevant teams • Fix misconfigurations, update policies, or set up automated remediation	• Assign monitoring tasks to relevant teams • Build team dashboards showing ongoing security compliance, and infrastructure data from all of the resources	• Gather and compile evidence for compliance audits, including logs, compliance framework reports, and list of remediation actions
Feeling		2-3 MINS / PER MISCONFIGURATION • Initial notification and acknowledgment	10-12 MINS / PER MISCONFIGURATION • Analysis and prioritization	15 MINS MAX 20 / PER MISCONFIGURATION • Remediation and follow-up		2-4 HOURS / PER AUDIT • Collecting and compiling audit evidence
Channel	• Slack / Pager Duty / Jira • External SIEM integrations or custom webhooks	• Slack / Pager Duty / Jira / ServiceNow + SMS and Email notifications			• Slack / Pager Duty / Jira	• Jira / Google Docs
Touchpoint	• Security teams' Datadog Alert Slack channels	• Datadog CSPM overview page	• Datadog CSPM compliance framework page	• Datadog CSM overview page, the CSM misconfiguration explorer page + CSPM detection rules	• CSM out-of-the-box dashboards or custom	• CSPM overview page + compliance framework detail page + dashboards
Goals	• Promptly receive and acknowledge misconfiguration notifications	• Quickly identify alerts and assess security risks in the cloud environment	• Efficiently prioritize misconfiguration issues based on severity and impact	• Accurately remediate misconfigurations to reduce risk surface and improve cloud security posture	• Continuously monitor compliance and show issues overall impact • Ensure adherence to security policies	• Comprehensively compile audit evidence to demonstrate compliance and security posture
Painpoints	• Potential delays in alert acknowledgment because of notifications or system issues	• Bombarded by alert signals within criptic security widgets • Limited visibility into the context of misconfigurations	• Inefficient prioritization due to lack of clear risk assessment criteria	• Challenges in remediation due to unclear instructions or lack of integration with remediation tools • UI confusion leading to incorrect prioritization or remediation actions	• Inefficient monitoring due to scattered or unclear security data and compliance information	• Difficulty in compiling comprehensive audit evidence (screenshots) due to disorganized data or lack of integration with documentation tools

Simplified User Flow



User Research

User Interviews

Redesigning the CSM compliance experience couldn't be done without understanding how our CSPM users behave from the CSM Overview page, which is the main entry point, through the CSPM Overview page. This deep user research effort was essential to understanding the specific needs of our users and how to address them accordingly. It has also been fueled by an ongoing competitive analysis that is continually nurtured to keep our features up to date with our competitors.

It required focusing on role-specific behavior patterns tied to the distinct responsibilities between security engineers and their leadership. This personalized approach streamlined workflows, rediscovering how our users accomplish their tasks more efficiently and effectively.

For instance, a CISO may prioritize accessing quickly comprehensive dashboards for an overarching view of organizational security issues to export, while a SecOps Engineer might focus on quick access to detailed granular and actionable security insights.

- UX Research Plan & Feedback: CSM Overview Page [ROUND 1] + UX Research Feedback: CSM Overview Page [ROUND 2]:
  The CSM Overview page is the entry point for CWS (CSM Threat) and CSPM (CSM Misconfiguration) products. Our assumptions and latest UX analytics pointed toward this page needing to be better optimized for users and actionable enough. We've researched to validate those assumptions by identifying use cases, usage patterns, and sentiment toward a few redesign explorations.

- UX Research Plan: CSPM [CSM Misconfiguration] Value Drivers: CSPM growth has stagnated over the last few quarters (i.e. churn and growth have matched each other) and we had to conduct this research to understand why and what improvements to target for that it is better positioned.

- UX Research Feedback: CSPM Overview Page: We conducted this research to learn more about customer sentiment and use cases with the existing/proposed CSPM overview page. The primary objectives of this research were to collect feedback on the current CSPM overview page, understand why customers come to the overview page, and overall, identify design improvement opportunities to leverage.

- UX Research Summary: How CSPM Customers Prioritize Findings : We conducted this research to better understand how users prioritize CSPM findings. These insights enabled us to start building a prioritization model to feed into the future redesigned CSM Overview Page. The primary objective of this research was to understand how CSPM customers prioritize findings and to uncover what role tags play in the prioritization process.

User Testing

We are currently conducting user testing sessions using Maze & Ballpark, and the outcomes have already been very positive. We will compile all the results as soon as possible in an appendix with recommendations for improvement.

Early Explorations

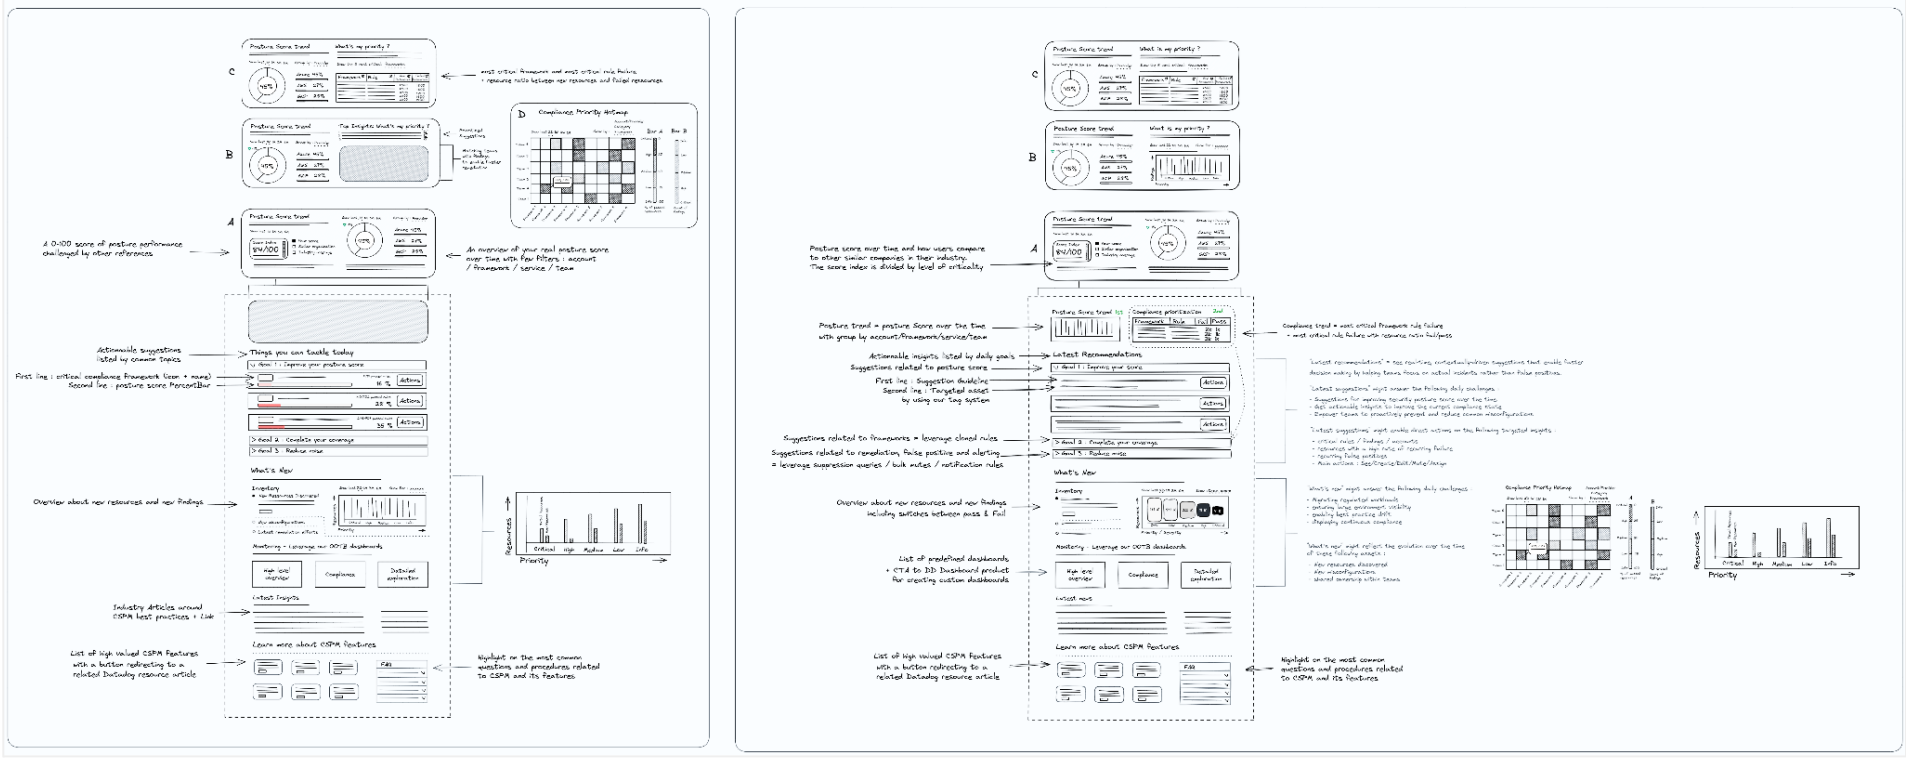
Trigger For Explorations

Our research revealed that security teams struggled to navigate complex compliance data across multiple cloud accounts. Users needed to quickly identify high-risk resources without drowning in noise. The initial sketches (Excalidraw link) reveal the foundational thinking behind transforming a cluttered and reactive compliance experience into a proactive, context-rich security tool, which is a direct response to user pain points around fragmented data.

Key Challenges

These low-fidelity prototypes prioritized three core challenges:

- **Information Overload** : early sketches highlighted the need to simplify complex and overwhelming (600+ resources to audit, in average) security data into actionable insights, replacing a "list-of-everything" approach with dynamic filtering, reliable risk scoring and innovative layout.
- **Contextual Drilling** : this project is connected to the CSM redesign and emphasizes linking misconfigurations directly to downstream risks (like toxic risk combinations), a concept validated through user testing with SREs who struggled to connect isolated findings to broader threats.
- **Compliance as a Workflow** : early sketches integrated compliance prioritization into team ownership details, ensuring team members could quickly map their remediation assignment.



Key Design Decisions

Those early sketches illustrate how we translated these pain points into actionable design decisions :

Resource-Centric Hierarchy : we reflected engineers' mental models by letting them drill down into specific services without losing account context. Grouping risks by cloud service, account and affected resources instead of generic severity lists seemed more adapted to track changes over time.

Compliance Data Density : we researched visual density by showing 50+ accounts and compliance frameworks in condensed views after users called out *navigation fatigue*. We explored a global compliance status into a color-coded heatmap and compliance frameworks under cards, replacing text-heavy tables.

Dynamic Filter Bar : we wanted to reduce the overall cognitive load by placing filters in a collapsible sidebar, but we pivoted toward a horizontal, sticky bar as early user interviews showed horizontal filters reduced vertical scrolling, preserving critical resource visibility.

Actionable Insights : we transformed raw findings into daily goals (like “Improve your posture score”) with progress bars and direct remediation actions.

Security Posture Score : we introduced a 0 →100% Posture Score that will replace multiple raw numbers as a digestible metric. Users will improve this score either by remediating misconfigurations or by fixing their underlying issues.

AS-IS → Cluttered & Data-Heavy

Compliance page

DATADOG

Posture Management

Home | Findings | Signals | Detection Rules | Compliance Events

Resources Scanned 269K
303K
View all resources >

Security Posture Score -
How is it calculated?

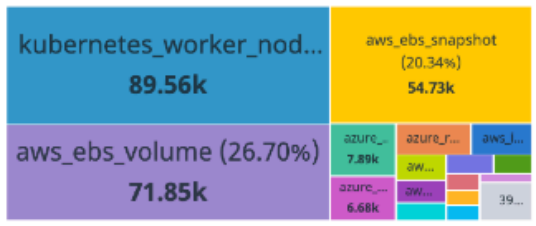
Resources At Risk 17.9K
24.3K
Review misconfigured resources >

Go to...
Watchdog
Events
Dashboards
Infrastructure
Monitors
Metrics
Integrations
APM
CI BETA
Notebooks
Logs
Security
UX Monitoring
Contact Support
Help
Invite Users
Jenil G.
Datadog HQ

Overview

No filters applied - click any widget to filter

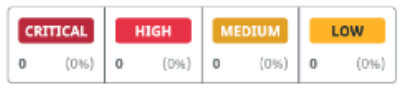
Resource Inventory



Evaluation Trend



Severity Trend



No data

Cloud Accounts

Search

ACCOUNT	PROVIDER	TOTAL RESOURCES	↓ FAILED RESOURCES	
013910733512	aws	29.2K	**21.0K**	(71.87%)
datadog-prod	gcp	14.0K	**14.0K**	(100.00%)
464622532012	aws	110K	**10.7K**	(9.79%)
c58b24d5-79ef-4ae6-8f02-266bb749d791	azure	7.86K	**6.53K**	(83.06%)
datadog-prod-us5	gcp	4.44K	**4.44K**	(100.00%)

← 1 2 3 4 5 →

Detection Rules

Search

RULE	↑ SEVERITY	↓ FAILED RESOURCES	
Default network access rule for Storage Accounts is set to deny	CRITICAL	**1.30K**	(99.62%)
Soft delete is enabled for Azure Storage	CRITICAL	**1.27K**	(97.17%)
RDP access is restricted from the Internet	CRITICAL	**241**	(21.69%)
SSH access is restricted from the Internet	CRITICAL	**164**	(14.76%)
All secrets in the Azure Key Vault have an expiration time set	CRITICAL	**84.0**	(59.57%)

← 1 2 3 4 5 6 7 ... 54 →

Compliance Frameworks

CIS - AWS - v1.3.0

Framework Overview | Explore Resources | Configure Rules

Rules evaluation

PASS	FAIL
1	23

Top 5 requirements by rule failures

IAM	1	12
Logging	0	5
Networking	0	4
Storage	0	2

Resource types with the most fail findings

aws_ebs_volume	69.6K	4.38K
aws_s3_bucket	451	1.93K
aws_network_acl	1	372
aws_iam_user	13.1K	289
aws_security_group	5.82K	259

CIS - Azure - v1.3.0

Framework Overview | Explore Resources | Configure Rules

Rules evaluation

PASS	FAIL
13	42

Top 5 requirements by rule failures

Database-Services	8	9
Logging	0	9
App-Service	0	7
Networking	0	5
Other-Security-Considerations	0	4

Resource types with the most fail findings

azure_managed_disk	1	8.48K
azure_storage_account	2.68K	3.96K
azure_app_service	191	753
azure_security_group	2.73K	424
azure_sql_server	132	210

CIS - Kubernetes - v1.5.1

Framework Overview | Explore Resources | Configure Rules

Rules evaluation

PASS	FAIL
30	26

Top 5 requirements by rule failures

API-Server	20	8
Kubelet	5	5
Pod-Security-Policies	0	4
etcd	2	4
Scheduler	0	2

Resource types with the most fail findings

kubernetes_worker_node	701K	329K
kubernetes_master_node	398	167
kubernetes_cluster	177	166

PCI - v3.2.1

Framework Overview | Explore Resources | Configure Rules

Rules evaluation

PASS	FAIL
46	163

Top 5 requirements by rule failures

Least-Privileged-Access	20	50
Firewall-Configuration	4	43
Monitoring	8	41
Cardholder-Data	4	30
Credentials	7	30

Resource types with the most fail findings

aws_ebs_snapshot	84.7K	23.4K
azure_managed_disk	1	8.48K
aws_ebs_volume	69.6K	4.38K
aws_s3_bucket	14.0K	4.04K
azure_storage_account	2.68K	3.96K

SOC 2 - v2

Framework Overview | Explore Resources | Configure Rules

Rules evaluation

PASS	FAIL
45	160

Top 5 requirements by rule failures

Logical-and-Physical-Access-Control	37	129
System-Operations	10	47
Communication-and-Information	8	40
Control-Activities	13	35
Change-Management	4	20

Resource types with the most fail findings

aws_ebs_snapshot	84.7K	23.4K
azure_managed_disk	1	8.48K
aws_ebs_volume	69.6K	4.38K
aws_s3_bucket	14.0K	4.04K
azure_storage_account	2.68K	3.96K

