2022

The State *of* Cyber Assets Report

SCAR

Jasmine Henry | Sounil Yu | Jennie Duong | Erkang Zheng | Mark Miller

Contents

JupiterOne

Contents

JupiterOne

The Challenge of Securing an Expanding Target

Analyzed

+372M

Cyber assets, findings, and policies

Analyzed

1,270

Organizations

Goal

Gain a better understanding of today's cyber asset landscape

Year after year, security research on data breaches shows that basic cyber hygiene is highly effective. Most security incidents with data loss fit several common attack patterns, and basic cyber hygiene offers protection against common threats such as ransomware and web application attacks.

The term 'basic cyber hygiene', however, implies that these activities are simple and leaves many business leaders wondering why security teams struggle with patching and enabling multi-factor authentication. What's missing from their assumption, however, is understanding the incredible volume and variety of assets that require cyber hygiene. "Doing the basics" is not so basic for security teams considering the sheer number and complexity of cyber assets in use today.

Not only is the attack surface growing, but the scale of the problem is now untenable. That's why we've set out to conduct and write this research.

To gain a better understanding of today's cyber asset landscape, we analyzed over 370 million cyber assets, findings, and policies across almost 1,300 organizations.

We believe the data uncovered reveals a lot about the state of cybersecurity in 2022 and why just "doing the basics" is so incredibly difficult for modern security organizations.

The technology shift towards cloud, software-defined, and everything-as-a-service have profoundly impacted security practitioners worldwide. The result is a massive growth in size of the enterprise attack surface and in volume of attacks against that exposed surface.

The evolving state of the modern cyber attack surface is the reason we created The (SCAR). It's the first of many annual reports to understand cyber assets, liabilities, attack surfaces, and their relationships to each other in the modern enterprise.

Jasmine Henry
Sounil Yu
Jennie Duong
Erkang Zheng
Mark Miller

JupiterOne

The Challenge of Securing an Expanding Target

Research Objectives

Executive Summary

Our research included an analysis of:

372.5M
Assets & attributes

1,270
Organizations

1.7M
Queries & alerts

210M
Cyber assets

151.5M
Alerts & findings

10.5M
Policies

JupiterOne

Executive Summary

Security Teams are Fatigued and Understaffed

Security teams have an unprecedented number of assets to secure and manage. The average security team is responsible for:

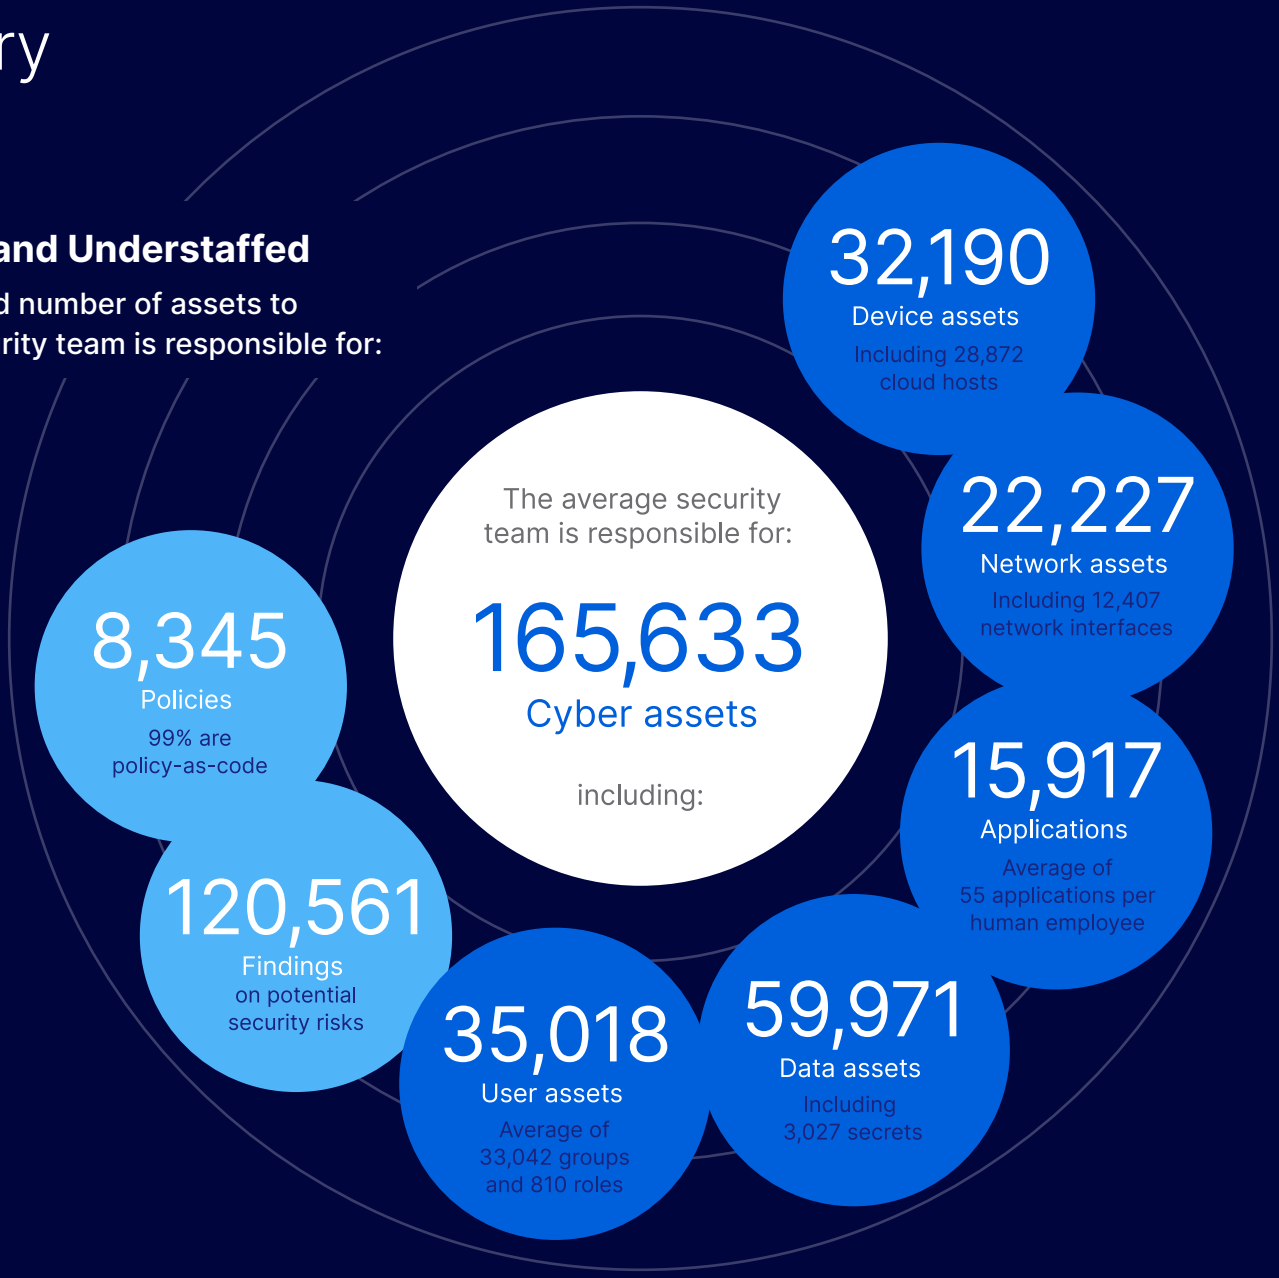

Cyber assets significantly outnumber employees in the enterprise. In fact, the ratio of cyber assets to human users is 564 to 1. Security teams are way outnumbered, especially considering that just 0.46% of the US workforce is security professionals.

JupiterOne

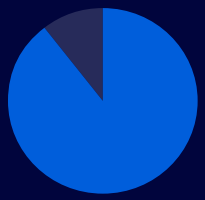

89.62%

Of devices in the modern organization are cloud-based

59:1

Cloud networks to physical networks

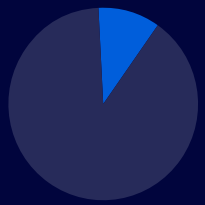

9.2%

Physical devices account for less than 9.2% of the total devices in the modern organization

110

Devices for every employee at the average organization

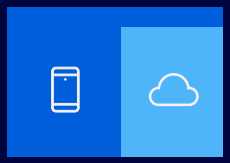

32,190

The average (mean) security team is responsible for 32,190 devices, including 28,872 cloud hosts

Legacy IT Security Skills Are Not Ready for Cloud-Native Realities

Actual human users make up only 18 of every 1000 assets (0.18%) and are a very small proportion of assets in the data set. The ratio of cyber assets to each human is 564 to 1.

Legacy IT skills don't reflect the realities of the cloud-native, serverless architectures of the modern enterprise. The industry is overdue for a serious look at how rising security pros are trained and certified.

Cloud assets generate over 97% of security findings, but cloud policies represent less than 30% of total guardrails. While sheer policy numbers are not a complete indicator of the state of cloud security, all security teams should consider whether their policy has scaled appropriately to mitigate cloud security risk.

JupiterOne

Distributed, Dynamic Architecture is Changing Security Forever

Today's enterprise infrastructure is engineered for extreme reliability (resiliency) via sophisticated, serverless architectures. This is great news for reliable software, since more pieces means there are less single-points-of failure. But, it also means that automation is a mandatory practice for security teams.

IP addresses comprise fewer than 1% of network assets, while network interfaces are the clear majority at 55.70% percent (267,029 IPs and 15,782,226 network interfaces, respectively). The average (mean) security team is responsible for 12,407 network interfaces.

> Modern DevOps teams use network interfaces to route traffic between subnets by hosting load balancers, proxy servers, and network address translation (NAT) servers.

Modern DevOps teams use network interfaces to route traffic between subnets by hosting load balancers, proxy servers, and network address translation (NAT) servers.