HIPAA - v1

Framework Overview | Explore Resources | Configure Rules

Rules evaluation

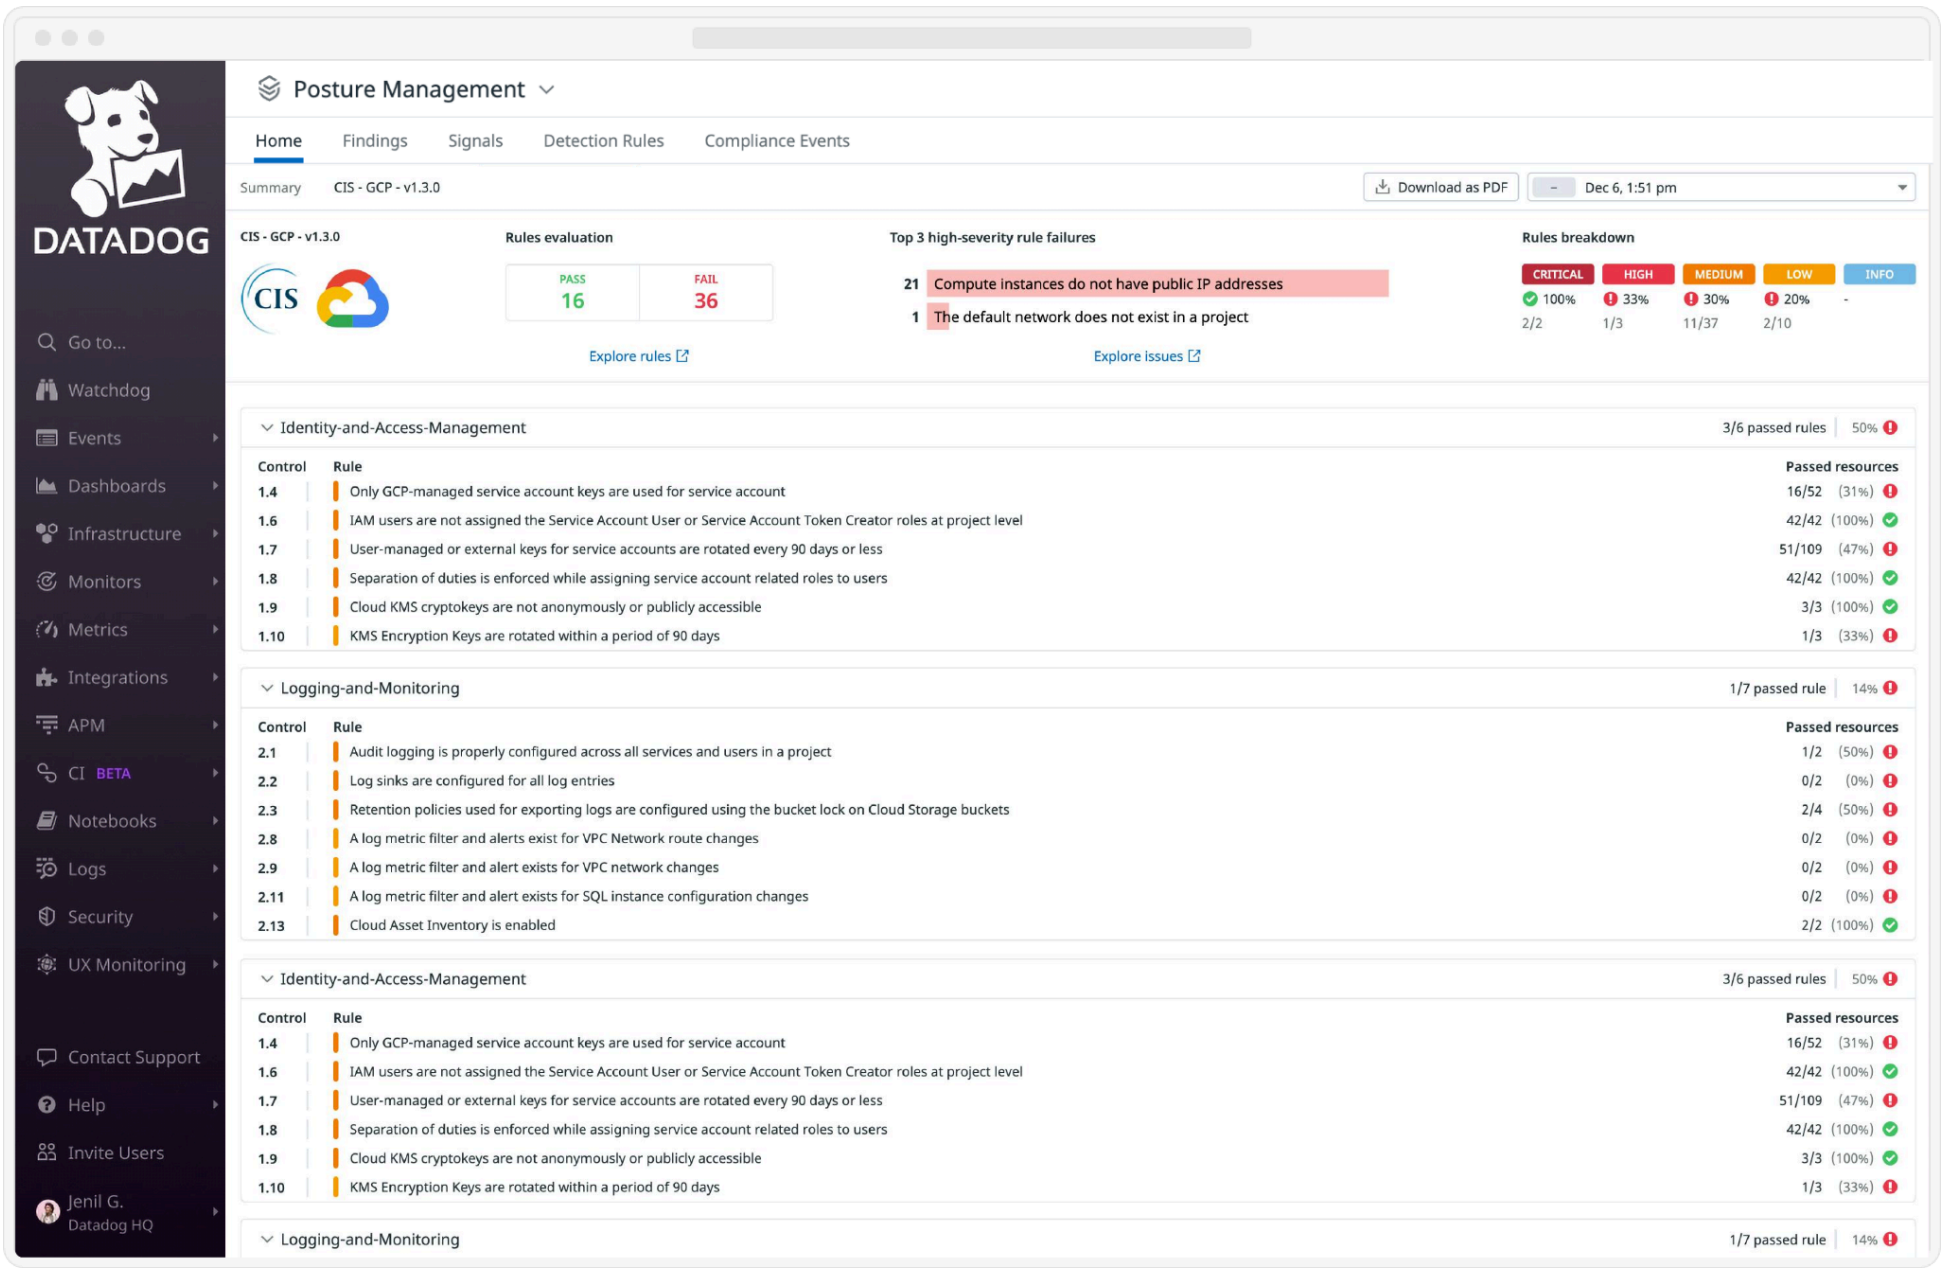
Top 5 requirements by rule failures

Security-Management-Process	18	81

Resource types with the most fail findings

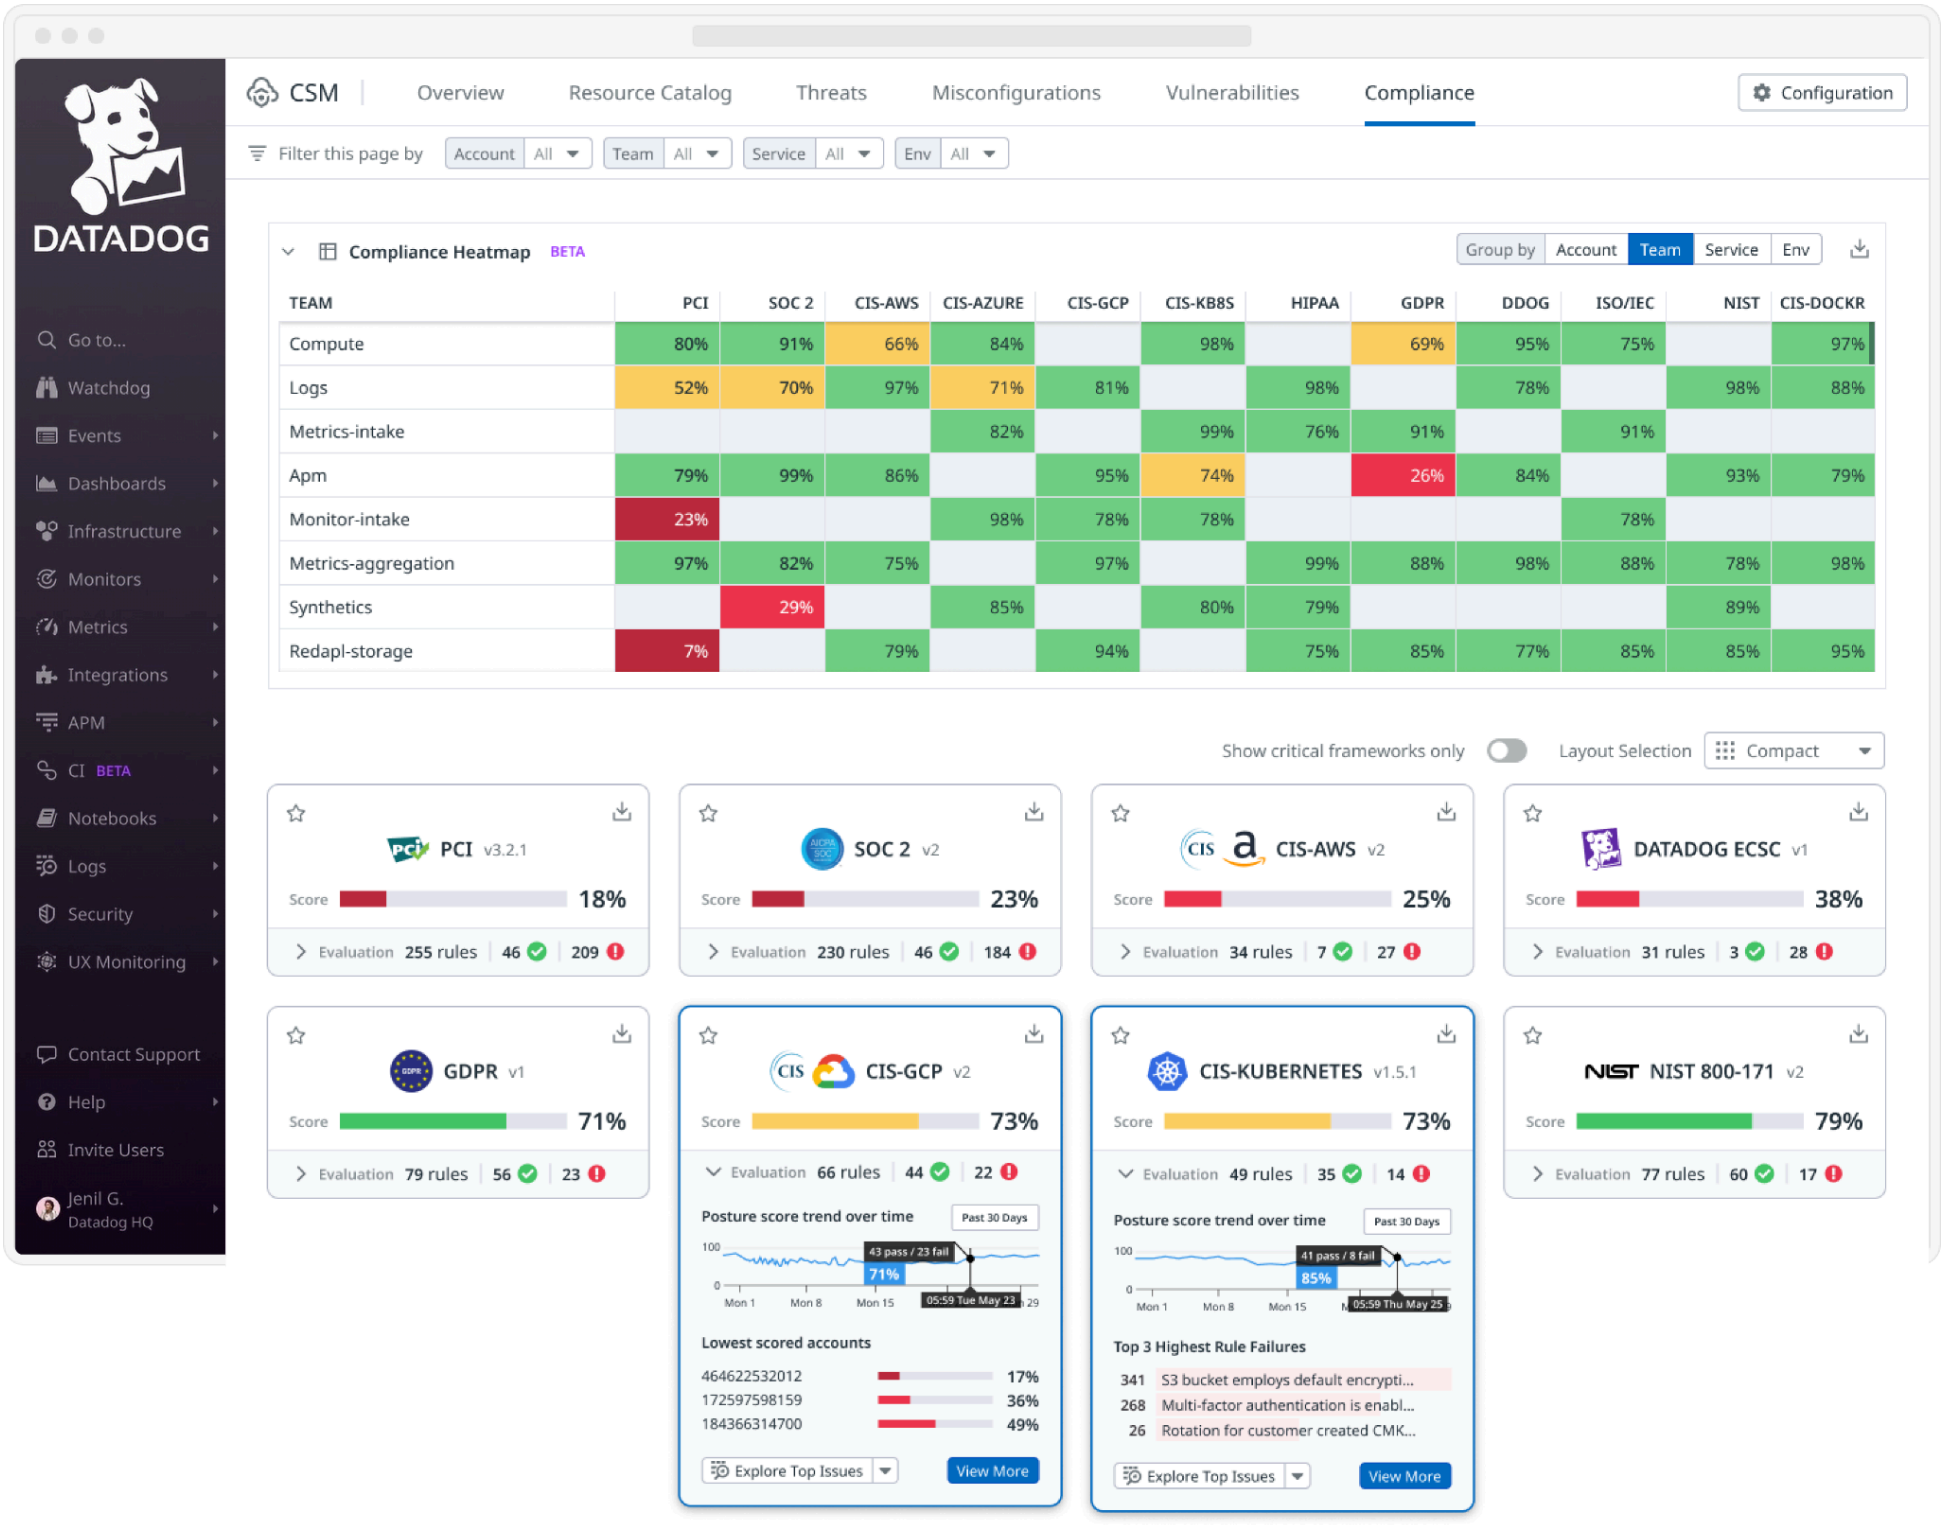
aws_ebs_snapshot	84.7K	23.4K

Compliance Framework Report page

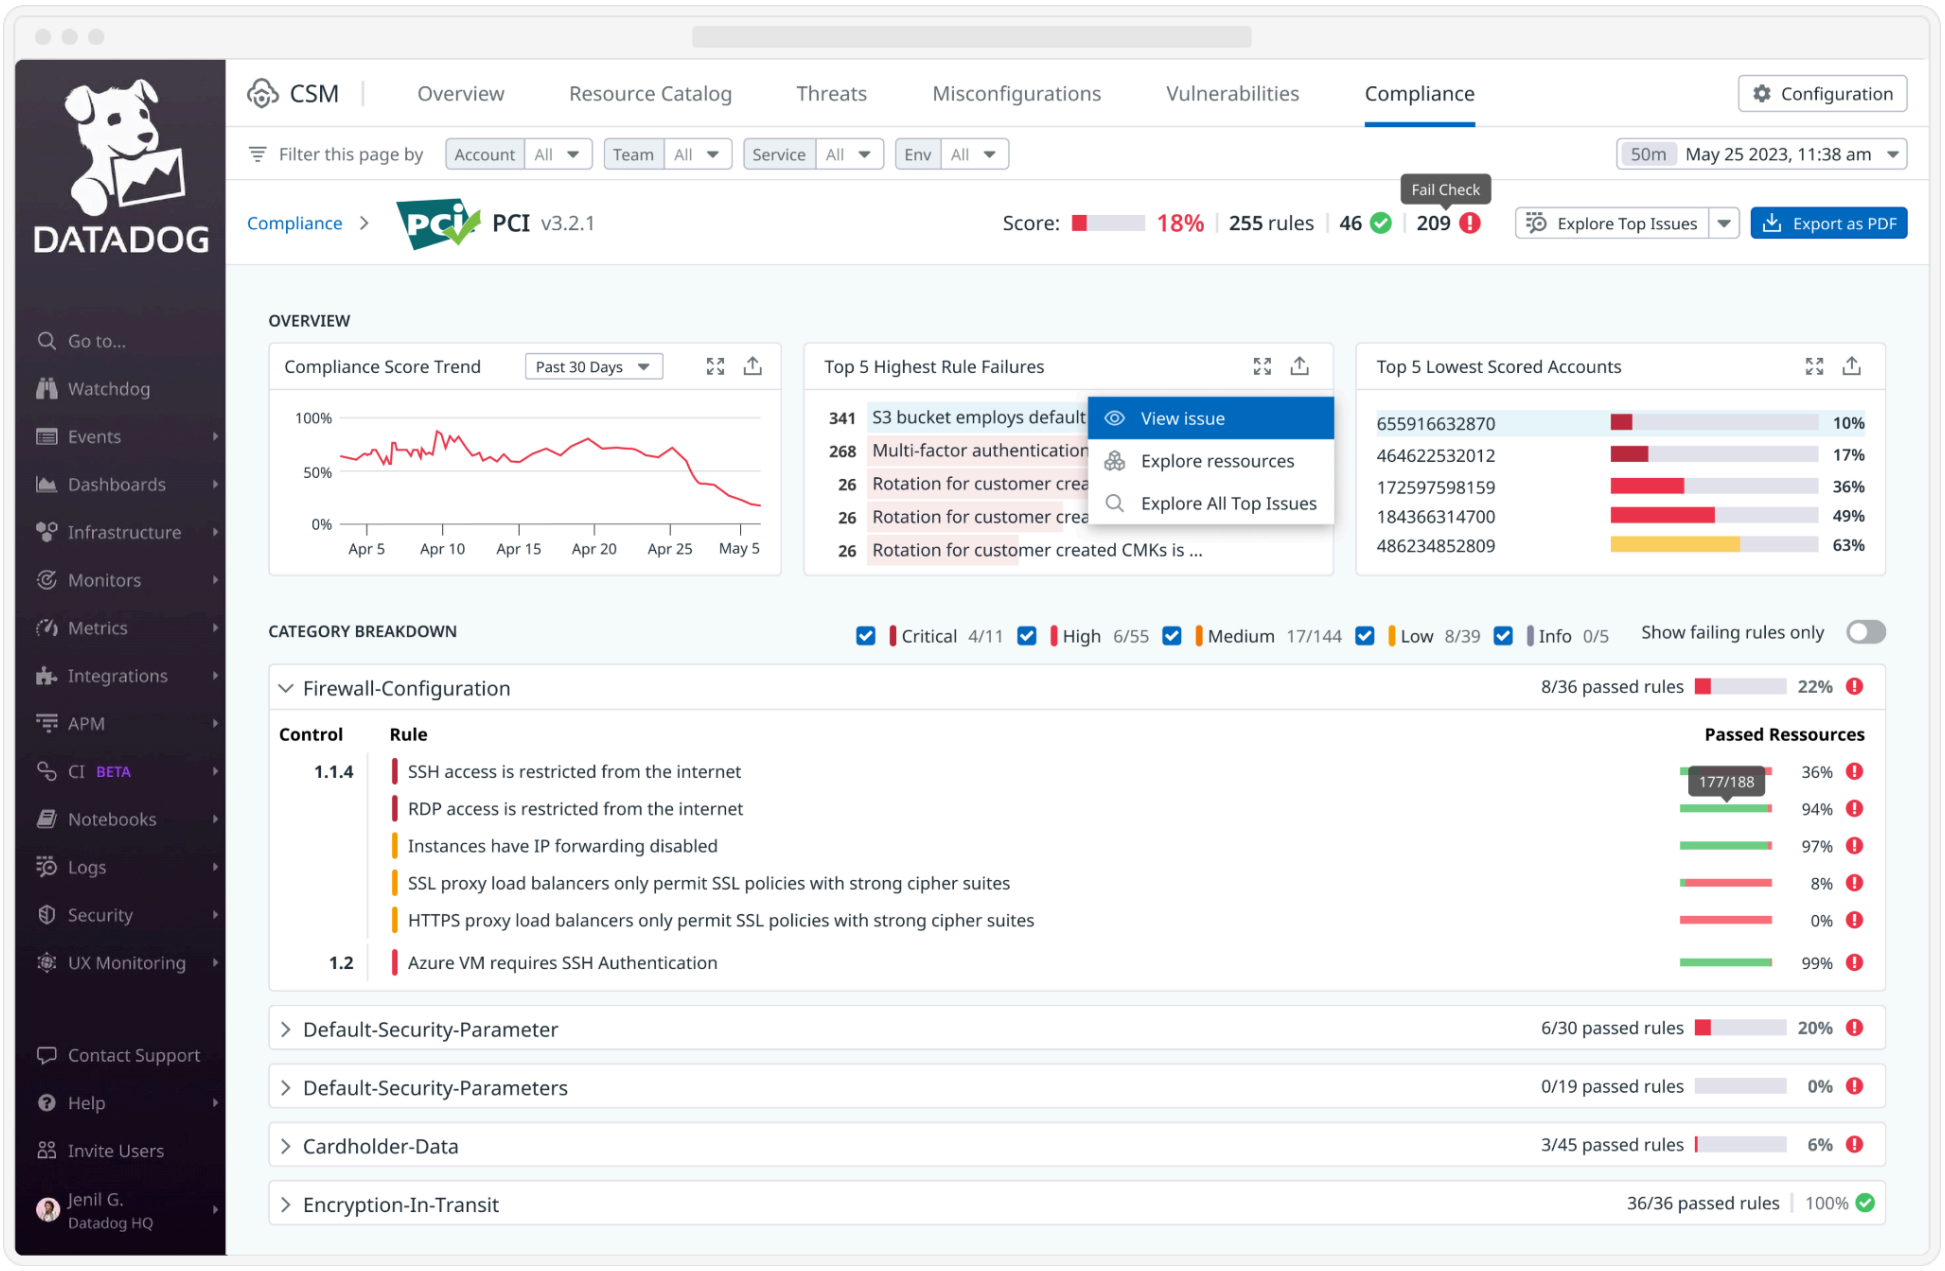


TO-BE → Actionable & Contextualized

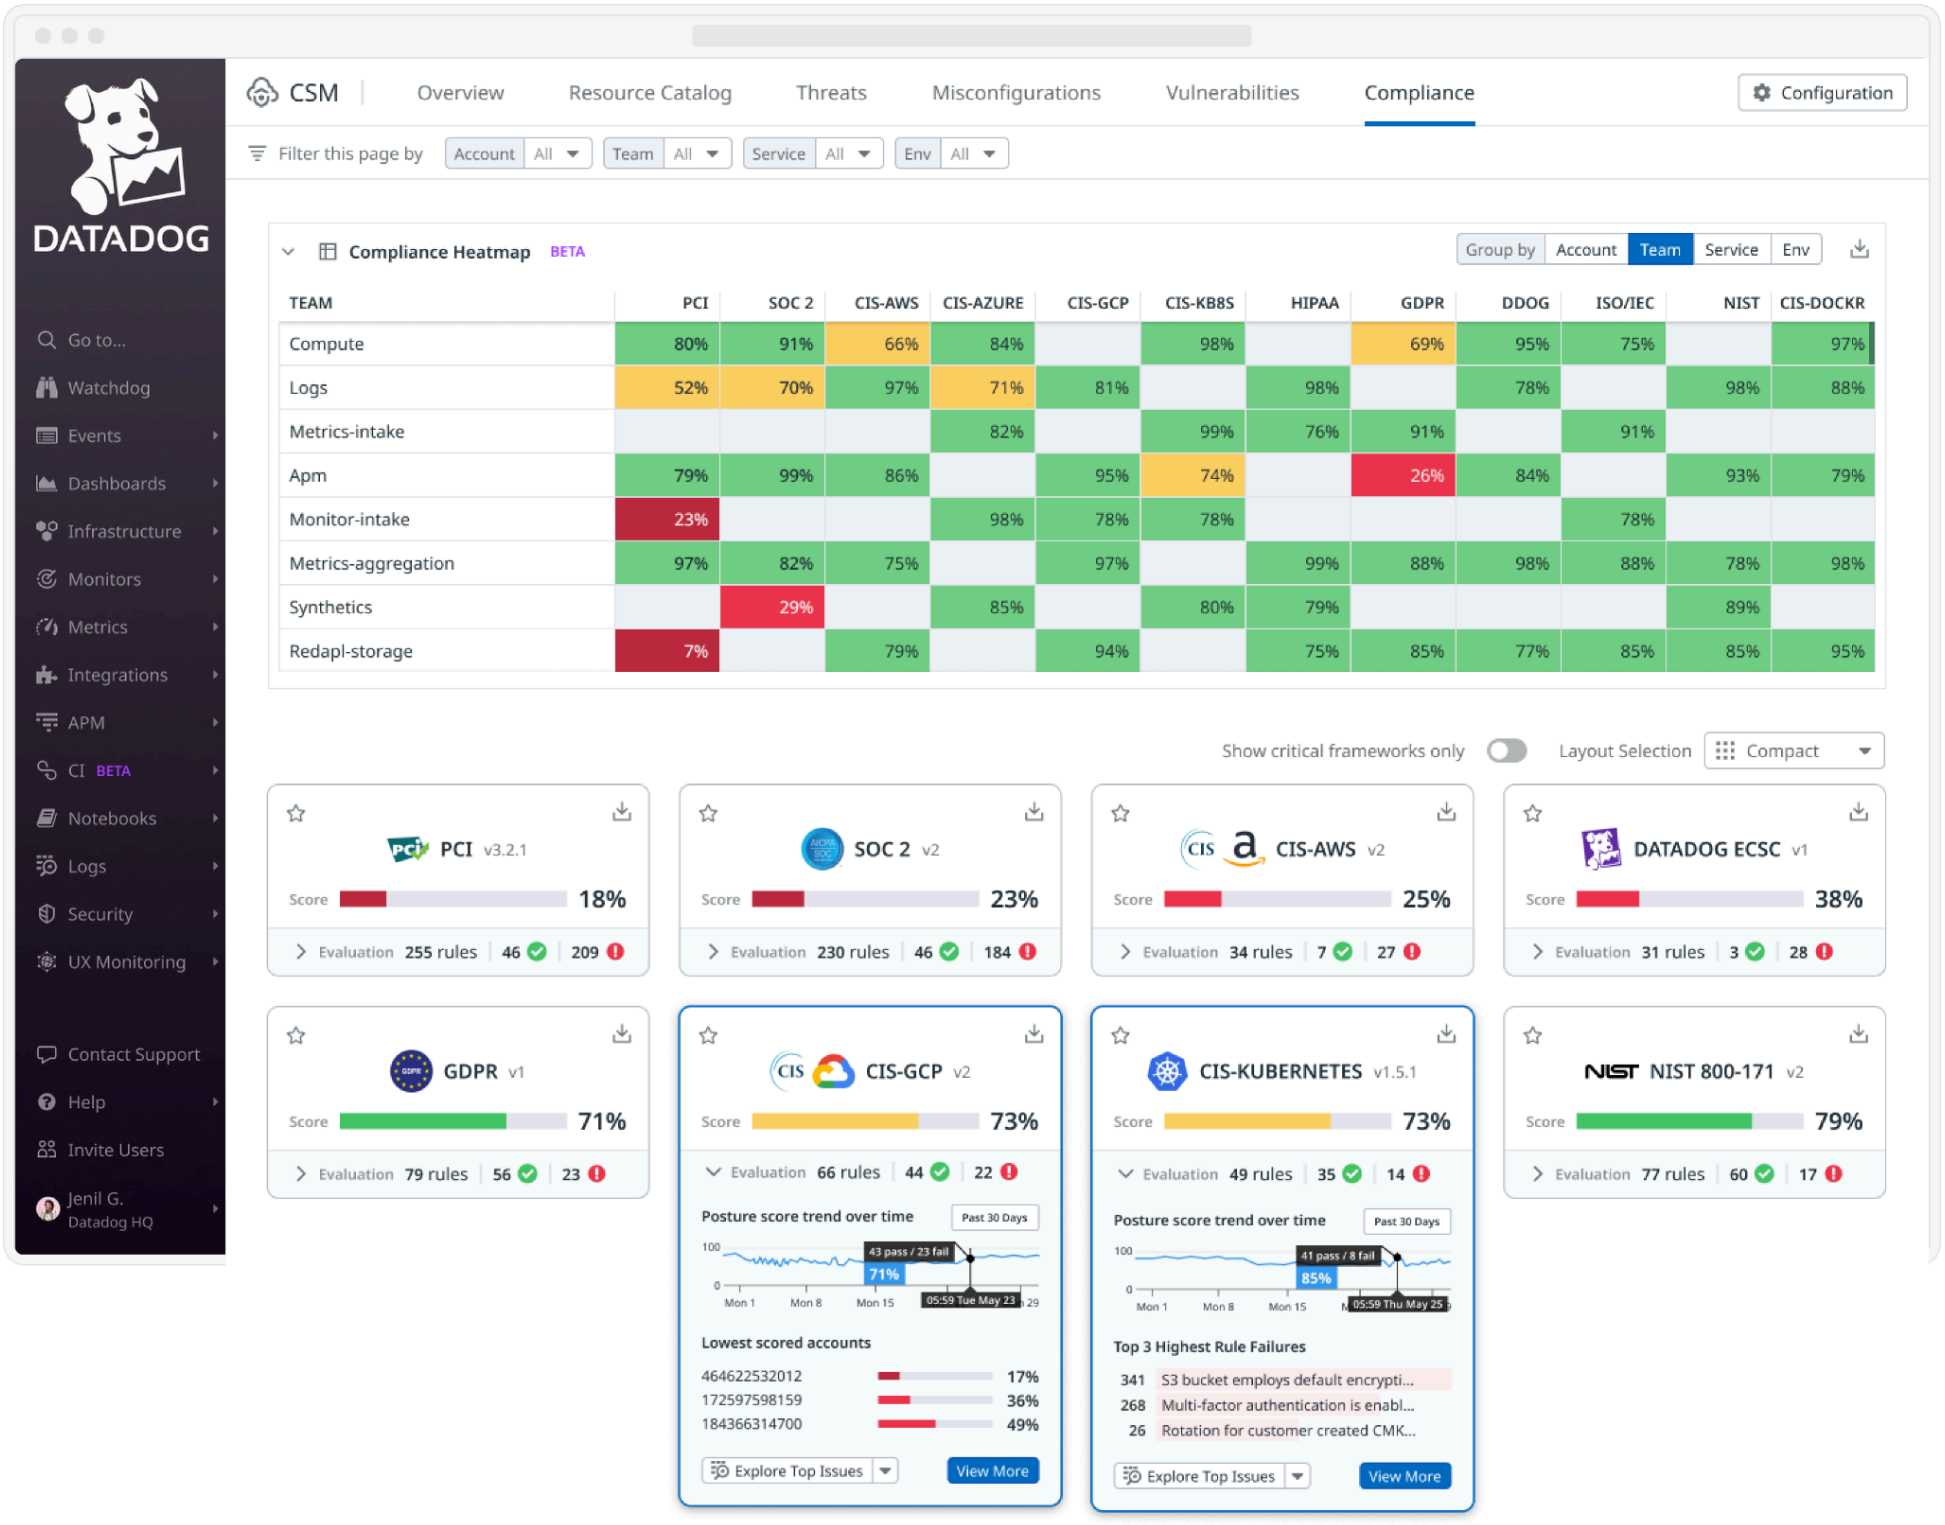
New Compliance Overview page



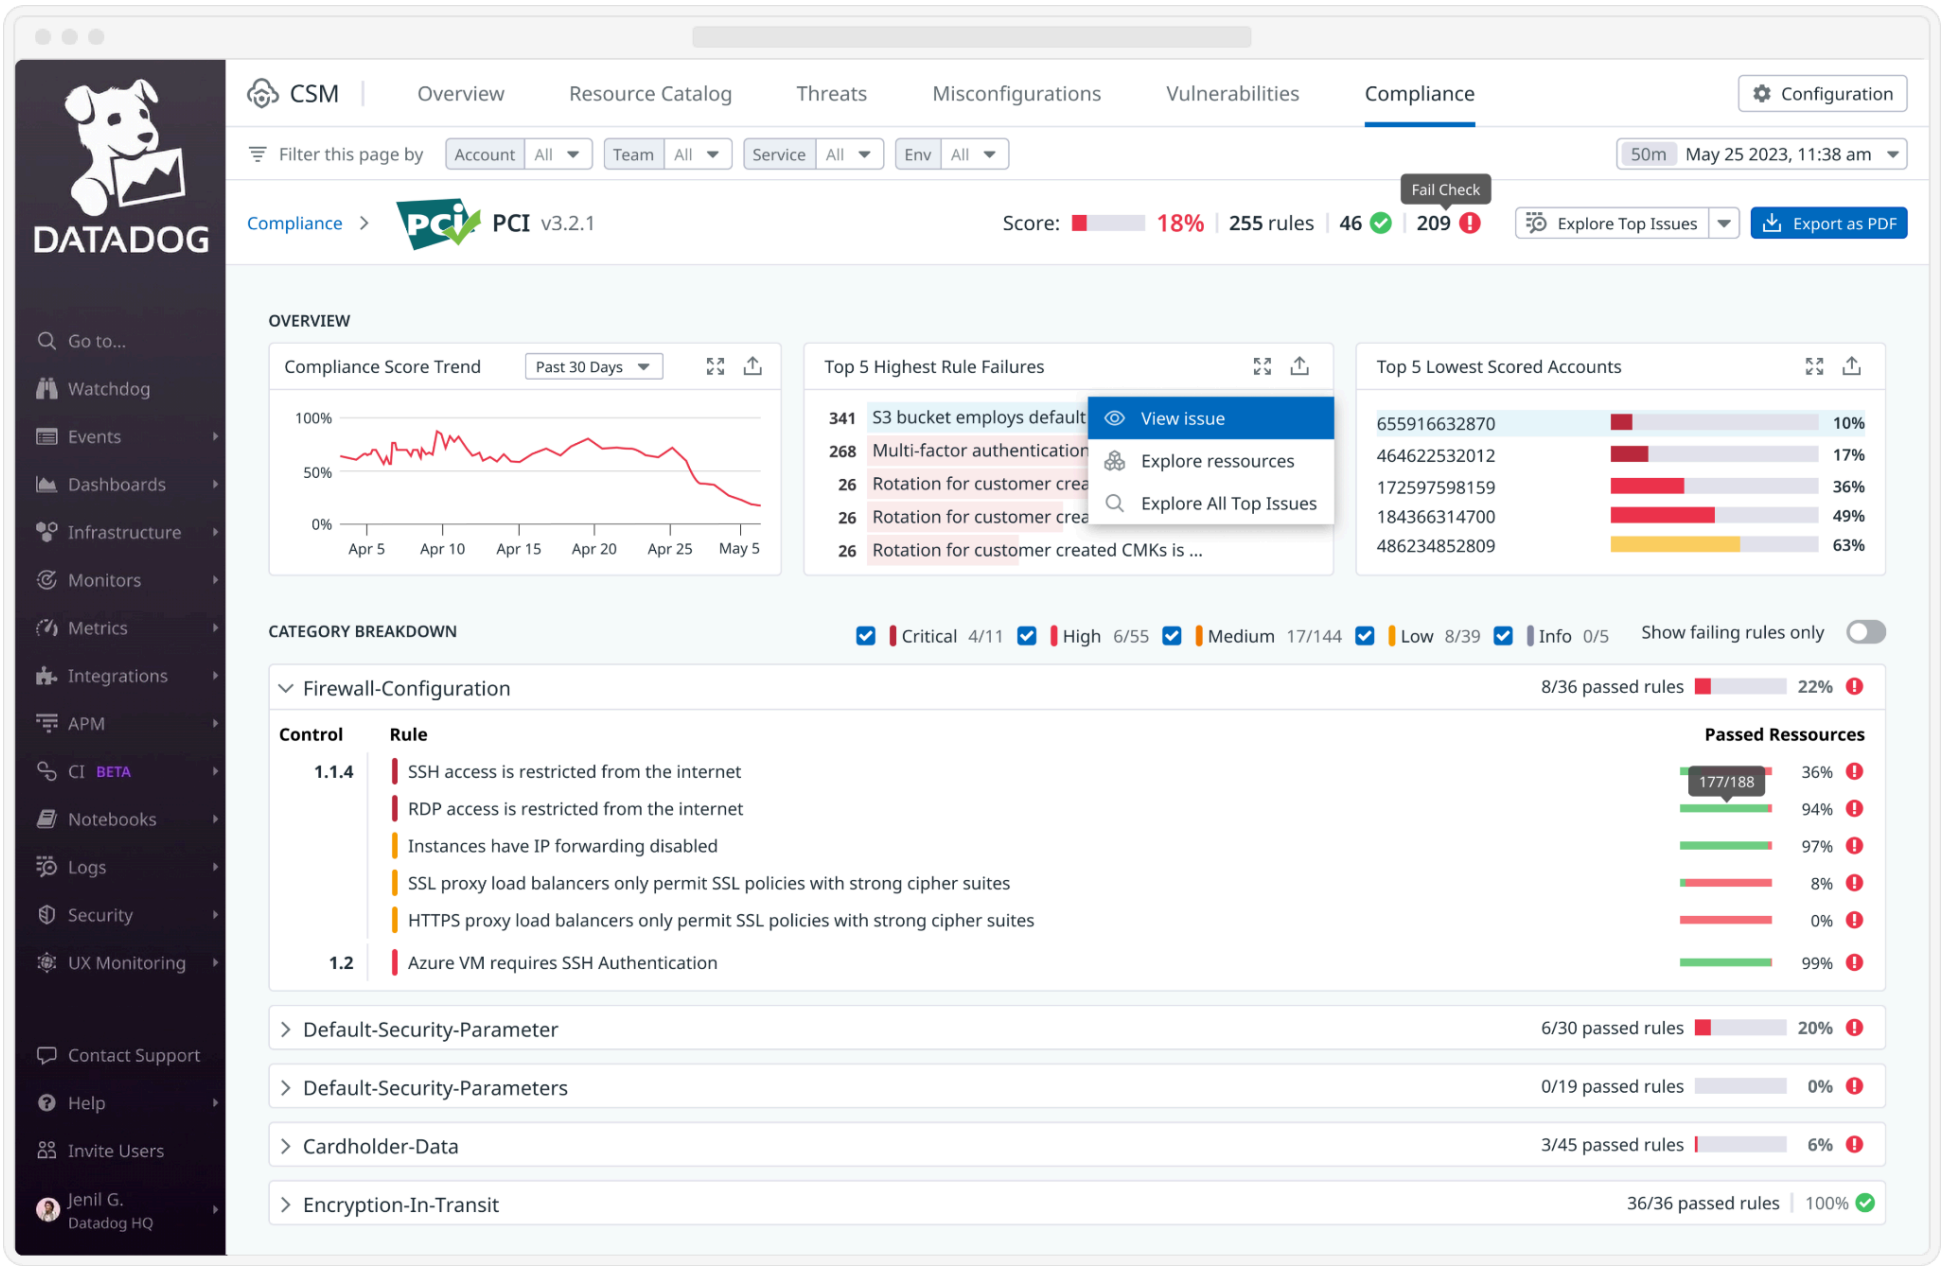
New Compliance Framework Report page



New Compliance Overview page



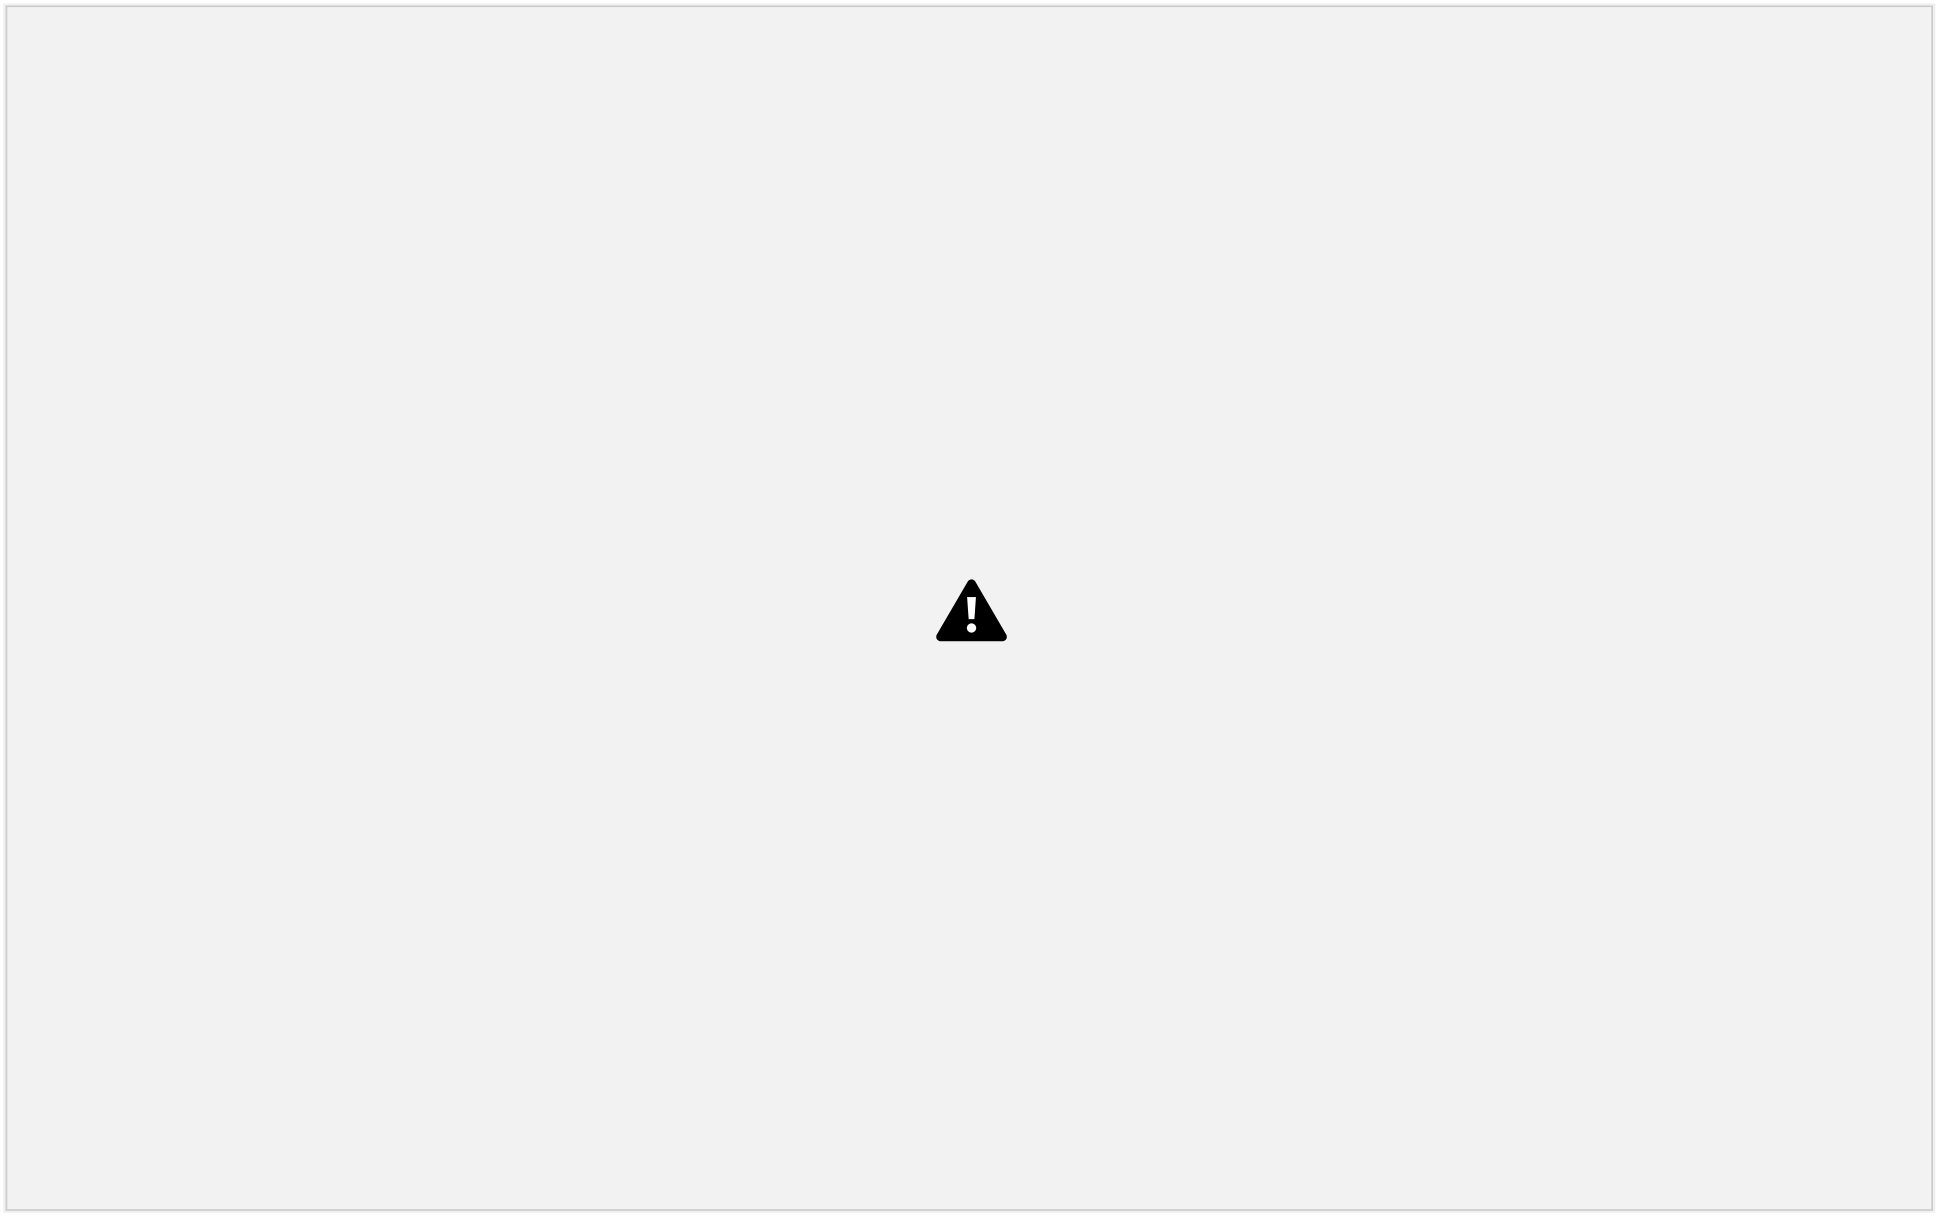
New Compliance Framework Report page



II/ Solutions

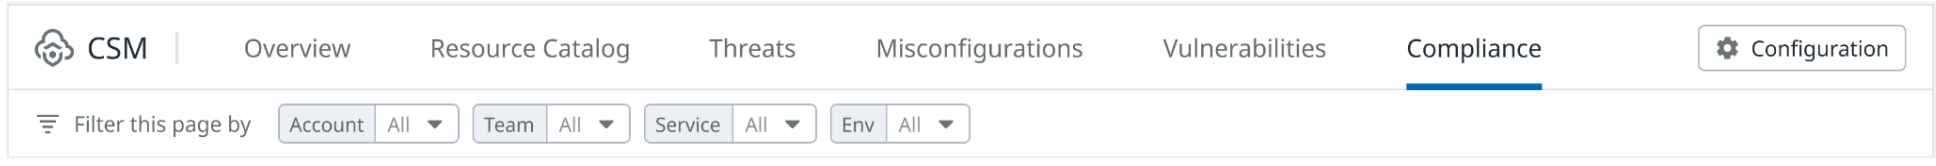