Third-Party Software Ate the World and Security Teams Pay the Price

Analysis of over 20 million application assets found that just 8.7% of applications are homegrown, or developed in-house. Only around 8.7% of code assets have change management trails to indicate it's developed in-house, like modules, functions, or pull requests (PRs).

91.3% of code running in the enterprise is developed by a third-party, meaning that modern organizations are incredibly vulnerable to supply chain attacks.

The average security team is responsible for 15,916.78 application assets, or an average of 54.54 assets per human employee. Nearly 16% (15.91%) of applications are services, or applications that run with minimal human touch, including web app firewalls, auto-scaling services, and event services.

Asset Relationships are a Blindspot for Many Security Teams

Analysis of nearly 400 million assets at almost 1300 organizations debunked the myth of "orphaned assets." Data, including critical data and sensitive personal records, is among the most-related types of assets, with 105 million first-degree relationships to users, apps, and devices. Analysis also uncovered nearly 45 million relationships between security findings, indicating that many security backlogs contain critical vulnerabilities or policy exceptions.

Debunked the myth of "orphaned assets"

3.8 million queries by security practitioners showed that practitioners are evaluating risks to devices, users, and apps far more often than networks or data. Just 8% of security queries consider second-degree or third-degree relationships between assets, indicating that security teams may be too under-resourced to fully understand the blast radius of potential compromises.

Section 1

The State of Cyber Assets: An Overview of Asset Classes, Attributes, & Trends

JupiterOne

Cyber Assets

Chart 1.1: Composition of Cyber Assets by Superclass

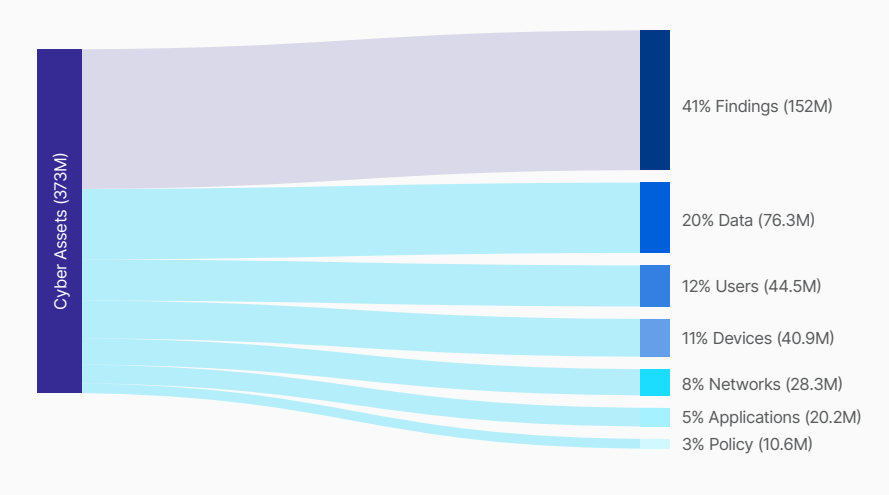

Chart 1.2: Frequency of Cyber Assets by Superclass

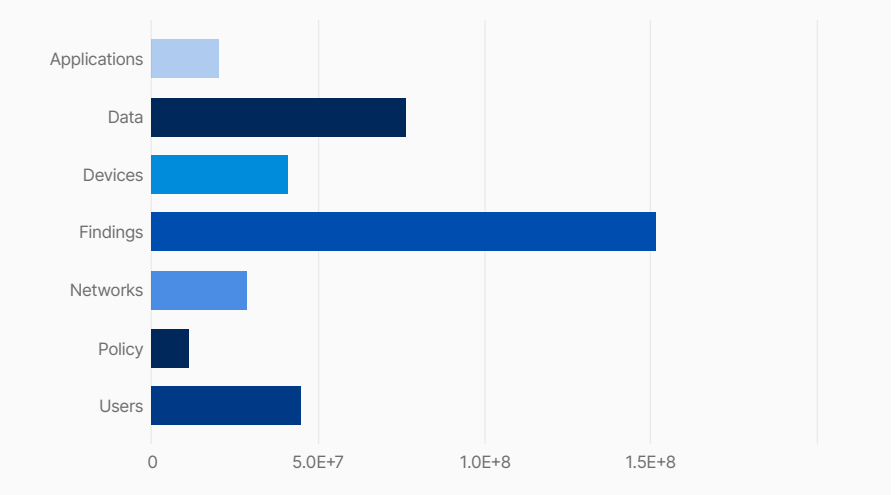

FREQUENCY

372,517,849 total entities at 1,270 organizations, or 210,354,473 cyber assets, 151,564,870 **FINDINGS**, and 10,598,506 **POLICIES**.

WHAT'S INTERESTING

Only 0.18% of cyber assets are humans, an approximate ratio of 1 human to every 564 non-human cyber assets.

ON AVERAGE

The average security team is responsible for 165,633 cyber assets, 119,342 **FINDINGS**, and 8,345 **POLICIES**.

Superclass	Number	Total
APPLICATIONS	20,246,148	5.43%
DATA	76,283,186	20.48%
DEVICES	40,945,907	10.99%
FINDINGS	151,564,870	40.69%
NETWORKS	28,336,413	7.61%
POLICY	10,598,506	2.85%
USERS	44,542,819	11.96%

JupiterOne

Cyber Assets

The immense number of cyber assets in the dataset paints a picture of an attack surface that is enormous. Nearly 7 in 10 organizations admit they have experienced at least one cyber attack that started through the exploit of an unknown, unmanaged, or poorly managed internet-facing asset.[3]

When organizations use automated asset inventory discovery tools, as opposed to manual asset inventory methods, they typically uncover a shocking number of cyber assets. Studies show that 31% discover sensitive data in an unknown location and 28% discover previously unknown SaaS applications. Other commonly-discovered cyber assets include misconfigured SSL certificates, weak encryption ciphers, code fragments, unknown third-party connections, and forgotten subdomains.[3]

Humans vs. Assets

Only 0.18% of cyber assets are actual human users.

This research study included 210,354,473 total cyber assets from 1,270 organizations with a collective total of 371,232 employees, or a mean of 292.31 employees per organization.

There Are 120,561 Cyber Assets For Every Security Practitioner

This presumes the 1,270 organizations included in this study have a total of 1,745 security practitioners. The latest research shows 0.46% of the US workforce is in security, or 4,600 security practitioners for every million U.S. workers.

We sincerely hope that all organizations represented in our data have a significantly above-average security team size. Nevertheless, the ratio of security practitioners to cyber assets is dire, even if your security team is 5% of total employee headcount.

Forget Manual Attack Surface Management

Traditional approaches to asset inventory like spreadsheets, IT asset management, or configuration management databases (CMDB) may still be a reasonable option in environments where humans have a hands-on approach to creating and provisioning all assets.

For the organizations represented in this study, however, asset creation is largely automated and asset inventory should be as well. Spreadsheets don't scale to the cloud.

Nearly

7 in 10

Organizations admit they have experienced at least one cyber attack

What it Means for Security

Today, an entire asset lifecycle can occur without human intervention or knowledge.

A world where assets are automatically created, managed, and destroyed demands that security teams automate everything, especially asset discovery and protection. After all, threat actors use automated tools to research the attack surfaces of their potential targets.

Section 2

Cyber Asset Terminology & Concepts

JupiterOne

Classification Patterns

Assets and attributes are grouped into superclasses for the purpose of visualization and analysis. These asset superclasses are largely drawn from Sounil Yu's Cyber Defense Matrix, with the addition of two superclasses to describe asset attributes: `FINDINGS` and `POLICIES`.

The definitions and inclusions for each category are included below, or described further in JupiterOne's data classification model on Github. This report will utilize all caps to easily identify the superclasses by name.

Asset Superclasses

Devices

The **`DEVICES`** superclass consists of workstations, servers, phones, tablets, containers, hosts, peripherals, storage devices, network devices*, web cameras, infrastructure, and more. It also includes operating systems, firmware, and any other software native to a device.

*Networking devices (like switches and routers) are included here because those devices are separate from the network communication pathways they create.

Networks

`NETWORKS` are communications channels, connections, and protocols that enable traffic to flow among **`DEVICES`** and **`APPLICATIONS`**. This superclass also includes DNS, BGP, VPCs, VPNs, CDNs, and certificates.

Data

`DATA` includes data-at-rest, data-in-motion, and data-in-use. This superclass includes databases, S3 buckets, storage blobs, and files. Also, the **`DATA`** superclass includes queues, logs, change records, tasks, and channels. Last and perhaps controversially, secrets are grouped with data, including encryption keys, key pairs, and vaults.

Users

`USERS` are people and machine identities who use the resources in other asset superclasses, and the identities associated with these users. The **`USERS`** superclass also includes groupings of users, including teams, organizations, and sites.

Attribute Superclasses

Findings

The **`FINDINGS`** category consists of alerts and results, incidents data, monitoring trails, threat intel, and vulnerabilities from both human and non-human sources.

Policy

Much like **`FINDINGS`**, **`POLICY`** falls outside the classification of traditional assets and should be considered an attribute of cyber assets in the sense they act as guardrails to protect assets.

IAM policies, control policies, configurations, requirements, and rulesets all fall within our classification model for **`POLICY`**. Human-generated policy and procedure documents are here, too, though they're a negligible percentage compared to other forms of **`POLICY`**.

The Graph Data Model

The underlying data model for this analysis is based on graph theory; specifically a reference model used to describe cyber assets and their complex interactions in a modern organization. The data model is defined by a set of entities and their relationships, or nodes and edges:

- An entity is a node or vertex in the graph that represents a cyber asset
- A relationship is the edge between two entity nodes in the graph

Additional information on the underlying graph data model can be found in this Github repo.