Compliance Posture At A Glance

Use Case Prototype

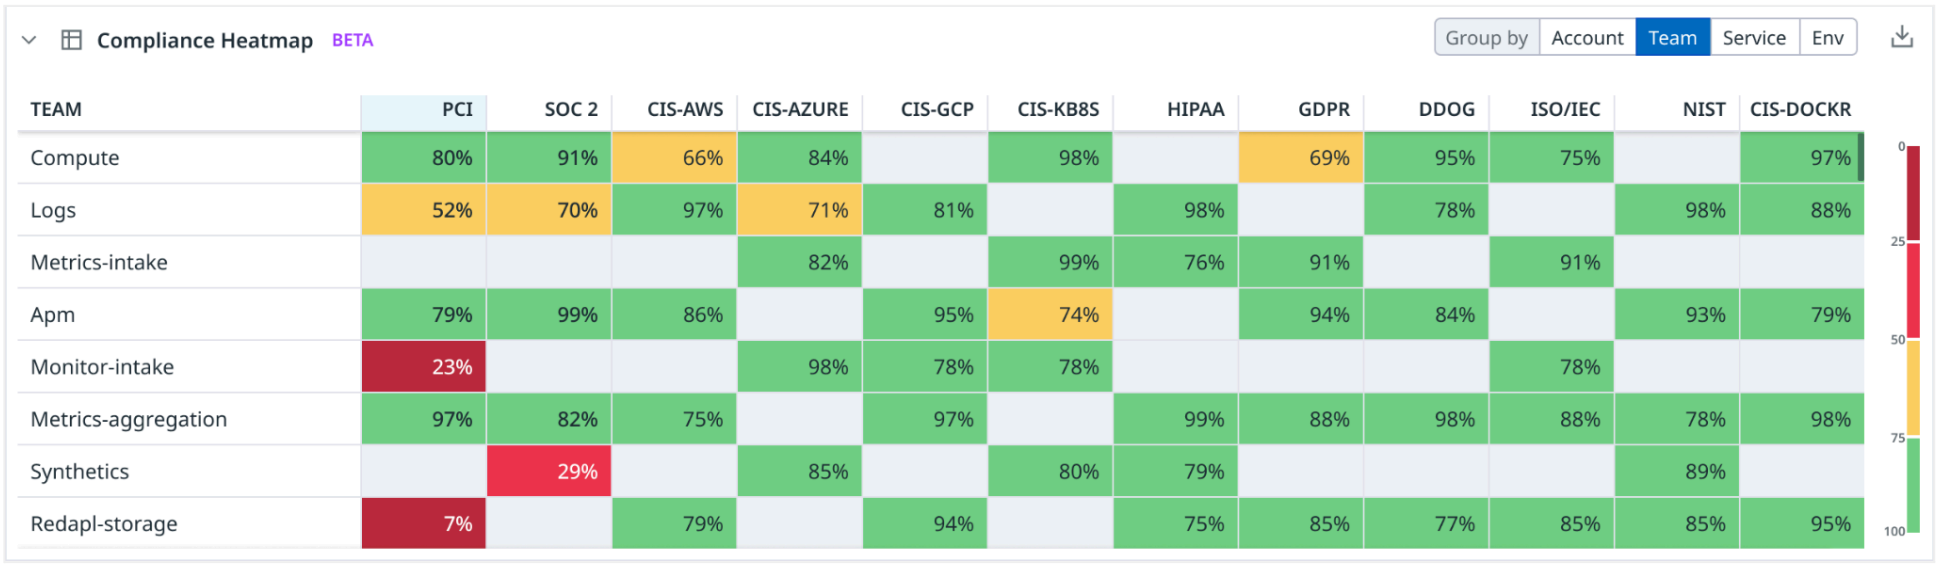


Solution

The top filter is inherited from the CSM Overview page and gives users more flexibility to scope what matters to them. Filter settings are stored per user to prevent any reconfiguration and allow a persistent experience across the compliance pages.

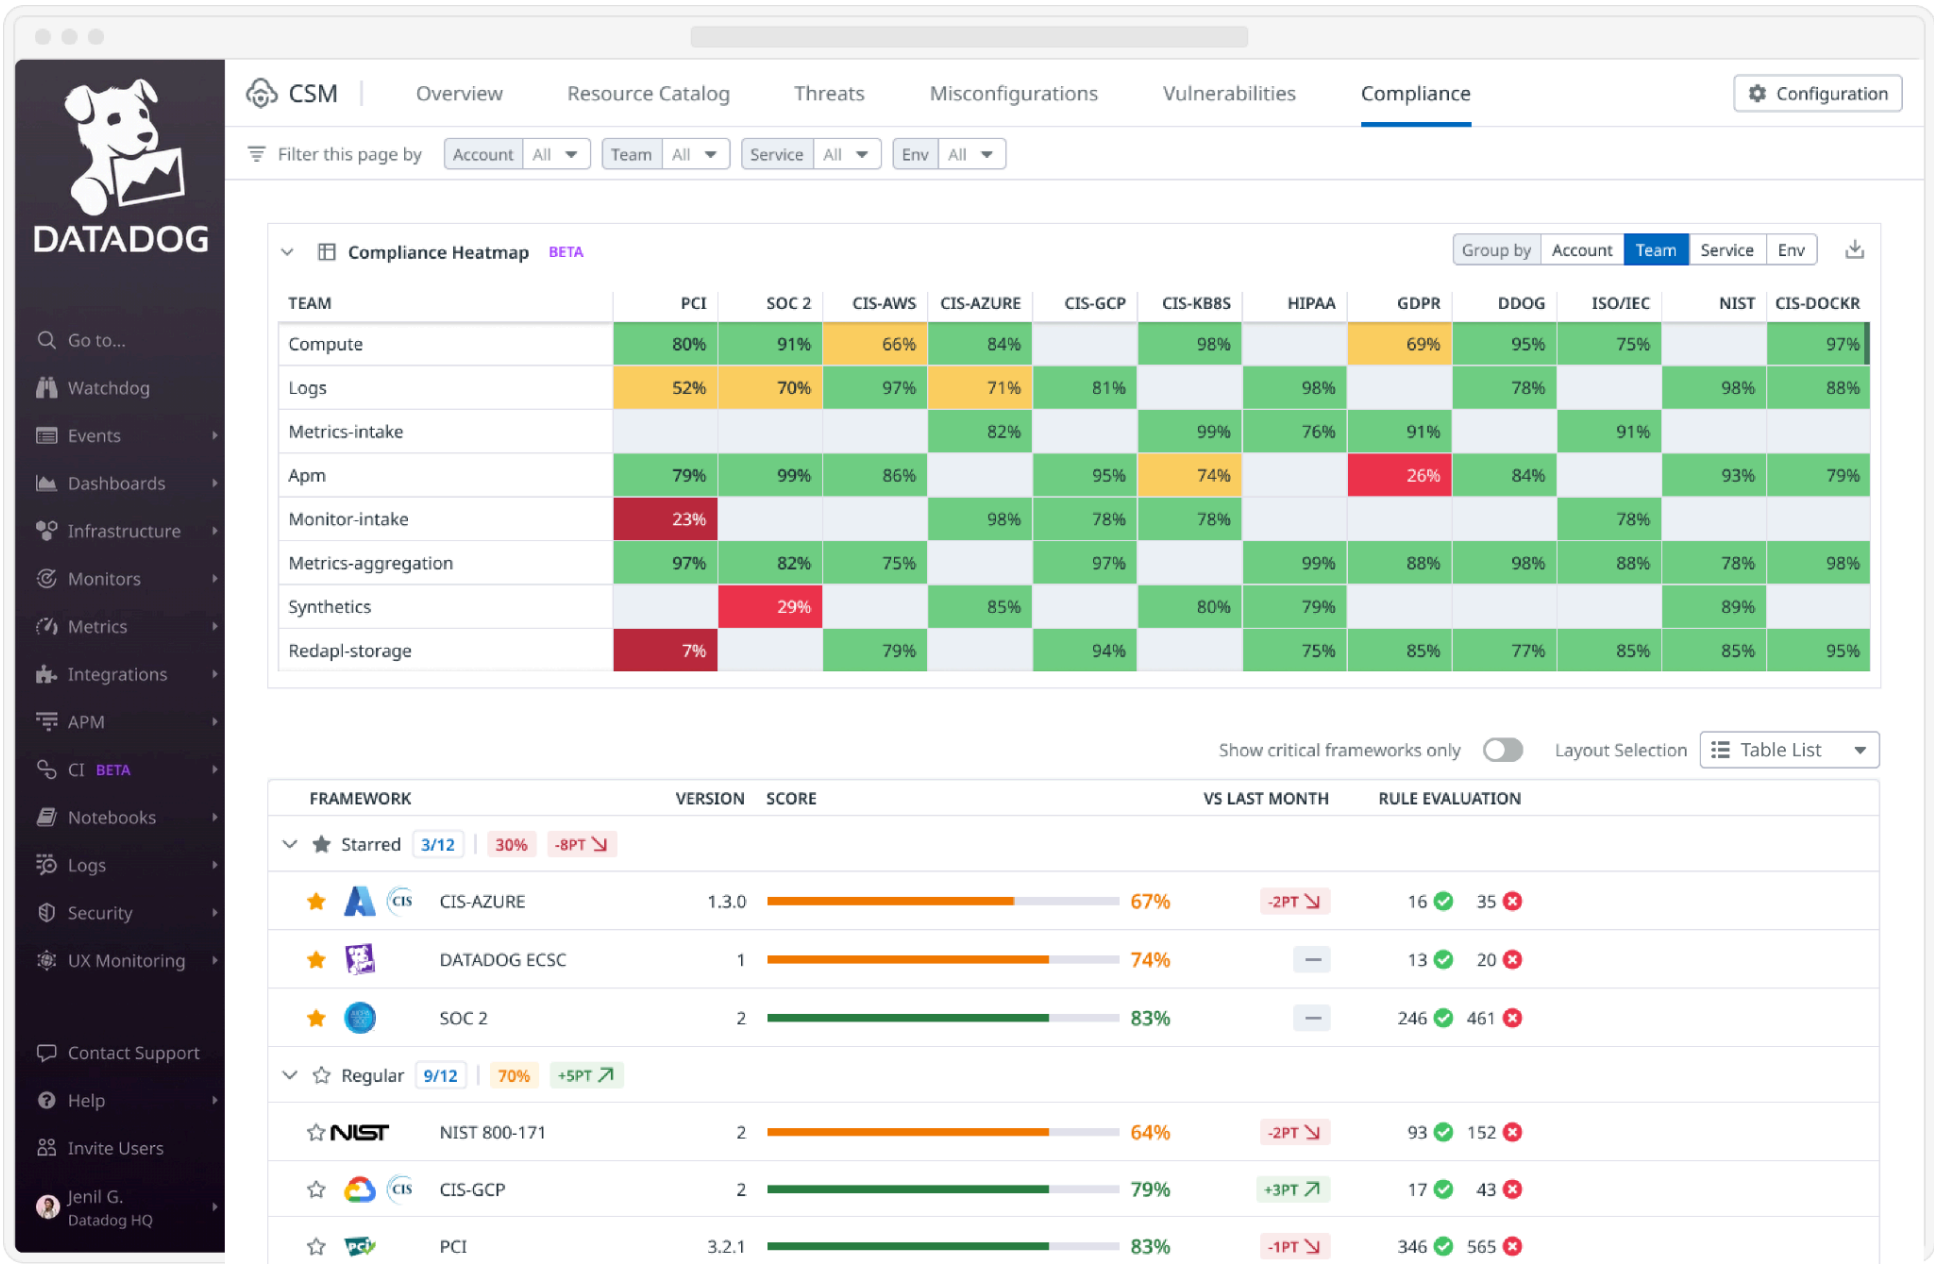


Providing users with a clear understanding of their compliance posture means allowing them to identify areas of improvement - at a glance - regarding a specific framework. That's why we introduce a compliance heatmap that helps them to understand where they stand about all frameworks quickly. The compliance heatmap will continuously assess cloud infrastructures against frameworks, and display posture scores matching their facet and "group by" selections. This way, we simplify their compliance visibility as it can be hard to understand how different fixes might improve their compliance posture.

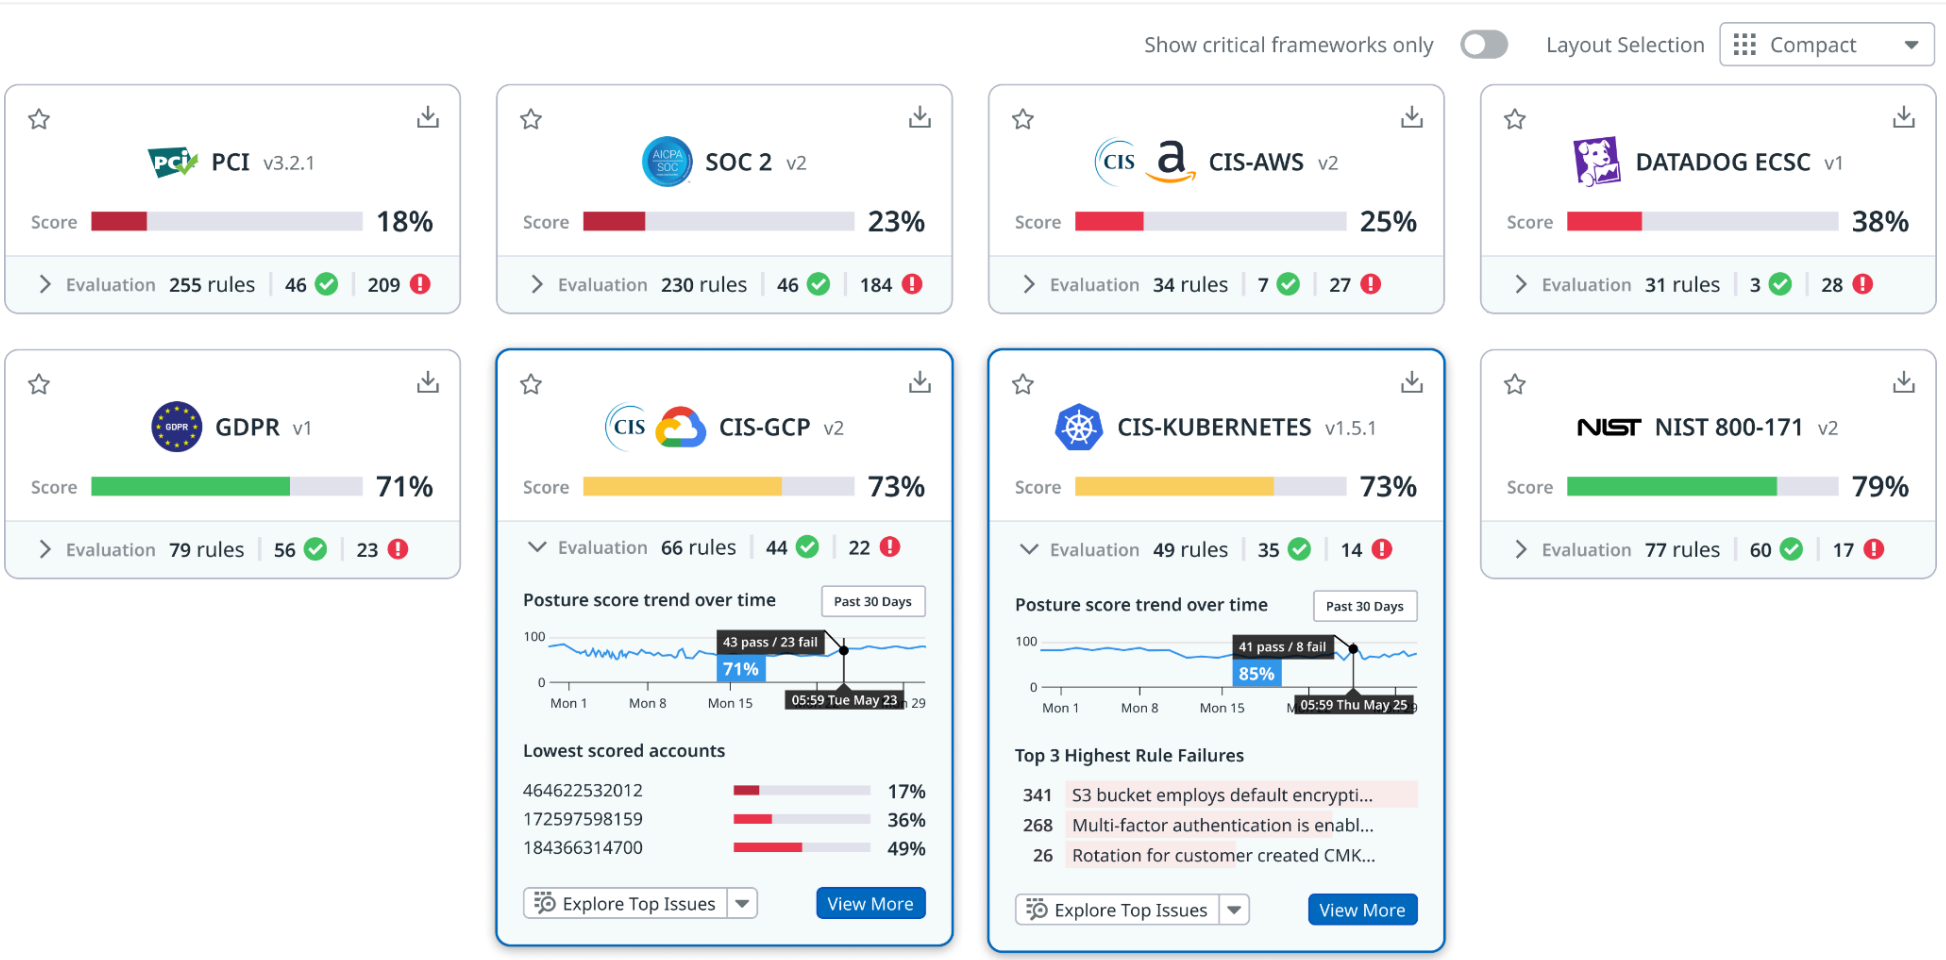


TEAM	PCI	SOC 2	CIS-AWS	CIS-AZURE	CIS-GCP	CIS-KB8S	HIPAA	GDPR	DDOG	ISO/IEC	NIST	CIS-DOCKR
Compute	80%	91%	66%	84%		98%		69%	95%	75%		97%
Logs	52%	70%	97%	71%	81%		98%		78%		98%	88%
Metrics-intake				82%		99%	76%	91%		91%		
Apm	79%	99%	86%		95%	74%		94%	84%		93%	79%
Monitor-intake	23%			98%	78%	78%				78%		
Metrics-aggregation	97%	82%	75%		97%		99%	88%	98%	88%	78%	98%
Synthetics		29%		85%		80%	79%				89%	
Redapl-storage	7%		79%		94%		75%	85%	77%	85%	85%	95%