JupiterOne

Additional Terminology

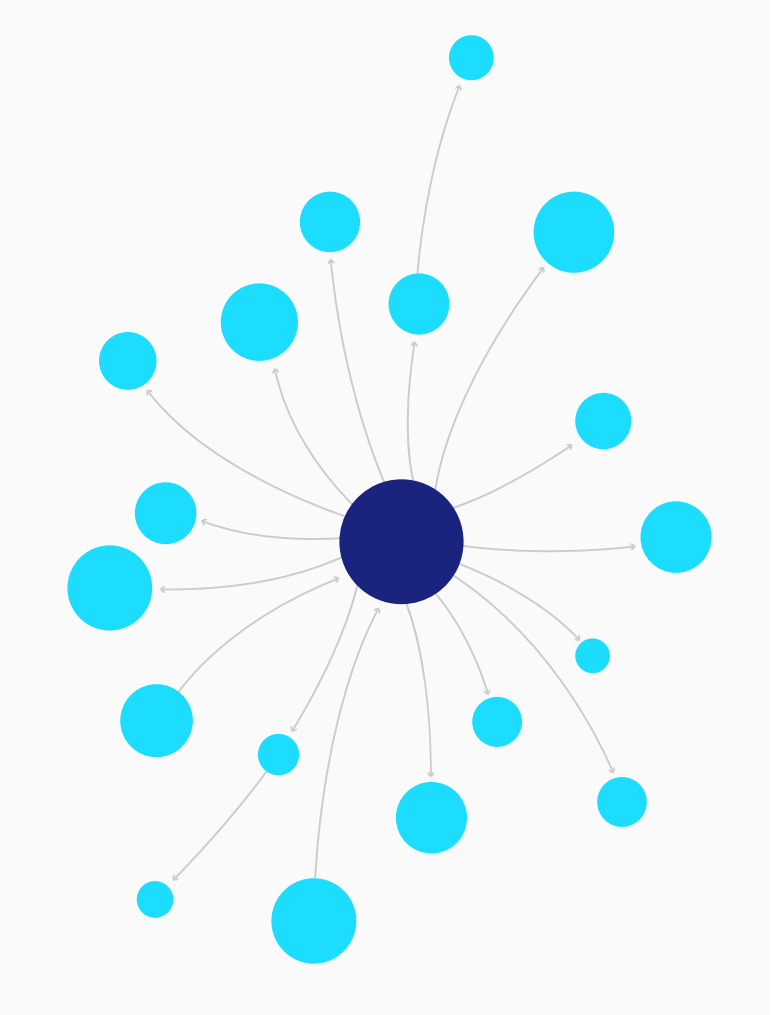

Cyber Asset Relationship

A relationship is the connection between two or more cyber assets. Assets in isolation don't tell us the entire picture, it's how they connect and work together to provide value or risk to a business and processes.

Context or 'Cyber Asset Context'

Context is metadata, data, or information that provides added perspective or attributes of any cyber asset(s). Context provides a broader understanding into a cyber asset, it's relationships, and how they relate to one another in the broader system or environment.

Critical Assets

Critical assets, as defined by CISA, are cyber assets that are essential to maintaining operations and achieving the organization's mission. If the confidentiality, integrity, or availability of a critical asset is breached, there are generally significant business consequences.

While critical assets often have common characteristics, the designation of a cyber asset can vary significantly between organizations. Criticality is largely dependent on the business environment, processes, risk appetite, and policy.

Security Practitioner

A security practitioner is an employee whose primary responsibility is the security of cyber assets. The term is inclusive of all employees who are part of the cybersecurity team, while also being agnostic of the variations in security job titles and ranks. Security practitioners are most often in roles such as CISOs, VPs, Directors, Security Engineers, Analysts, Testers, and Auditors.

Section 3

Cyber Assets by Superclass

Devices Superclass

Chart 3.1: Composition of DEVICES Superclass

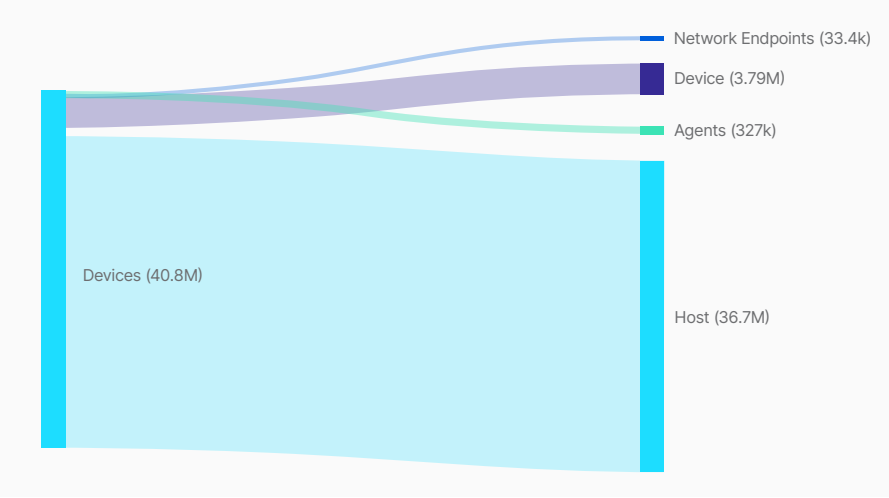

Chart 3.2: Composition within DEVICES Superclass

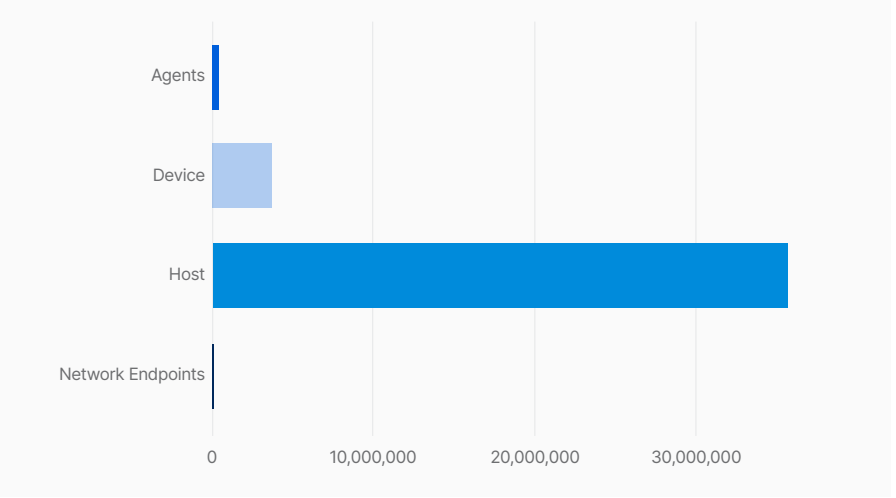

FREQUENCY

40,945,907 **DEVICES**

WHAT'S INTERESTING

Nearly 90% of enterprise **DEVICES** are cloud-based.

ON AVERAGE

The average security team is responsible for 32,190 **DEVICES** – including an average of 28,872 cloud hosts per organization.

Superclass	Number	Total
Agents	327,296	0.80%
Device	3,788,378	9.25%
Host	36,695,493	89.62%
Network Endpoints	33,378	0.08%

JupiterOne

Devices Superclass

Tangible Devices Are Less Than 10% of DEVICE Attack Surface

5G proponents have been talking about the Internet of Things takeover for years, but physical devices aren't anywhere near the greatest concern to security teams.

We certainly don't mean to diminish the security struggles caused by remote work and bring-your-own-device programs (BYOD), since perimeter-based security really was a much simpler era. According to our data, tangible devices only account for 9% of the **DEVICES** attack surface: there are bigger, more cloud-based assets to secure.

Enterprises Have 110 DEVICES Per Human

Out of 40,945,907 **DEVICES**, an incredible 89% are cloud hosts. That's nearly 37 million cloud hosts or an average of 28,872 software-defined endpoints that security teams need to secure.

DEVICES significantly outnumber the humans in an enterprise. The average enterprise has around 110 **DEVICES** per human when the entire asset superclass is taken into consideration. When you only consider the ratio of physical devices to employees, the enterprise has 10.2 physical devices per person.

Physical Network DEVICES Are a Tiny Percent of Total

Network endpoints bear inclusion, but there's an incredibly small number of physical network devices in proportion to everything else at 0.8% of the **DEVICES** superclass. The average modern enterprise has around 26 physical network components. This highlights the massive scale disparity between cloud and premises-based networks, and amplifies the industry's desperate need for cloud security talent.

What it Means for Security

The enterprise now has significantly more **DEVICES** than humans, but the majority of these **DEVICES** are not physical smartphones or laptops. The dominance of cloud hosts in the device mix demands new approaches to security training and upskilling that prepare rising talent for careers that are mostly cloud-focused.

The relatively tiny proportion of physical devices probably doesn't mean we're overspending on endpoint protection for tangible devices like laptops and smartphones. Physical endpoints are disproportionately risky since they're used by unpredictable humans. Chances are, we're probably way underspending and under-resourcing cloud security.