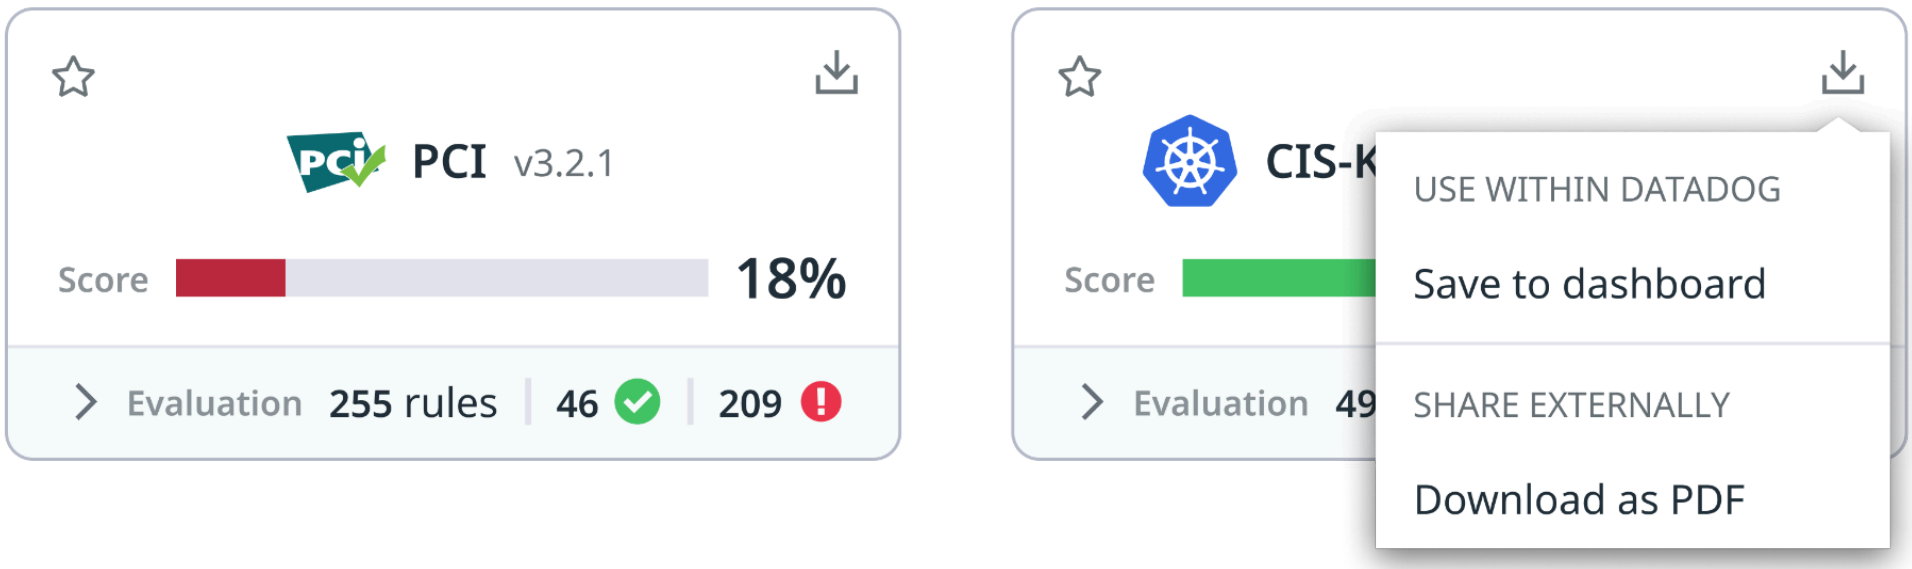
We’re also preventing users from scrolling over a long suite of compliance framework rows by rethinking how to display them efficiently. First, we choose a table list pattern as a standard view to ensure visual consistency between the compliance page and the misconfiguration explorer page and lighten the visual weight of framework rows. This Table List layout answers the call from some users to get a minimalistic overview of their compliance frameworks: they can pin /starre the frameworks they care the most at the top of the table and click on each line to access the dedicated compliance framework report page.



We complete the layout alternative pattern with an option allowing users to discover further compliance insights under framework cards. These expandable cards showcase just enough security context to help our users analyze what happened over time, compare compliance framework activities (as some compliance rules are shared between different compliance frameworks), prioritize what to tackle, and kickstart remediation. Their granular posture score under a percent bar enables fast readability and the rule evaluation section allows users to drill down into more data. They can also filter this card grid with a switch button to show only the frameworks with the lowest scores (under 50%).

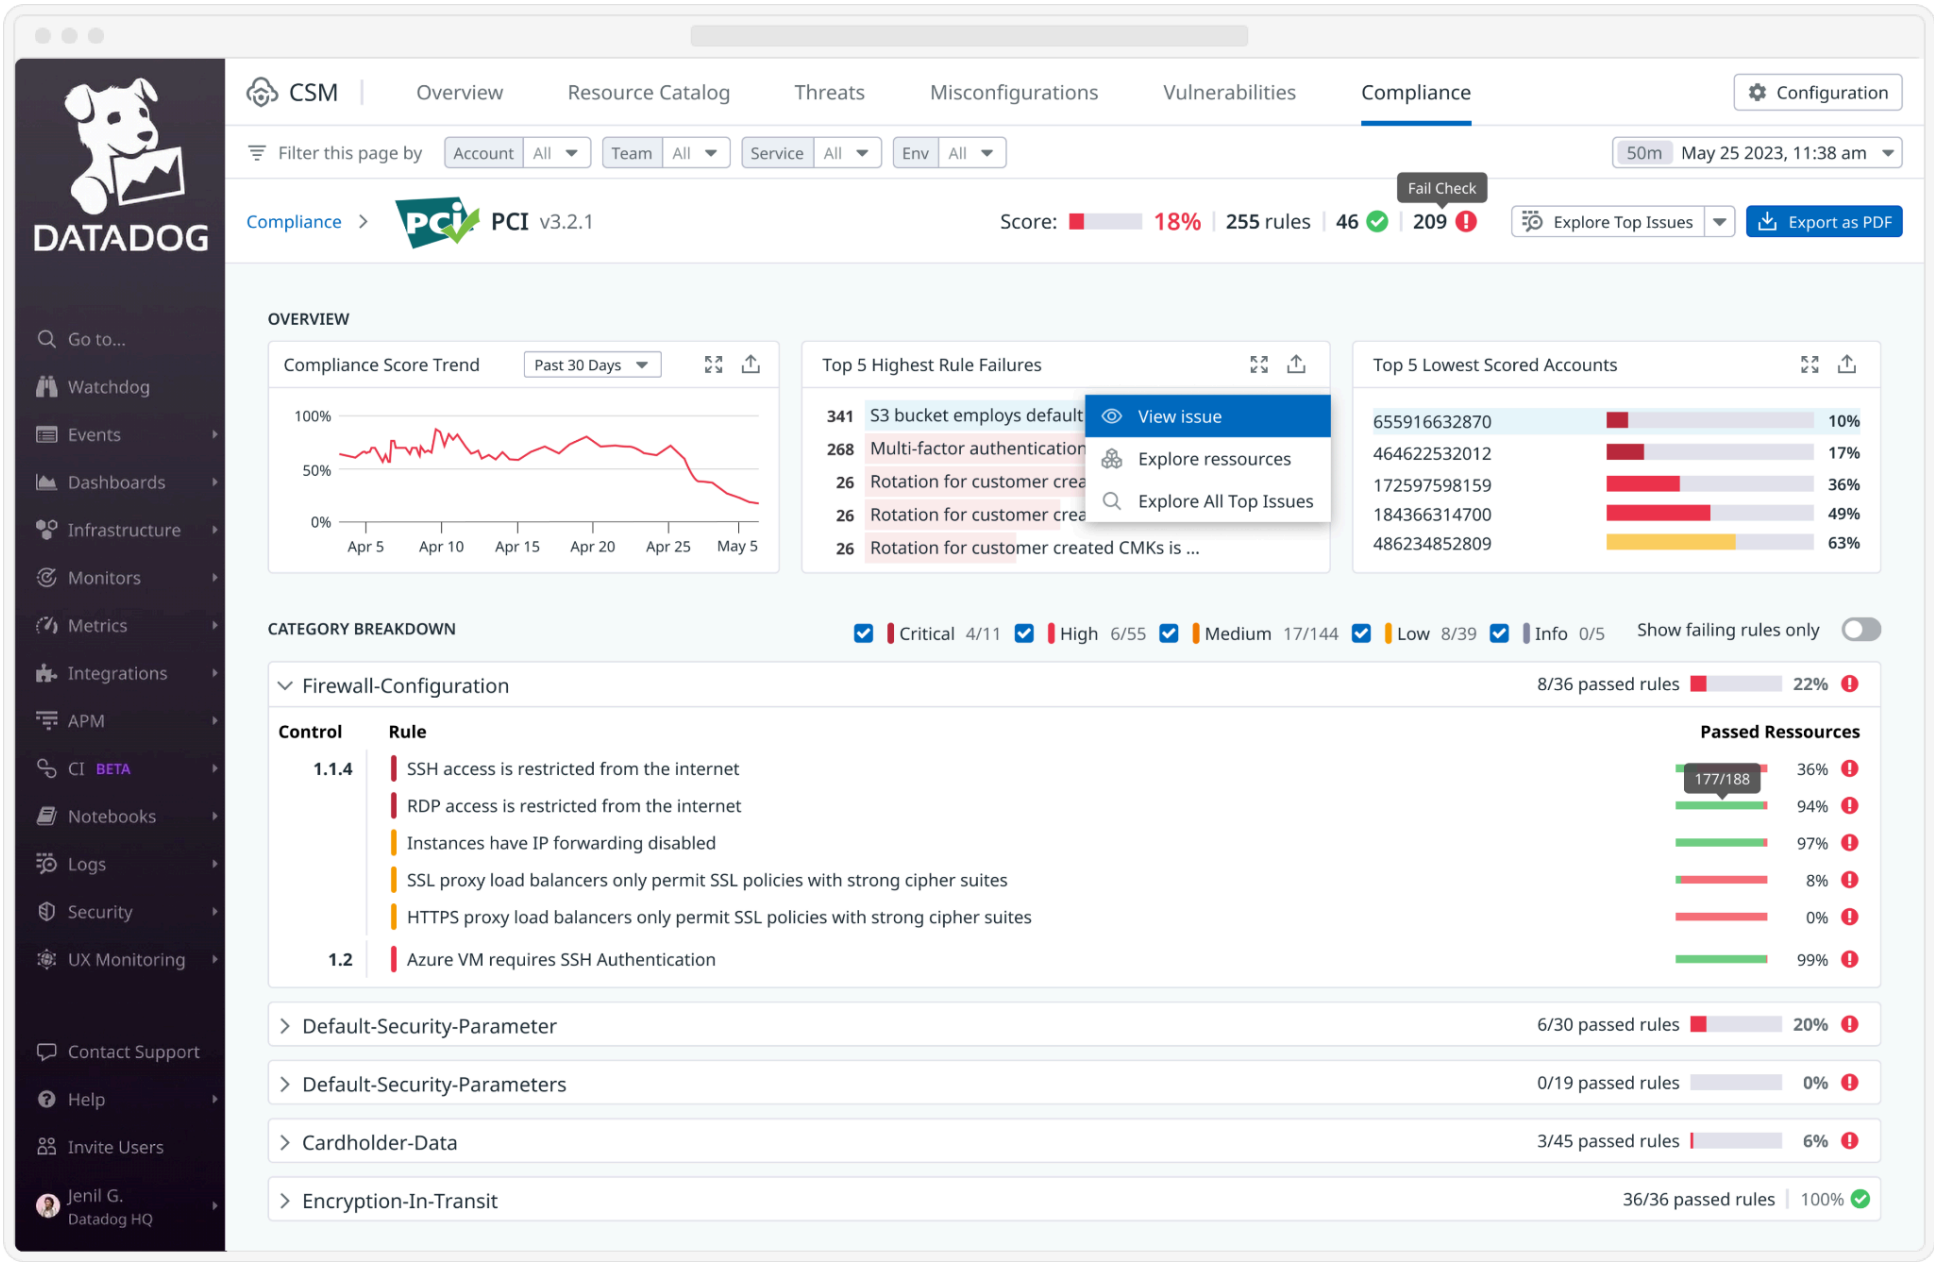


Currently, when preparing ahead of an audit, users struggle to generate proof of improvement over time for auditors. Taking screenshots is a common workaround that doesn't ease their reporting tasks. We are helping them by allowing a granular reporting approach through an "export" button placed in each widget of the compliance page and the framework page. Therefore, users can export the compliance heatmap and each framework card but also the whole framework page and each overview card. This "Export" button reveals the most useful reporting actions: "Save to a dashboard" and "Download as a PDF".



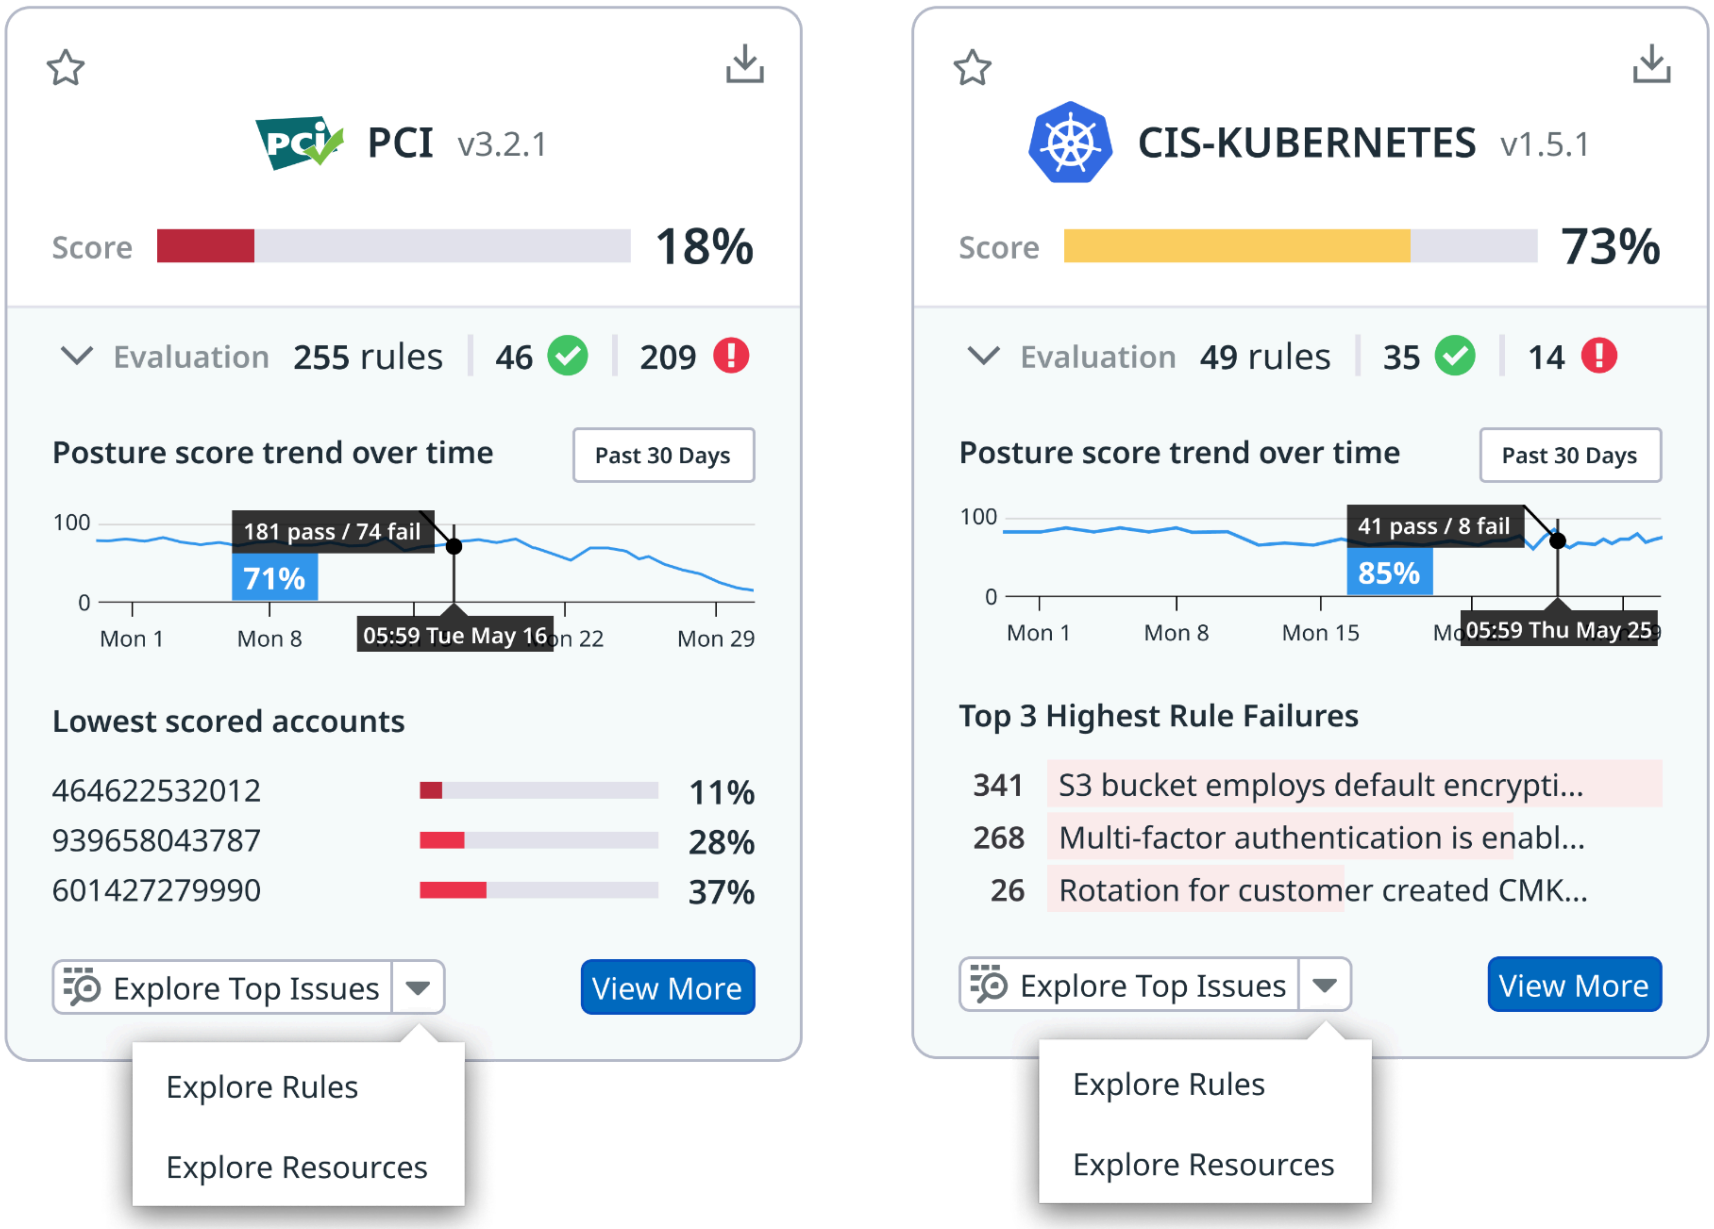
Actionability For Faster Prioritization/Remediation