Networks Assets Superclass

Chart 3.3: Composition of **NETWORKS** *Superclass*

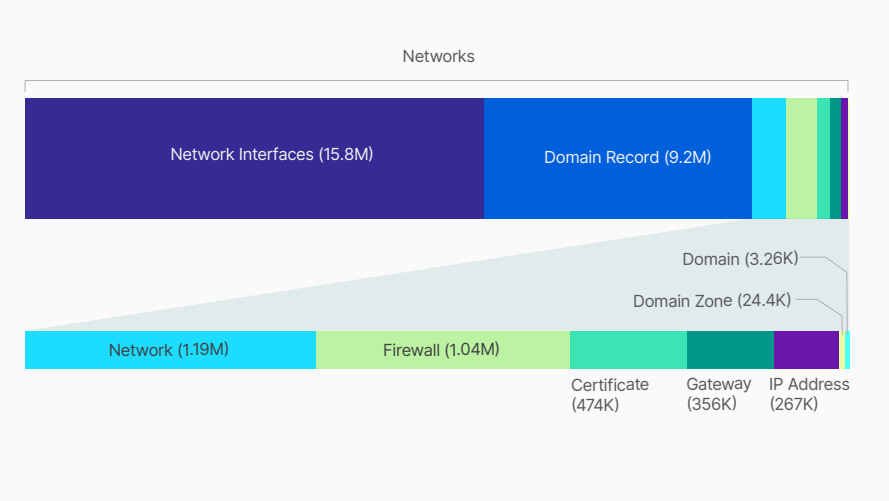

Chart 3.4: Composition within **NETWORKS** *Superclass*

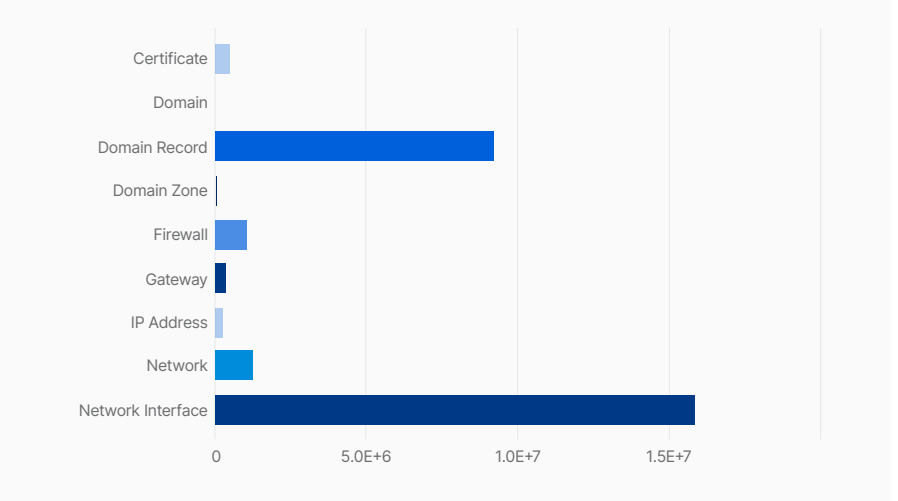

FREQUENCY

28,336,413 **NETWORKS** assets

WHAT'S INTERESTING

IP addresses comprise fewer than 1% of **NETWORK** assets, while network interfaces are the clear majority at 55% percent (that's a respective total of 267,029 IPs and 15,782,226 network interfaces).

ON AVERAGE

The average security team is responsible for 22,277 **NETWORK** assets, including 12,407 network interfaces.

Superclass	Number	Total
Certificate	474,311	1.67%
Domain	3,262	0.01%
Domain Record	9,204,506	32.48%
Domain Zone	24,429	0.09%
Firewall	1,038,008	3.66%
Gateway	356,247	1.26%
IP Address	267,029	0.94%
Network	1,186,395	4.19%
Network Interface	15,782,226	55.70%

JupiterOne

Networks Assets Superclass

Network Interfaces Significantly Outpace IP Addresses

267,029 IP addresses is not an insignificant number, but what is really significant is the ratio of IPs to network interfaces. There are 59 network interfaces for every IP address, which reveals the state of modern network architecture and how network security has changed. The vast majority of pathways traveled are dynamic, elastic, and involve no static network components. Traditional network security concepts for static, perimetered networks have limited value for modern practitioners.

The rate at which premises-based networks have become proportionally insignificant is the first of several findings which calls into question the industry's approaches to cybersecurity skill, talent, and education. It may be time to start considering the role of cloud network security in cybersecurity certifications and skills initiatives. Skills programs pay a disproportionate amount of attention to premises-based networks compared to the actual cyber asset mix discovered in our research.

Over Half of `NETWORK` Assets Are Network Interfaces

The dramatic shift towards resilient network architecture is abundantly clear, especially given the proportion of network interfaces to the network assets superclass. DevOps teams use network interfaces to route traffic between subnets by hosting load balancers, proxy servers, and network address translation (NAT) servers.

The fact that over 55% of **`NETWORK`** assets are designed for elasticity is a triumph for availability, and perhaps even a natural reaction to a year in which at least one major cloud infrastructure provider had extended downtime incidents. On a less delightful note, the idea of assigning policy to over 12,407 network interfaces is a lot to wrap your head around, especially compared to simpler times when security practitioners were responsible for a small handful of networks.

Certificates Demand Modern Lifecycle Management Approaches

While the relative proportion of certificates is fairly low, it's still noteworthy considering that 474,311 certificates at 1,272 organizations is a mean of 372 certificates for each DevSecOps team to manage. Let's all agree that legacy approaches to certificate management, like spreadsheet tracking for certificate expiry, are officially over.

What it Means for Security

So, does the mean of 12,407 network interfaces per organization indicate that the glass is half-full or half-empty? The answer depends entirely on whether the reader comes from a site reliability engineering (SRE) or a cybersecurity background.

The incredible number of network interfaces is different, depending on where you sit in the "confidentiality-integrity-availability" triad. Proliferating network interfaces are great news for availability, but a new challenge for those of us whose primary interest is confidentiality.

JupiterOne

Applications Superclass

Chart 3.5: Composition of the **APPLICATIONS** *Superclass*

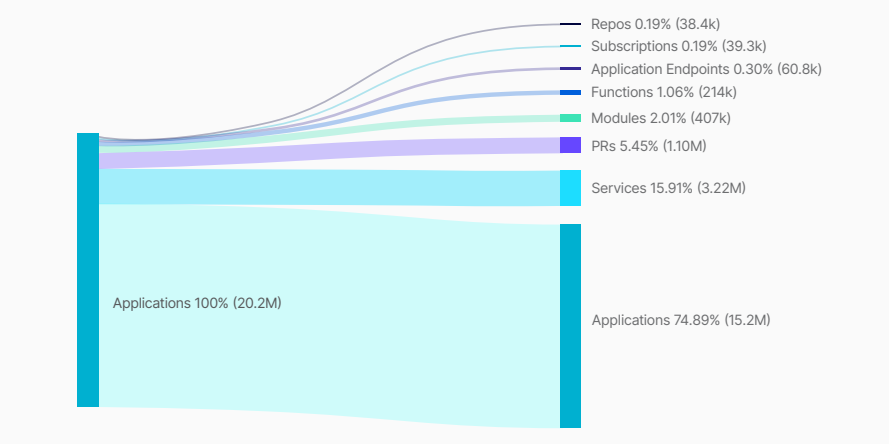

Chart 3.6: Composition within **APPLICATIONS** *Superclass*

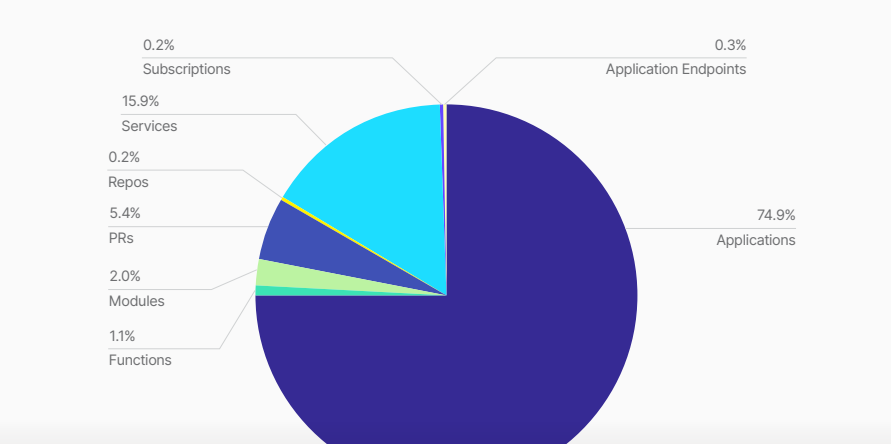

FREQUENCY

20,246,148 **APPLICATION** cyber assets

WHAT'S INTERESTING

Only 8.7% of **APPLICATIONS** assets have the change management trails associated with homegrown code, such as modules, functions, pull requests (PRs), or code repositories. The remaining 91% of **APPLICATIONS** running in an enterprise are built on third-party code.

ON AVERAGE

The average security team is responsible for 15,916 **APPLICATIONS** assets, or an average of 54 **APPLICATIONS** per human employee.

Superclass	Number	Total
Application Endpoints (e.g. APIs, dedicated web app URLs)	60,828	0.30%
Applications	15,162,398	74.89%
Functions	213,813	1.06%
Modules	406,969	2.01%
Pull Requests (PRs)	1,103,047	5.45%
Repos	38,384	0.19%
Services	3,221,446	15.91%
Subscriptions	39,263	0.19%

JupiterOne

Applications Superclass

There's a Lot of Third-Party Code

In today's world, every organization competes on software, but don't infer that homegrown applications power enterprises. In fact, the asset subclasses associated with internal software development lifecycle (SDLC) collectively comprise only 8% of the total. The subclass breakdown is functions (1%), Modules (2%), pull requests (5.4%) and repos (0.2%).

The security events of 2021 made all of us acutely aware of third-party code security risks and their downstream impact. The composition of the **APPLICATIONS** superclass only serves to underline the existing panic we all feel about our exposure to code we didn't write ourselves.

Did We Mention There's a Lot Of Code?

When humans are viewed in proportion to everything else, the incredible scale of the enterprise attack surface becomes painfully clear. The average enterprise has 54 **APPLICATIONS** per human employee. We all use a staggering number of apps to do our jobs!

It's little wonder that applications are among the fastest growing attack vectors for threat actors, considering their prominent role in the enterprise asset mix.[4]

APPLICATIONS Are Often Autonomous

Notably, 16% of the **APPLICATIONS** superclass is "services," which are applications that run a task with minimal human involvement, including web application firewalls, auto-scaling services, and event services.

The average security organization must create effective guardrails for an average of 2,532 services. To manage this volume of services requires modern cybersecurity efficiency at a scale previously unimaginable.

What it Means for Security

Software supply chain security is becoming untenable. The sheer number of **APPLICATIONS** and the relatively narrow proportion shows why the software supply chain is so difficult to manage. The recent Log4Shell vulnerability is a crucial example of how challenging it can be for many organizations to trace popular libraries across **APPLICATIONS**, **DEVICES**, and legacy systems—especially when it comes to understanding business dependencies and relationships between vulnerable assets.

Visibility into assets and the software bill of materials (SBOMs) are obvious solutions to the challenges faced by security teams, but security functions also must minimize unneeded complexity. Vendor consolidation and end-of-life procedures for legacy systems become a necessity in resource-constrained environments.