Use Case Prototype

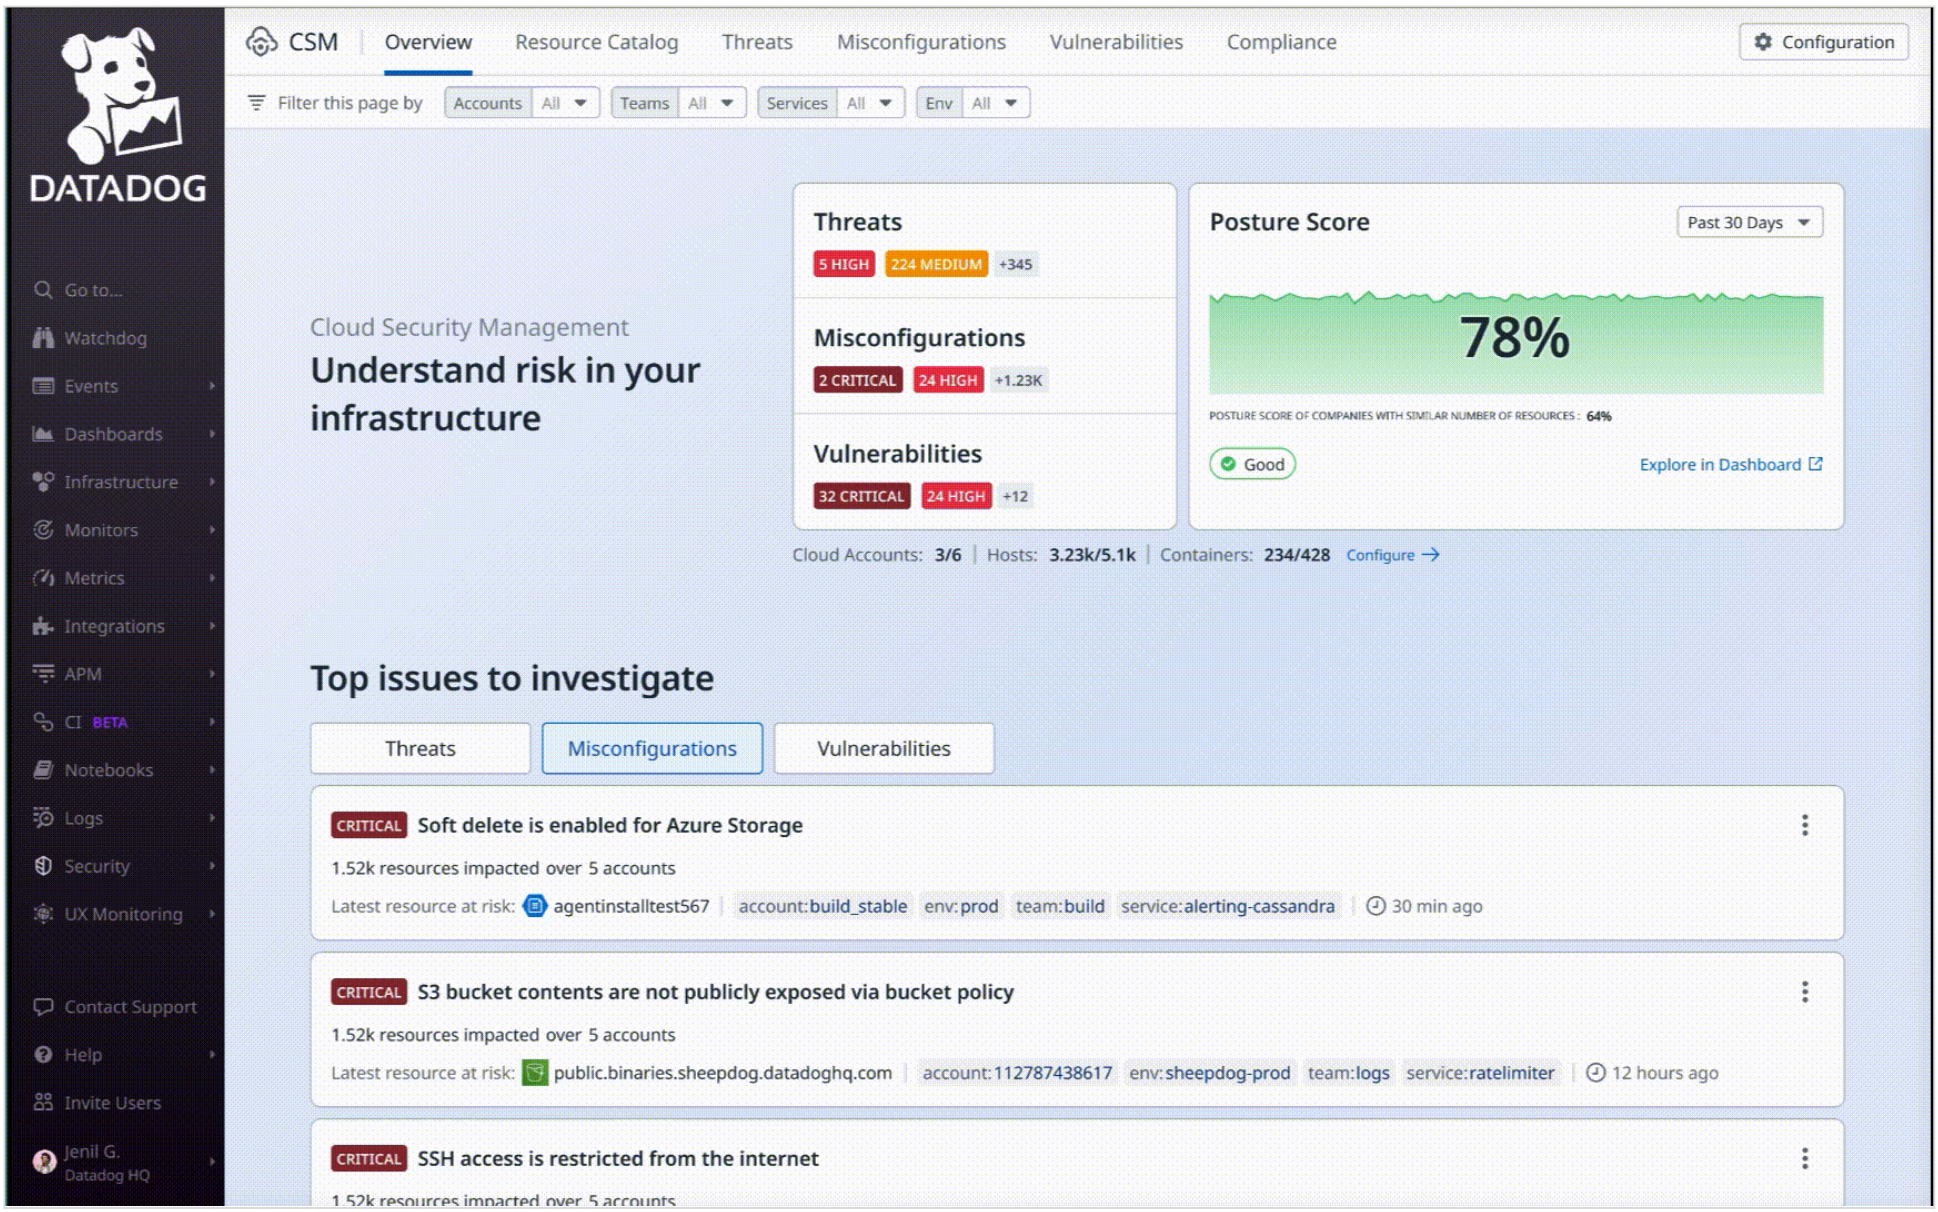


Solution

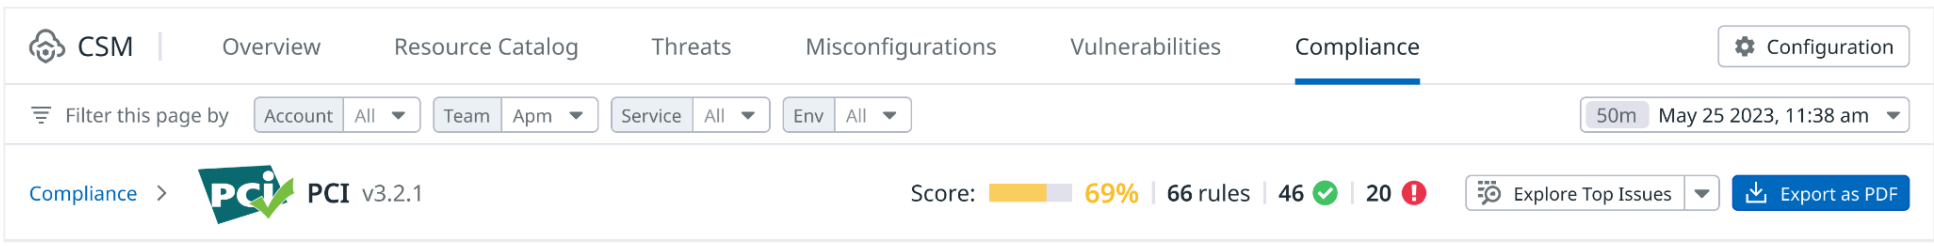
From a prioritization standpoint, the new compliance experience makes it easy for users to investigate critical compliance spots. At the framework card level (in the compliance page) the score trend over time and the most pressing issues (under a dynamic top list) give a sense of prioritization. A split button gathers 3 investigation actions : “Explore top issues”, “Explore Rules”, “Explore Resources" allowing users to tackle remediation from the compliance page and to jump into a prefiltered explorer.

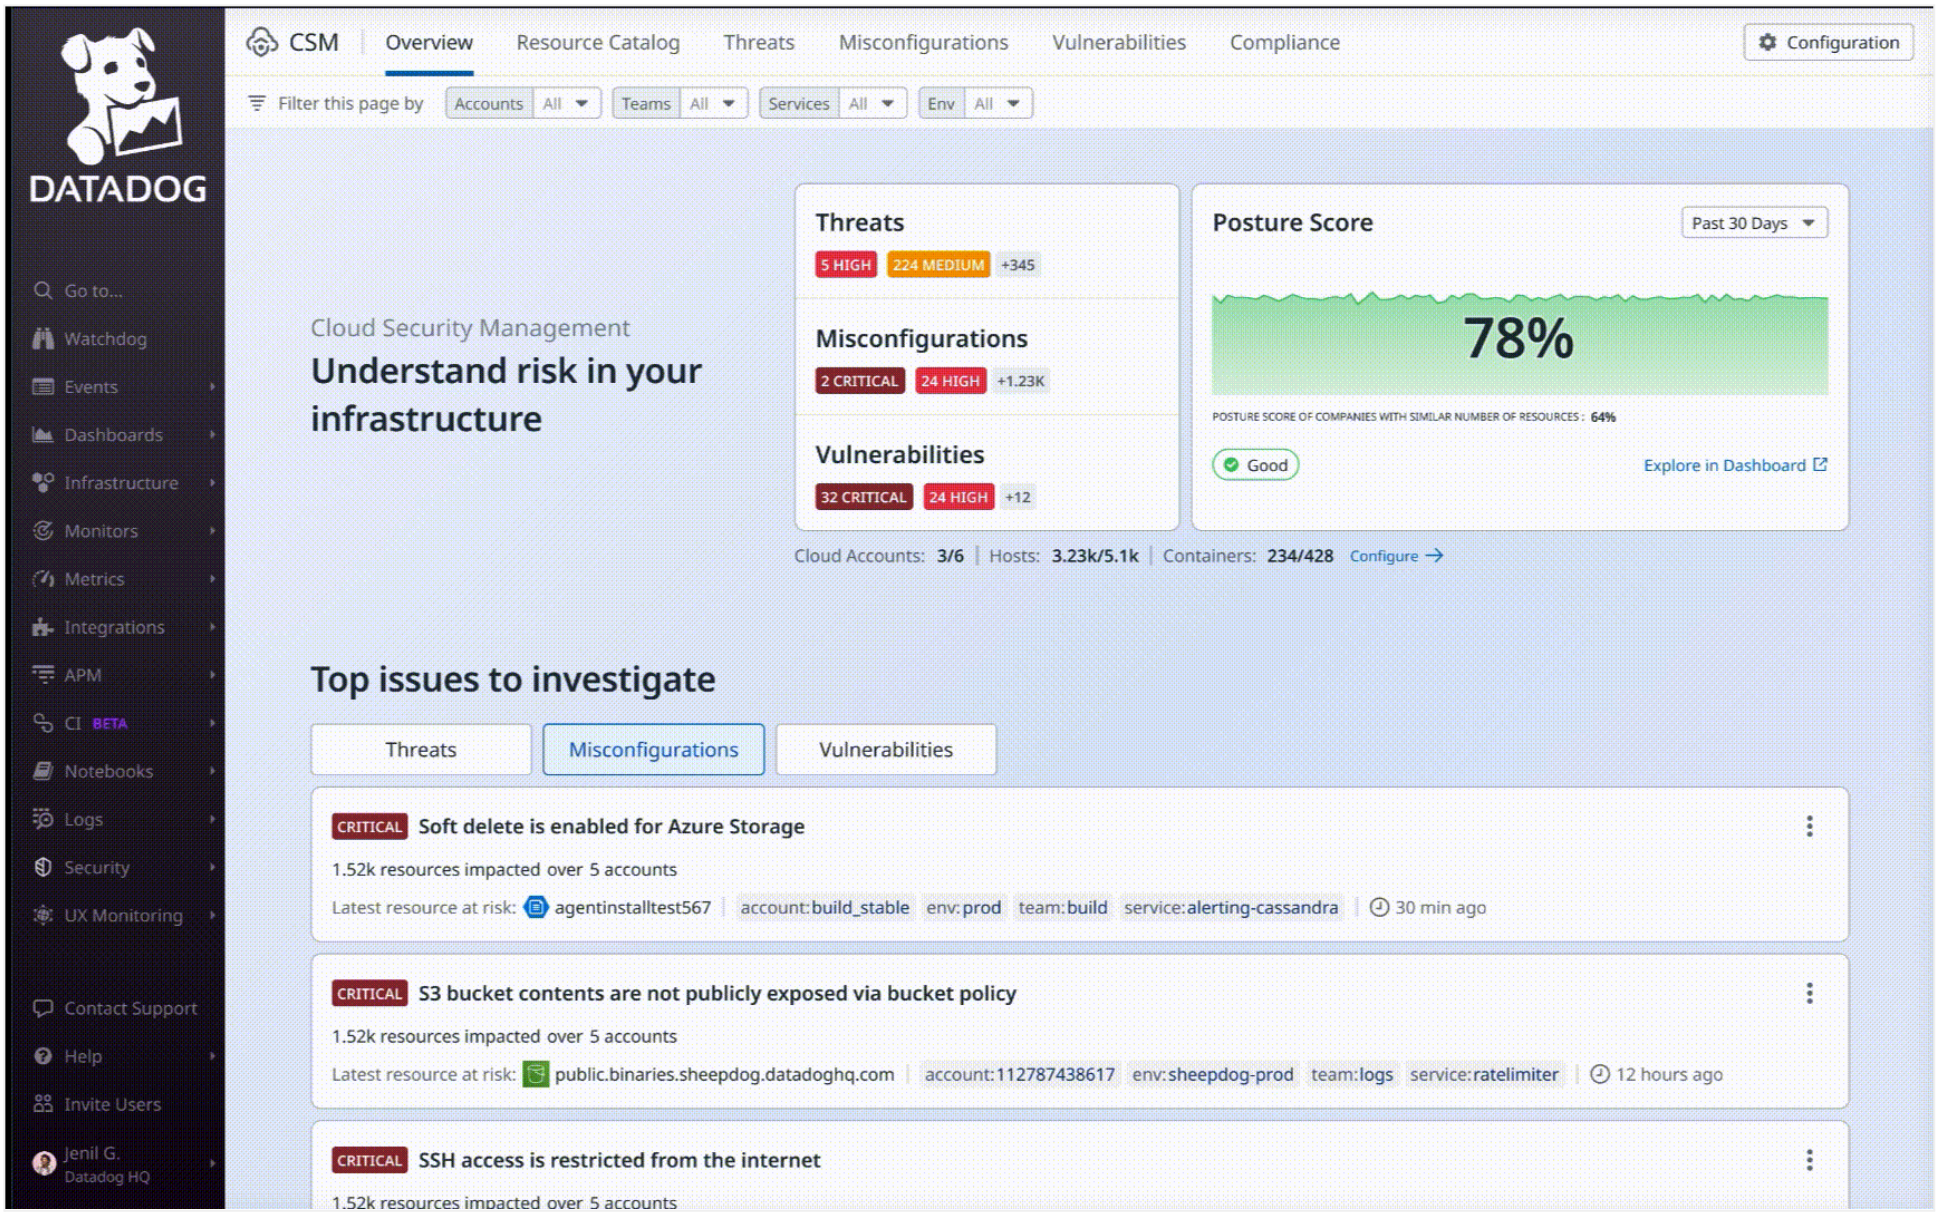


From the compliance heatmap, users can access a pre-filtered view of the matching framework page by clicking on a critical cell. The attributes of the targeted cell populate the facets of the inherited top filter and automatically adapt all the framework page's data (posture score, rule evaluation, score trend, top issue lists). Furthermore, users can reveal all the failing rules with a click.