Data Superclass

Chart 3.7: Composition of the **DATA** *Superclass*

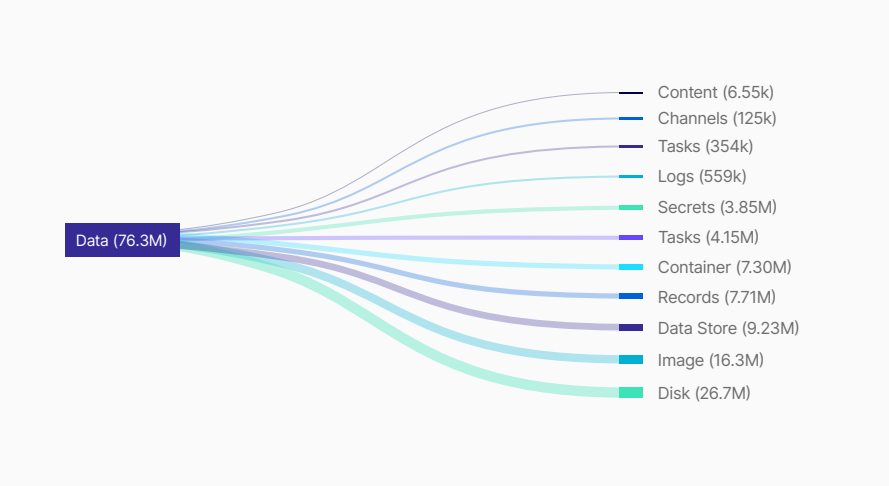

Chart 3.8: Composition within **DATA** *Superclass*

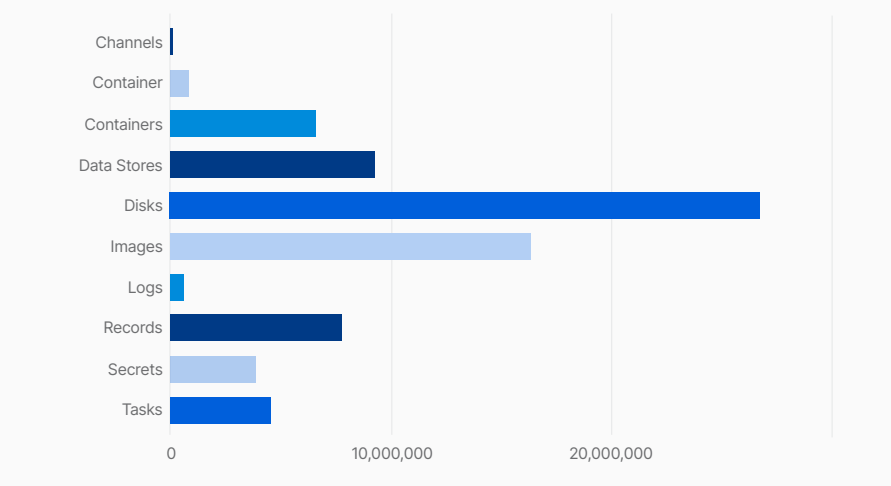

FREQUENCY

76,283,186 **DATA** assets

WHAT'S INTERESTING

Tasks now account for nearly 6% of **DATA** assets.

ON AVERAGE

The average security team is responsible for 59,971 **DATA** assets, including 3,027 secrets.

Superclass	Number	Total
Channels	124,783	0.16%
Containers	7,304,830	9.58%
Data Stores	9,233,081	12.10%
Disks	26,675,011	34.97%
Images	16,314,814	21.39%
Logs	558,603	0.73%
Records	7,717,804	10.12%
Secrets	3,849,916	5.05%
Tasks	4,504,344	5.90%

Data Superclass

Big Data, Bigger Problems

If you are thinking "that's a staggering number of cloud data stores," you are correct. Disks and images collectively comprise 56% of enterprise data, all of which must be stored securely on cloud devices. Disks and images significantly outpace the number of logs and records.

Secrets Are No Fun

The tasks subclass comprises secrets, vaults, and keys. The relative significance of this category is worth noting, especially considering there are 3,027 secrets for the average security team to manage. This challenge is compounded by the extra time and resources needed to ensure all of these secrets have appropriately granular access, transfer protocol, a defined lifecycle, and loads of encryption.

Queues

The "tasks" subclass includes both tasks and queues, and is well worth interpretation. While it's far from the most stressful finding in this report, it's a clear sign showing the ways in which the enterprise works have changed significantly. Queues are a strong indicator of the decoupling that results from a microservices architecture, which can be anything from a request or reply, to an error message or metadata. It shows that today's enterprise architectures are vast, decoupled, and built for resilience.

This finding has interesting security implications when viewed in the context of everything else. The stakes for security teams continue to rise as their software-driven organizations adopt more and more cloud-driven components. The enterprise has embraced the idea of availability, which means there are more components to inventory and secure.

What it Means for Security

Assets could be a major expense for organizations that store backups long past the bare minimum required for data retention obligations (contractual, legal, and regulatory).

To be clear, we are not in the business of providing legal advice or compliance consulting around how long you should retain your data. Data retention schedules are completely between you and your compliance department since obligations vary between organizations.

However, we can advise you to consider the fact that big, unnecessary backups lead to more assets to store and manage, especially if you're exceeding minimums or not using long-term storage.

Users Superclass

*Chart 3.9: Composition of the **USERS** Superclass*

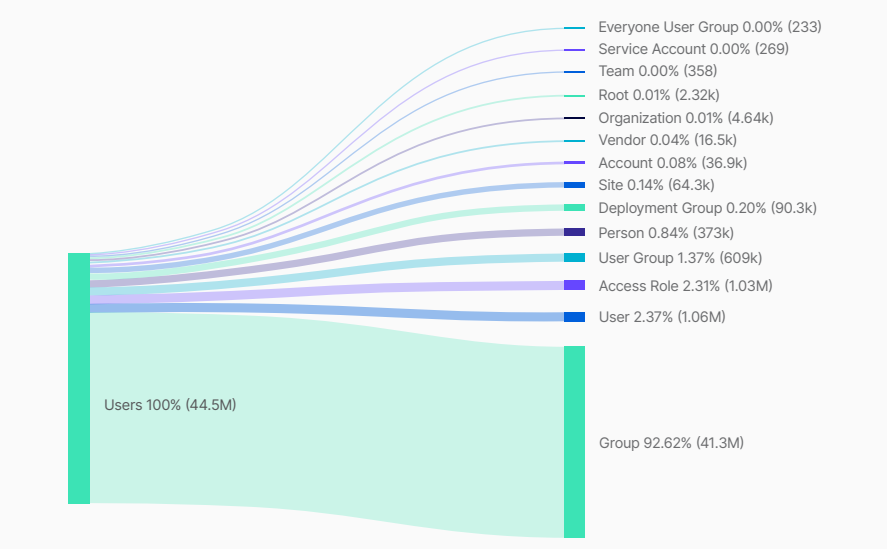

*Chart 3.10: Composition within the **USERS** Superclass*

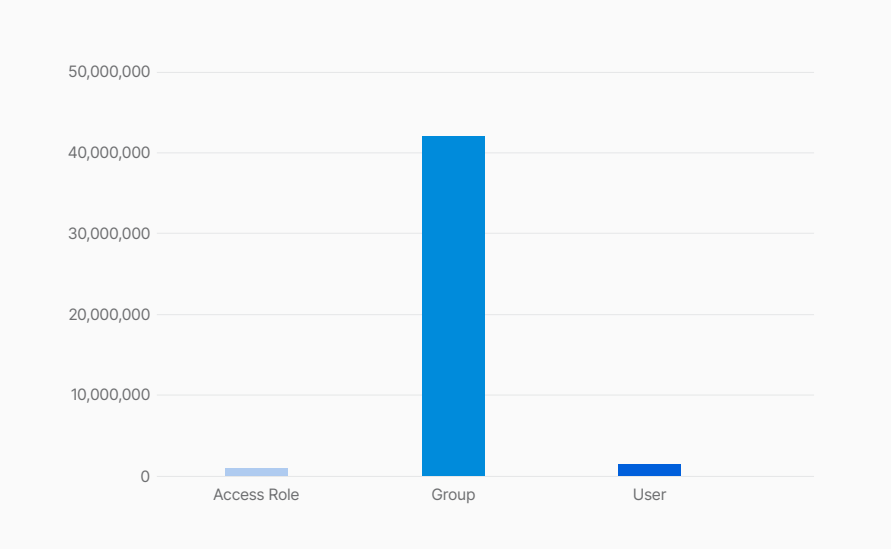

FREQUENCY

44,542,819 users, groups, and roles

WHAT'S INTERESTING

USERS definitely aren't the same as employees. On average, there are a total of 120 user identities per employee.

User identities can include each user account assigned to an employee, contractor, or authorized third party to access the organization's resources.

ON AVERAGE

The average security team is responsible for 35,017 user assets, including an average of 33,041 groups and 809 roles.

Superclass	Number	Total
Access Role	1,030,033	2.31%
Group	42,028,908	94.36%
User	1,483,878	3.33%

JupiterOne

Users Superclass

Security Teams Have Surprisingly Good Control Over Human Users

The total number of accounts belonging to humans included in analysis was not wildly different from employees. On average, there were 1.0051 human accounts for every employee, or fewer than 2,000 accounts belonging to former employees or other unauthorized users. If you're keeping track, that's an average of only 1.51 inactive employee accounts total per organization. Well done, security teams!

Further, the user subclass of this superclass includes a total of 16,467 users who are marked as "vendors," or accounts that belong to a third-party such as consultants or contractors. This means that only 4.2% of human users work for a third-party organization—a significant percentage but nothing terribly frightening.

Groups Abound, But IAM Security Is on The Right Track

A high number of groups is typically evidence of loads of applied identity and access management (IAM policy) in the wild. More groups indicate more efforts by enterprise security teams to slice and dice their users according to the principle of least privilege and business need. Based on this logic, we can conclude that security teams are trying to scale and apply least privilege, especially considering there are 28 groups for each user.

On The Flip Side, There Are Fewer Defined Access Roles Than Users

Precisely speaking, there's 0.69 roles for each user. One would expect an enterprise to have fewer roles than users, even given the proliferation of systems, so this indicates that identity and access security is headed in the right direction.

Users Are a Bit of a Paradox

Users are an infinitesimal component of the assets that security teams must inventory and manage. All identities are less than 12% of total assets, and actual employee identities are only 0.10% of the nearly 400 million assets included in this analysis. Our research suggests humans will continue to be a diminishing proportion of enterprise assets, with an outsized impact on security outcomes.

What it Means for Security

Humans are unpredictable and prone to error, so they will continue to be a stressful variable in the CISO's world, even as their relative proportion slips from 0.10% to 0.001% of total assets.

It will be hard for security teams to answer seemingly basic questions about who can access a cyber asset since the idea of access is incredibly ephemeral. Identity sprawl is only one piece of the puzzle. The cloud has layered on more complexity to manage, such as assumed roles and run-as identity configurations. Trying to reconcile identity and access with concepts like the blast radius of an asset is important, but also nearly impossible to do manually.

Section 4

Cyber Asset Attributes by Class

Findings Superclass

Chart 4.1: Composition within **FINDINGS** *Superclass*

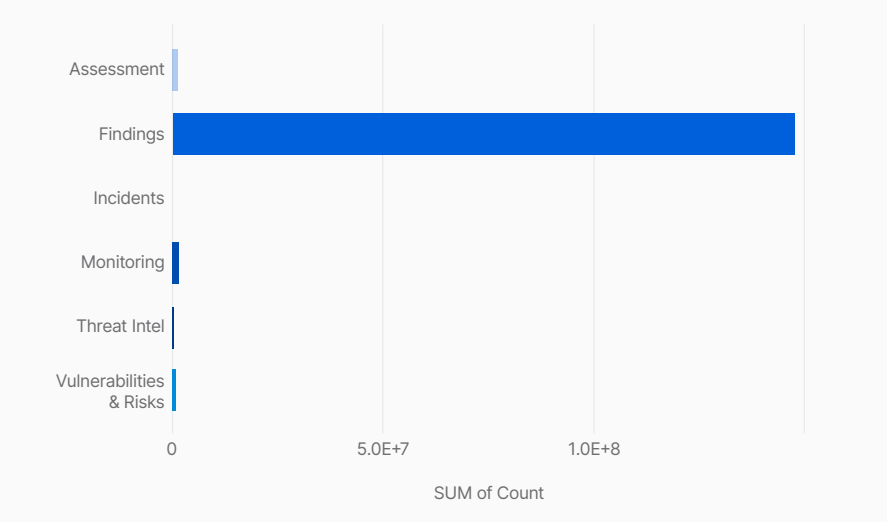

Chart 4.2: Composition of "Finding" Subclass within **FINDINGS** *Superclass*

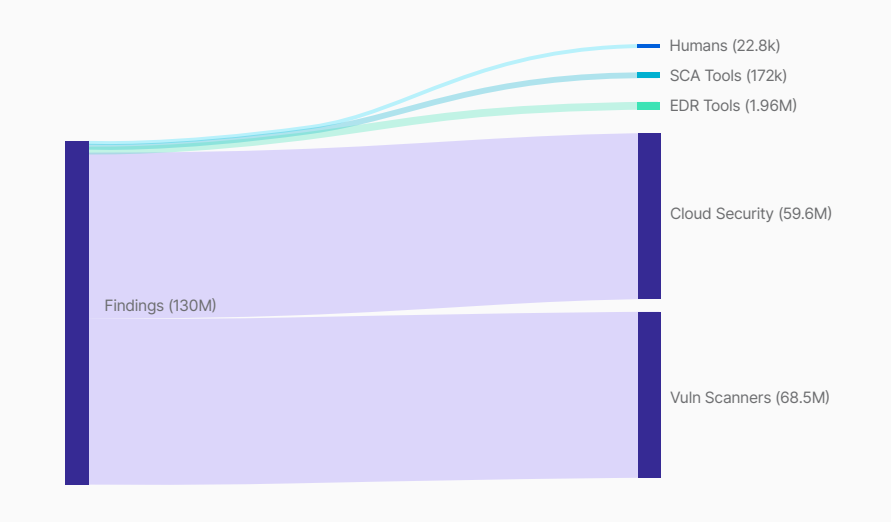

FREQUENCY

151,564,870 **FINDINGS**

WHAT'S INTERESTING

Only 2.7% of security findings have classification as a risk, threat intel, or assessment. The remaining 97.3% of security findings are like a firehose of noise.

ON AVERAGE

The average security team is responsible for a backlog of 120,561 security findings.

Superclass	Number	Total
Assessment	1,250,804	0.83%
Findings	147,432,773	97.27%
Incidents	113,442	0.07%
Monitoring	167,0848	1.10%
Threat Intel	144,568	0.10%
Vulnerabilities & Risks	952,435	0.63%

JupiterOne

Findings Superclass

151,564,870 Findings Is Not a Small Number

Humans are not responsible for the majority of `FINDINGS`, especially when you drill into the 97% of assets in this superclass that are simply labeled "findings." Deeper analysis within this subclass confirms that cloud security tools and vulnerability scanners are creating an unprecedented volume of alert noise.

As many practitioners know, a high volume of `FINDINGS` can make it difficult for any security team to effectively identify, triage, and remediate the most critical risks to your organization. The ability to surface and understand critical findings is subjective and often based on a company's policies or risk appetite.

What's Lurking in Your Backlog of `FINDINGS`?

Cloud security posture management tools are really, really loud. In fact, we dare to say that the enterprise security team is probably not whiny enough about how they're dramatically under-resourced and drowning in alerts.

Is there any other enterprise team that's tasked with a mean backlog of 120,561 issues? Security teams are not exaggerating the fact that they face deafening amounts of alert noise.

The Ratio of `FINDINGS` to Assets is Staggering

`FINDINGS` aren't cyber assets, they're attributes of cyber assets. Still, the fact this superclass represents nearly 41% of the total data in our analysis is staggering. That's a ratio of over 7 `FINDINGS` for every 10 cyber assets.

What's in a finding?

The "finding" subclass comprises 97% of this superclass, or a total of 147,432,773 findings (that's a total of 39.5% of all data included in this report). The majority of these findings are traceable to cloud security alerts, monitoring, and solutions.

We Probably Aren't Worried Enough About Alert Fatigue

Alert fatigue is a phenomenon of busy workers who become desensitized to safely alerts and ignore important warnings. Desensitization is a normal human response to constant alerts, especially in environments where the majority of alarms are false positives.

Alert fatigue has been studied extensively within healthcare as an unintended consequence of digital transformation in patient care settings. Researchers acknowledge that alert fatigue has sufficient impact on clinician behavior to create significant risk to patient safety.

Cybersecurity practitioners are generally familiar with the concept of alert fatigue. Security leaders should advocate for visibility and resources to manage findings; it's the most significant part of the enterprise cyber asset inventory. As an industry, we should collectively advocate for the resources necessary to tune our alerts, validate findings, and triage response efforts.

Security Teams Have Not Been Crying Wolf

The number of `FINDINGS` in this report is staggering and vindicates what security teams have been reporting for years: they are drowning in alerts. One recent study indicates that 70% of teams have seen alert volume double since 2015. According to the same report, 99% of security leaders admit alert volume is causing issues for their team.[6]

Automation isn't the entire answer. There's no industry standard for what constitutes a "mature" SIEM implementation, or an appropriate, safe use for automation in a cloud-native security operations center. The answer also isn't "more AI." Security teams already have too many black box algorithms that leave an unmanageable trail of findings and false positives. There's a long road ahead for the industry and creating visibility into the state of cyber findings will be necessary for us to find a way out.

JupiterOne

Findings Superclass

Cloud Security
59,555,755
45.72% of total findings
"Cloud security" includes alerts over AWS, Azure, or Google Cloud environments, as well as findings from cloud posture management (CSPM) tools.

EDR Tools
1,964,090
1.51% of total findings
Endpoint detection and response tools for both physical endpoints and cloud endpoints, like Carbon Black, Crowdstrike, and SentinelOne.

Human Findings
22,821
0.02% of total findings
This includes a small amount from security change management systems like Jira, plus all of the findings from pentests and bug bounty hunters.

Software Composition Analysis (SCA) Tools
172,499
0.13% of total findings
Findings from application security testing programs, specifically software composition analysis tools such as Snyk, Gitlab, and Black Duck.

Vulnerability Scanners
68,543,690
52.62% of total findings
Includes vulnerability scanner findings from Qualys, Nessus, and similar scanners.

Everything Else
171,739,180
11.65% of total findings
11.65% of total findings

What it Means for Security

The solution to the state of alerts is likely a mixture of recognition, human talent, and technology. CISOs and other security practitioners should educate executives about the sheer volume of findings and advocate for additional resources when their `FINDINGS` queues become unmanageable.

At many organizations, additional hires are likely a necessity. Security teams must consider tools that allow them to learn from false and true positives, and adjust their alerts to avoid fatigue.