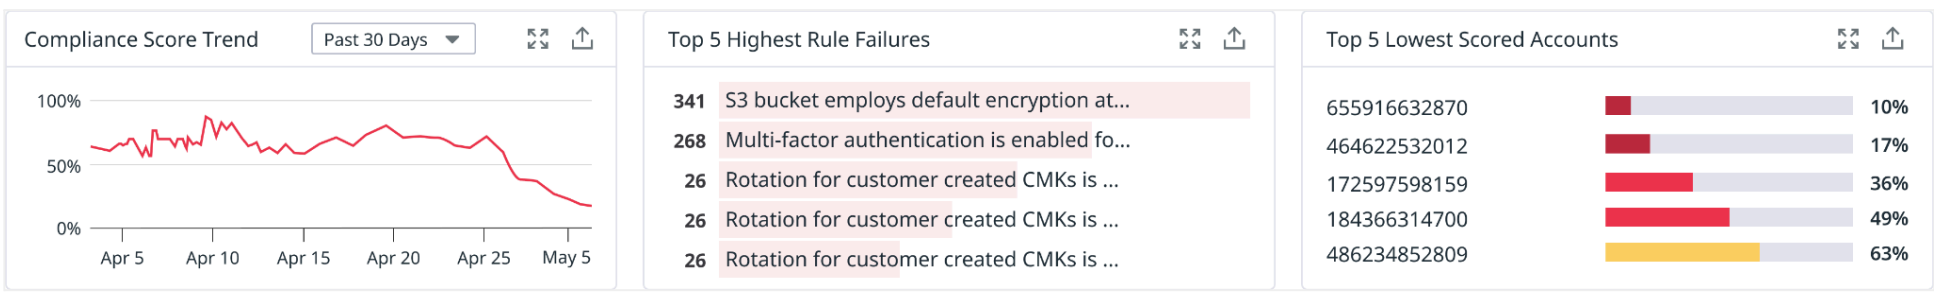


Now, users can easily scope their investigation and remediation efforts, from the compliance overview page to the compliance framework report page. The header gathers the inherited top filter of the compliance page (completed by a timeframe selector) and framework overview details (posture score, number of rules, rule evaluation). From the compliance overview page, users can enter their scope attributes in the top filter, which will be reused to adapt the data of the framework page. The same split button included on each framework card of the compliance page enables users to investigate further top issues, rules and resources. Regarding the export options, the only option mentioned is downloading as a PDF, which may not be sufficient for further analysis. However, we initiated discussions with engineering and are pushing to offer multiple export options, including CSV / Excel formats.

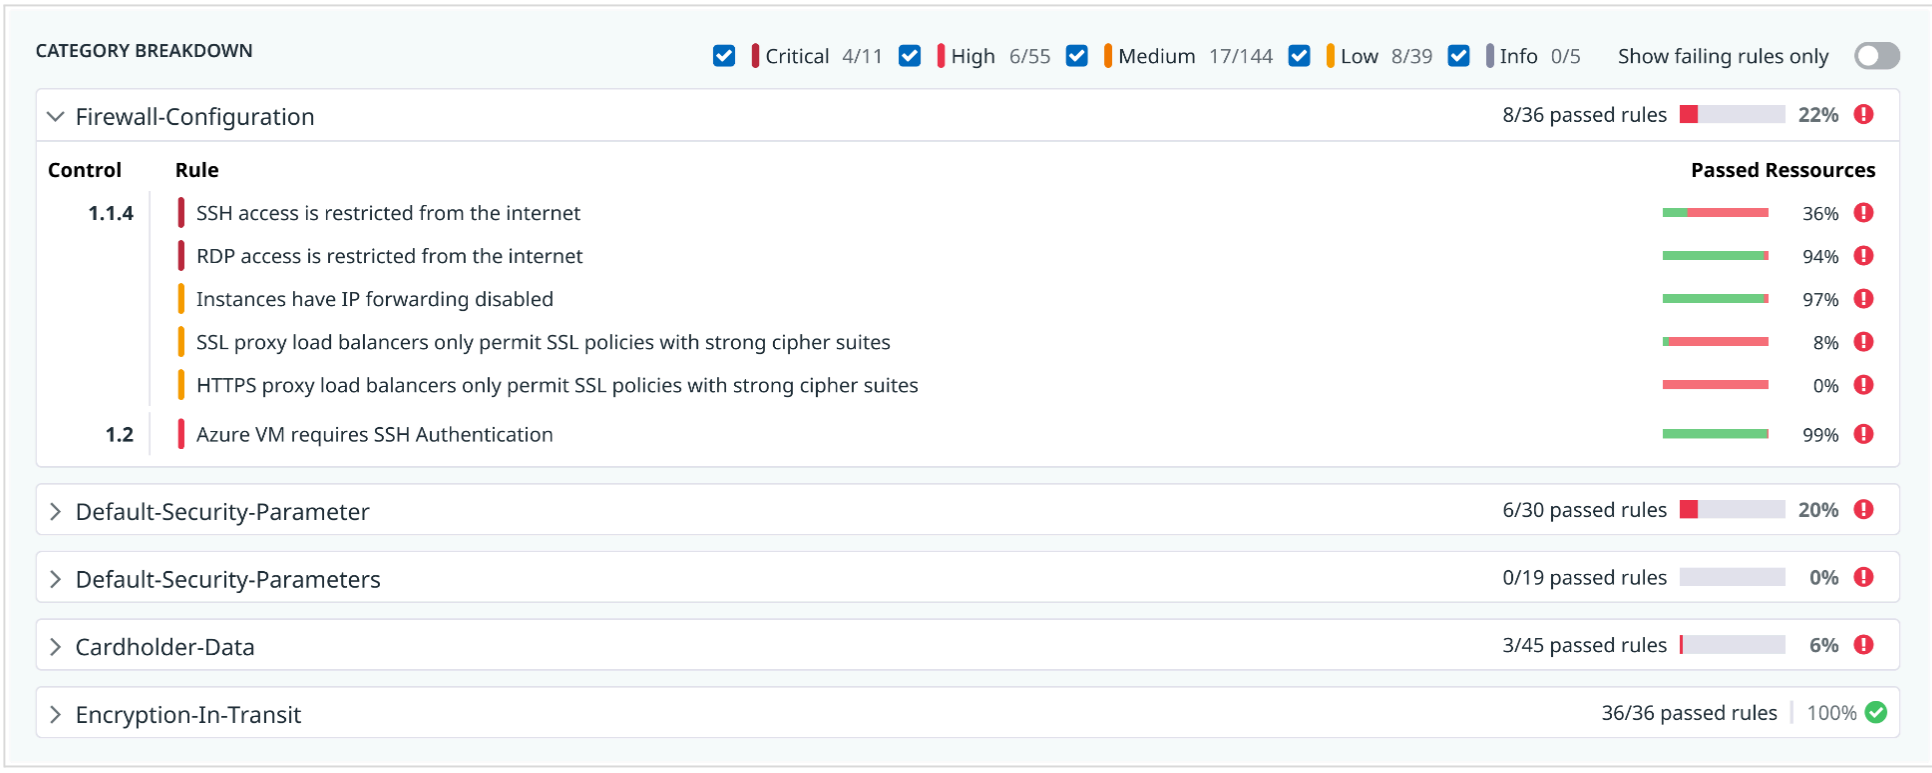




On another hand, we improved the compliance framework page's top overview under a triptych of compliance insights linking the framework posture score trend to the most pressing issues (top 5 of the highest rule failures and top 5 lowest scored accounts). Each of these compliance widgets offers further contextual information through hover interactions and tooltips. The compliance score trend graph is a valuable addition to track progress over time and fills the lack of historical data. It will help users identify patterns and measure the impact of their actions. This way, we give users the capability to define their next remediation efforts right on the compliance framework report page.



We standardize the posture score data presentation using percentages only, as applied on the compliance overview page. This new posture score pattern has been added to each rule category and each rule itself, users can filter all framework rules by severity level and highlight those that are failing with a click on a switch button.



Currently, demonstrating compliance performance to leadership and external auditors is a tedious process for our users who have to do it manually by exporting a PDF version of the compliance framework report page. We are pushing for letting our users generate ongoing evidence records through automated compliance reports, so they don’t need to create manual reports anymore. It will require moving from a simple button (“Export as PDF”) to a split button offering to generate an automated compliance report. We have already worked on a first report template and are discussing with engineering how we could soon generate automated reports. Later, we would like to confirm the need to make it as granular as the header filter options introduced above. We might also need to work on a disclaimer privacy policy notice to show when our users try to generate a report.

Page 1

DATADOG

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Cloud Conformity Infrastructure Report: NIST 800-171 (Rev 2)

AWS Account: Staging

Sun Apr 12 2020 15:33:31 GMT+0000 (Coordinated Universal Time)
NIST 800-53 (1 Account)

Summary

Datadog Score 99%

14 Requirements **110** Tested Configuration Controls

260k Evaluated Resources 1% **2.6k** Failed 99% **257.4k** Passed

Filters applied: **None**

Standards & Frameworks: National Institute of Standards and Technology

Compliance Score and percent success metrics are dependent on (1) active / selected filters, (2) data access you have provided on your account(s) or provided to you by your Organisation admin, and (3) controls which Datadog's Cloud Security Management is able to check for on your cloud infrastructure - this would exclude for example workload across accounts, and organisational processes. Compliance score, standard and framework control to rule mapping represents the expert opinion of Cloud Security Management and not necessarily that of the standard or framework authority. Rule severities and categories apply to Cloud Security Management's rules and not the controls they are mapped to. Your account(s) compliance with any Framework or Standard should be assessed in conjunction with your own internal review.

Artboard 72

Access Control

No:	Control			Priority	Total counts
AC-1	ACCESS CONTROL POLICY AND PROCEDURES			P1	0

Rule	Service	Categories	Risk level	Counts	
No rule to display					

No:	Control			Priority	Total counts
AC-2	ACCOUNT MANAGEMENT			P1	FAIL: 1 PASS: 5

Rule	Service	Categories	Risk level	Counts	
IAM Policy Changes Alarm	CloudWatchLogs	Security	High		
Root Account Usage Alarm	CloudWatchLogs	Security	High	PASS: 1	
IAM User Present	IAM	Security	Medium	FAIL: 1	Remediate
Account Alternate Contacts	IAM	Security	High	PASS: 1	
Root Account Usage	IAM	Security	High	PASS: 1	
Support Role	IAM	Security	High	PASS: 1	

No:	Control			Priority	Total counts
AC-3	ACCESS ENFORCEMENT			P1	FAIL: 141 PASS: 1

Rule	Service	Categories	Risk level	Counts	
Root Account Usage Alarm	CloudWatchLogs	Security	High	PASS: 1	
IAM User Present	IAM	Security	Medium	FAIL: 1	Remediate
Root Account Usage Updated	IAM	Security	High	PASS: 1	
Support Role	IAM	Security	High	PASS: 1	

III/ Next Steps

Learnings

This 2024 CSPM redesign proposal is based on a user-centric approach to bridge the gap between CSPM's technical complexity and user's expected clarity. By addressing CSPM's high churn rates with enhanced actionable insights, we propose a more intuitive and empowering experience for our users.

Our user interviews and testing confirmed that clarity, explainability and path-to-remediation matter more than anything else. Security engineers, SREs and CISOs are willing to adopt a new compliance experience as long as they can see, in one place, which frameworks and accounts are at risk, why the score changed, and how to move from insight to remediation without getting lost in separate explorers / tools / experiences.

On the impact side, the new compliance experience helps us track a clearer set of UX KPIs. We can now track time-to-first scoped view (from page load to landing on a meaningful, filtered compliance slice), posture score improvement for key frameworks over time, export and automated report creation rates, and the share of compliance sessions that lead to a remediation path (pre-filtered explorer, ticket creation, or configuration change). These indicators will help validate if we actually reduce churn and improve accountability, ownership, reporting, and CSPM's perceived value (engagement and trust).

As we move forward, we're already testing those prototypes with "internal customers" and selected customers through Ballpark and will leverage those first results to advocate for this new approach in our next meetings with Leadership. On another hand, we need to align further with PM and ENG for a P(0) and define how we could properly track the metrics described above in dashboards. We also need to confirm the customization limits of Karl Sluis' dataviz heatmap widget (APM) as a possible alternative to creating a new compliance heatmap component.

First Tradeoffs

From day one, the inception of Cloud Security Management and the redesign of CSPM meant navigating constantly into uncertainty and ambiguity. It isn't a sleek vision: we constantly have to make concrete tradeoffs with PM and ENG.

After presenting the initial early sketch of the new compliance heatmap, PM made it clear that we need to enhance the compliance posture's clarity without adding a new costly custom visualization component. ENG confirmed its feasibility but highlighted the need for a dedicated team and potential conflicts with other ongoing CSM features. Both pushed to reuse an existing DRUIDS component to ship faster and reduce maintenance. As a middle ground, we consider Karl Sluis' dataviz heatmap for P(0), if customizable. The new heatmap component is planned for P(1), if the redesign meets the UX KPIs listed above.

The top filter that is inherited from the CSM Overview page triggered another round of trade-offs. PM wanted persistent filters shared between the CSM Overview page and the CSPM overview page, while engineering warned about state management's complexity and deep links, since persisting filter settings per user across multiple pages meant additional storage, caching and potential edge cases. So we agreed on a few core facets like account, environment and service in the first iteration, as reducing friction on Compliance was our main goal.

The table list and framework cards as a single layout also required negotiation. PM doesn't want to introduce two views that could confuse users and slow down implementation. ENG wants to avoid too many components to support and test. But, our user research and interviews clearly uncovered two distinct mental models: some users wanted a minimal, classical view to quickly scan and export, while others preferred the cards' narrative that connected posture trends and top issues at a glance. We explained that the framework cards were designed from existing DRUIDS primitives (cards, badges, progress bars and inline charts) to be implemented as a composition rather than a new custom component. We push for the table list as default view for (P0) and the cards as an improvement for P(1).

Tracking/observability is usually a shared responsibility between PM and ENG, our role being to propose relevant UX KPIs. PM needs reliable UX metrics tied to the OKRs, such as time-to-first scoped view and export usage while ENG needs to avoid metrics that would add significant overhead to achieve. We reviewed together the few events that were essential for design and they agreed between them to reuse existing analytics hooks and how they would derive UX metrics from raw events. This way, we avoid metrics that are impractical to track by ENG, or metrics that don't make sense from a UX perspective.

We already know that some other subjects will be trimmed or postponed as DASH will quickly knock at our doors (like automated compliance reports in PDF implying technical challenges on the backend). But we're confident that the core of this redesign proposal could be shipped. We will keep you posted.

If you have any questions or would like to discuss further, please don't hesitate to reach out through Slack.

Thank you for reading and commenting on this ongoing design review.