Table 4.2: Composition of "Finding" Subclass within FINDINGS Superclass

Policy Superclass

Chart 4.3: Composition of **POLICY** *Superclass*

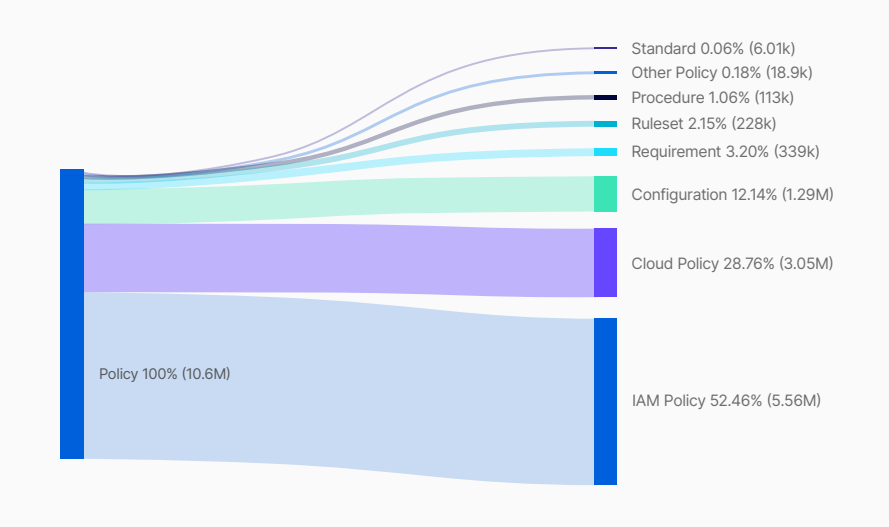

Chart 4.4: Composition within **POLICY** *Superclass*

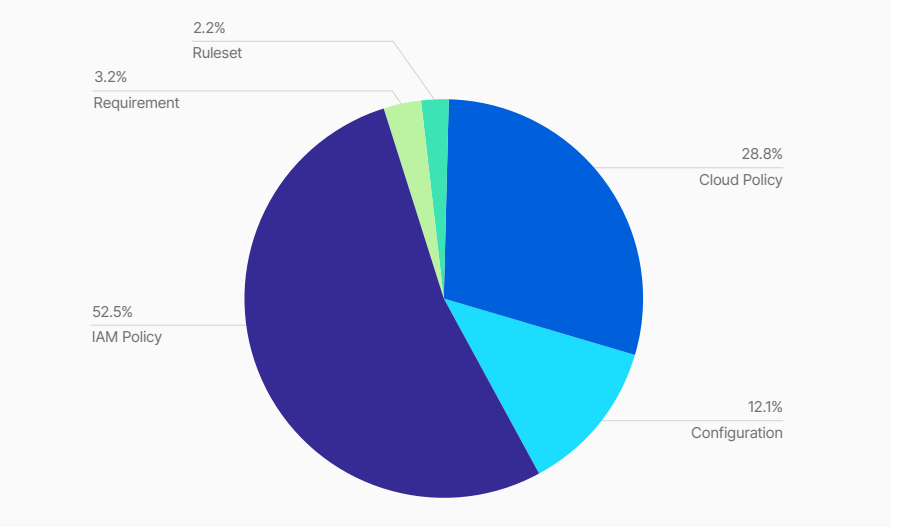

FREQUENCY

10,598,506 **POLICIES**

WHAT'S INTERESTING

Only 0.24% of **POLICY** is authored by humans.

ON AVERAGE

The average security team is responsible for 8,345 policies, but luckily, 99% of these policies aren't entirely written or enforced by humans.

Over 99% of **POLICY** is configurations, rulesets, procedures, and other policy-based enforcement.

Superclass	Number	Total
Cloud Policy	3,047,897	28.76%
Configuration	1,286,404	12.14%
IAM Policy	5,560,038	52.46%
Other Policy	18,895	0.18%
Procedure	112,533	1.06%
Requirement	338,624	3.20%
Ruleset	228,106	2.15%
Standard	6,009	0.06%

JupiterOne

Policy Superclass

The Trouble with End-User Credentials

Automation isn't optional for security teams, especially considering the incredible scale of the cyber asset landscape. The gradual death of security perimeters has many implications, including a shift to an attack surface which has fewer firewalls and more dependencies on users, configurations, and procedures for identity and access management (IAM).

Most security practitioners agree that enterprise security is far too reliant on undependable factors like human behavior and end-user credentials. Luckily, the `POLICY` superclass reveals that enterprises recognize identity and access as a significant pathway for compromise. Over 5.5 million IAM policies in total is a mean of 4,377.98 IAM policies per organization.

Cloud Policy and the Shared Responsibility Model

Cloud policies and configurations collectively make up 41% of the `POLICY` superclass, or a total of 4,334,301 cyber assets. The percentage of cloud policies is much lower than the corresponding percentage of cloud assets in inventories. This suggests that security policies have not scaled to the cloud in many organizations.

What is troubling, however, is that current cloud policies and configurations aren't working well enough. Recent studies show that 36% of organizations have experienced a data breach due to a cloud misconfiguration. Analysts predict that at least 99% of cloud security failures in 2022 and 2023 will be the fault of enterprise security teams due to cloud resource misconfiguration. The most common causes of cloud security failures include IAM policy failures, poor object storage access policies, or a lack of encryption.

As organizations continue to adopt policy-as-code, cloud policies and configurations are two crucial classes of cyber assets to watch. The number of cloud policies is less important than how effectively these policies cover the attack surface. A policy can have a single statement managing a single cyber asset, or 100 statements providing coverage for thousands of entities.

Relationships Matter Once Again - Policy Edition

Understanding the types of relationships between cloud policy and cloud assets can yield a richer understanding of how large the cloud security gap really is. The relationship class is important. For example, relationships classed as ALLOWS or DENIES around policy can show vulnerabilities or guardrails.

The number of direct and indirect relationships also matters. An S3 bucket exposed publicly via EC2 instance can create a chain reaction of problems based on the first, second, and third degree relationships around the S3 bucket. The misconfiguration is definitely a problem, but it's not the full problem, and understanding blast radius requires a full analysis of cyber asset relationships.

POLICY Can Also Benefit From Automation

It is impossible for security teams to manually manage hundreds of thousands of cloud resources, especially in an environment that practices CI/CD. Instead, it will become increasingly important for DevOps and security teams to partner on effective policy-as-code and adopt automated methods to discover misconfigurations or exploitable vulnerabilities.

What it Means for Security

End-user credentials demand a lot of attention from security teams, and the `POLICY` superclass reveals that a great deal of policy effort is applied to governing identity and access. It is crucial for organizations to use nuanced IAM policy that increases the efficacy of how access permissions, privileges, and asset entitlements are governed.

IAM policy should be proportional to the incredible risk of human behavior, and granular enough to limit the ways in which employees can view, modify, or update critical assets.

Section 5

Why Cyber Asset Relationships Matter

JupiterOne

Relationships

A relationship is the connection between two or more cyber assets. Assets in isolation don't tell the complete story—it's how they interoperate and work together that provides value.

As companies move more of their assets and activities to the digital environment, relationships have become more complex and understanding the connections between assets more important. As a result, when a security incident happens, the data you need often lives in unrelated systems and tooling.

Threat actors have long recognized the importance of relationships. The relationship between an over-privileged user and sensitive assets is how and why social engineering and account takeover are highly successful tactics for threat actors. The blast radius of a realized risk is the product of relationships around a vulnerable asset and includes user permissions, configurations, and integrations.

For the purpose of this analysis, we examined the most common types of relationships between all seven superclasses of cyber assets.

We discovered that 13 relationship types account for 97% of asset relationships within the enterprise. The "everything else" category consists of relationships with less than 0.5% of total relationships, including "runs," "manages," "triggers," "trusts," and so on.

Relationship	Number	Total
Allows	20,334,200	6.22%
Assigned	4,779,331	1.46%
Connects	3,534,677	1.08%
Contains	11,765,864	3.60%
Defines	1,735,270	0.53%
Evaluates	4,334,061	1.33%
Everything Else	9,096,976	2.78%
Has	126,407,445	38.66%
Identified	5,927,587	1.81%
Installed	4,717,232	1.44%
Is	45,086,350	13.79%
Protects	23,619,622	7.22%
Scans	17,564,998	5.37%
Uses	48,059,285	14.70%

JupiterOne

Relationships

Relationships aren't particularly interesting when viewed in isolation. Relationships between assets, on their own, are not inherently secure or risky. After all, even the most critical data store should have a small number of trusted relationships with **`DEVICES`**, **`NETWORKS`**, and **`USERS`**. When viewed in aggregate, however, relationships can reveal patterns of misallocated concern, over-provisioning, or excessive trust.

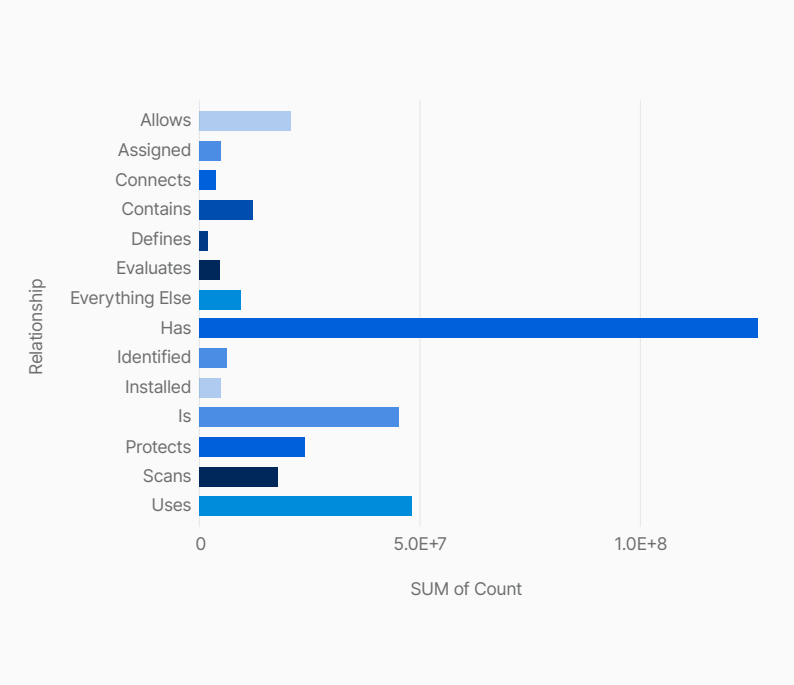

Chart 5.1: Composition of Asset Relationships by Relationship Class

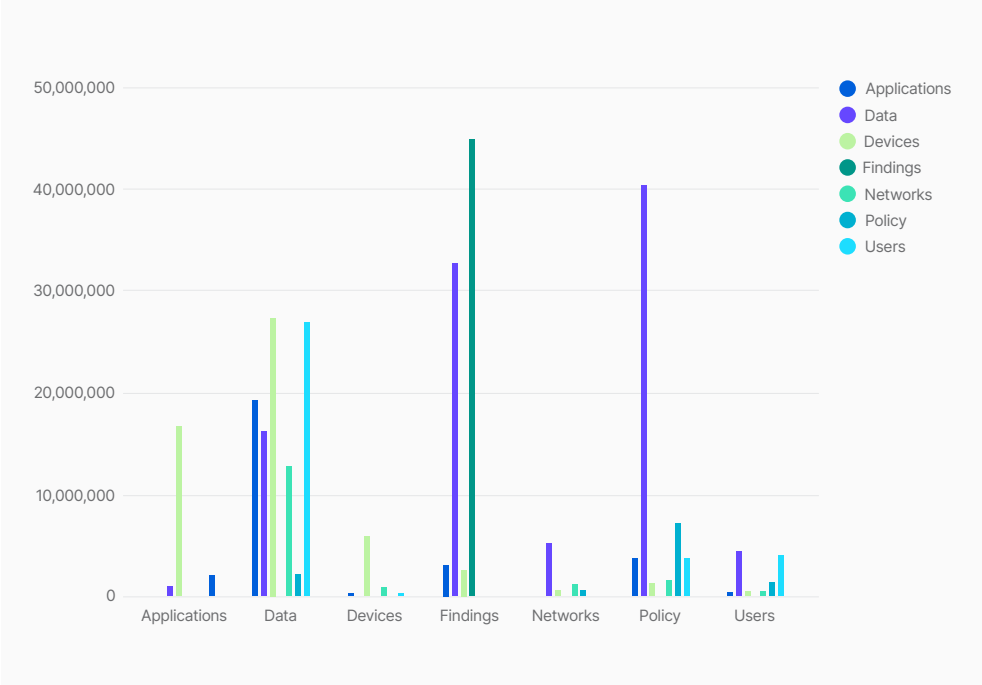

Chart 5.2: Relationships Between Assets and Attributes by Superclass

JupiterOne

Relationships

Traditional security practices warned practitioners of “orphaned assets” and “shadow IT,” or heavily-isolated assets that exist in a cloak of secrecy, outside the reach of administrators and policy. While there is some truth in this mindset, our research shows that the reality is much scarier than the industry has traditionally believed.

Security practitioners significantly underestimate the number of cyber assets that exist beyond the reach of policy and administrators. These assets are not, however, an island. For example, a single user’s unauthorized browser extension connects to the user, and is a few networks and missing IAM policies away from the company’s most crucial assets.

No cyber asset is truly an island; think about how assets in the shadows share **`USERS`**, **`NETWORKS`**, and **`DEVICES`** with the most mission-critical data.

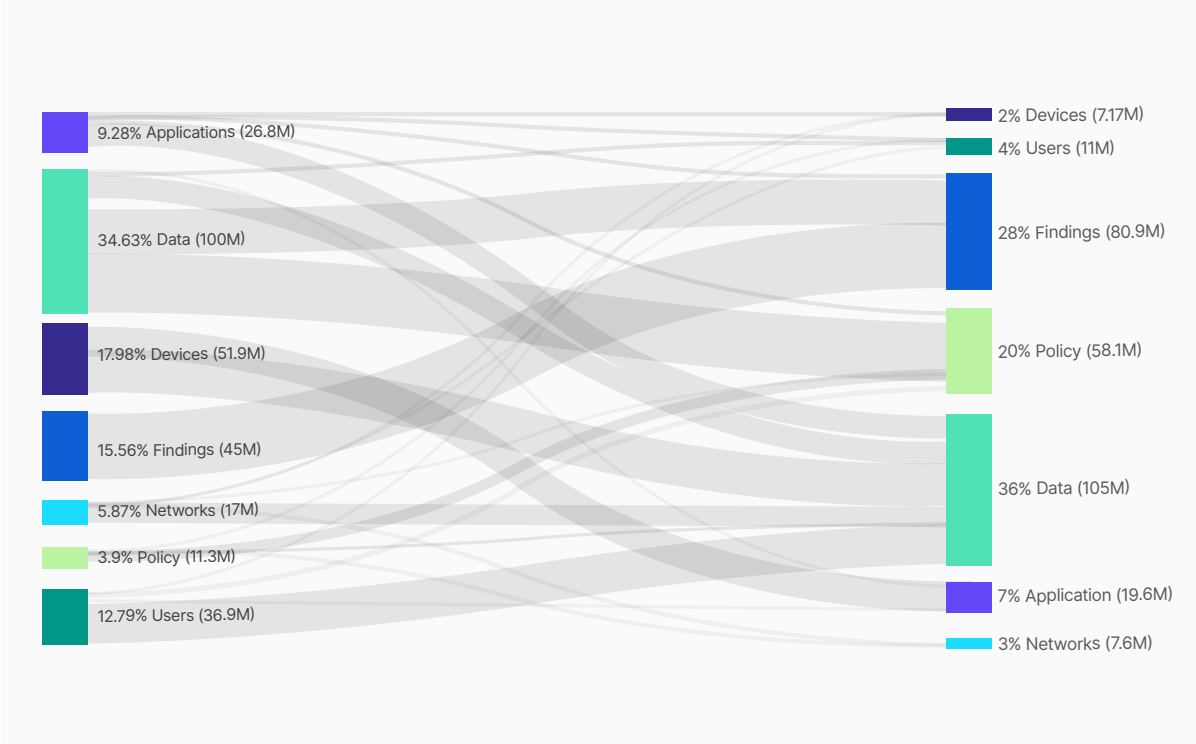

Chart 5.3: Relationships Between Over 370 Million Cyber Assets and Attributes by Superclass

JupiterOne

The Most-Related Assets and Attributes

DATA is the most-related superclass of cyber assets or attributes, followed by **FINDINGS** and **POLICY**. **NETWORKS** have the fewest number of relationships to other assets, along with **DEVICES**, **USERS**, and **APPLICATIONS**.

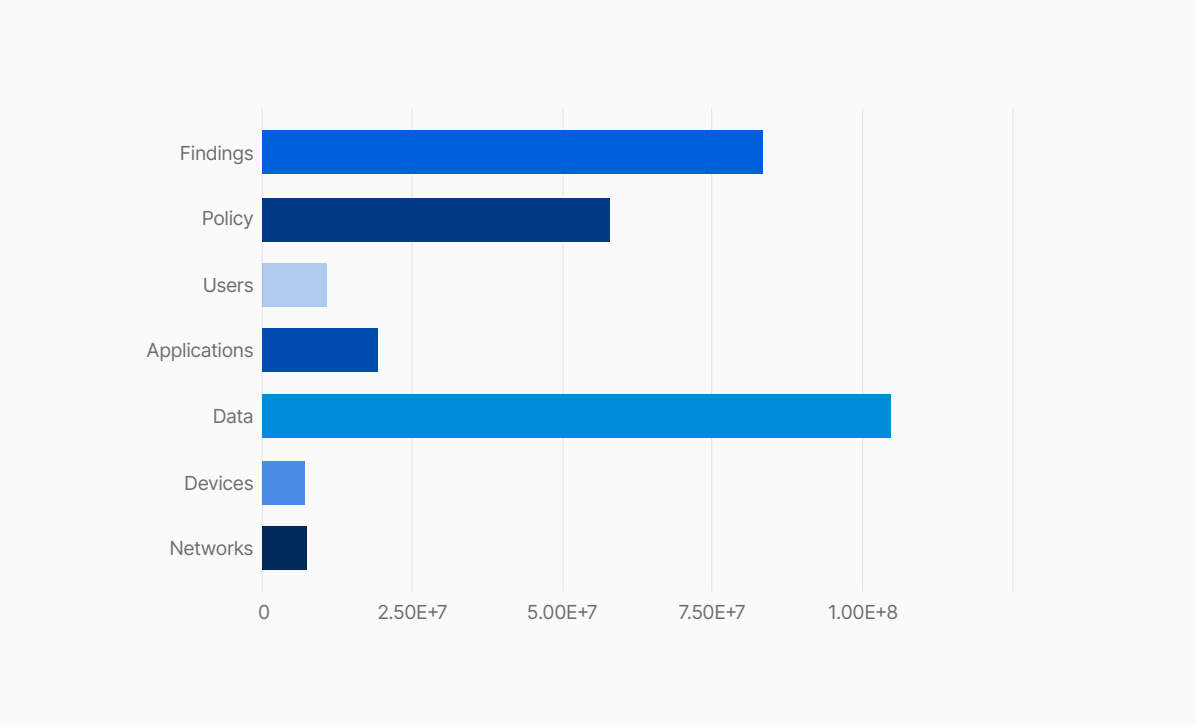

Chart 5.4: Aggregate of All Relationships by Asset and Attribute Superclass

JupiterOne

Findings Relationships

All asset relationships can reveal some insight, but the relationship between **FINDINGS** and other assets are more noteworthy as they offer understanding into the source of alerts, alarms, and findings.

Relationships between **FINDINGS** was the most common category of finding relationship, especially relationships that identify vulnerability scan results as a finding, or a finding as a critical vulnerability (CVE). This shows security teams have made significant progress in integrating their findings with numerous sources of insights to better understand which findings have potential impact or demand action.

Relationships between **FINDINGS** and **DATA** are, however, less positive since they indicate that organization's critical assets are generating an enormous amount of alerts.

*Table 5.2: Count and Relative Frequency of Relationships between **FINDINGS** and Other Classes*

Type	Number	Total
Relationships between **FINDINGS** and **APPLICATIONS**	3,089,171	3.70%
Relationships between **FINDINGS** and **DATA**	32,814,612	39.33%
Relationships between **FINDINGS** and **DEVICES**	2,574,777	3.09%
Relationships between **FINDINGS** and **FINDINGS**	44,964,683	53.89%

*Chart 5.5: Relative Frequency of Relationships between **FINDINGS** and Other Classes*

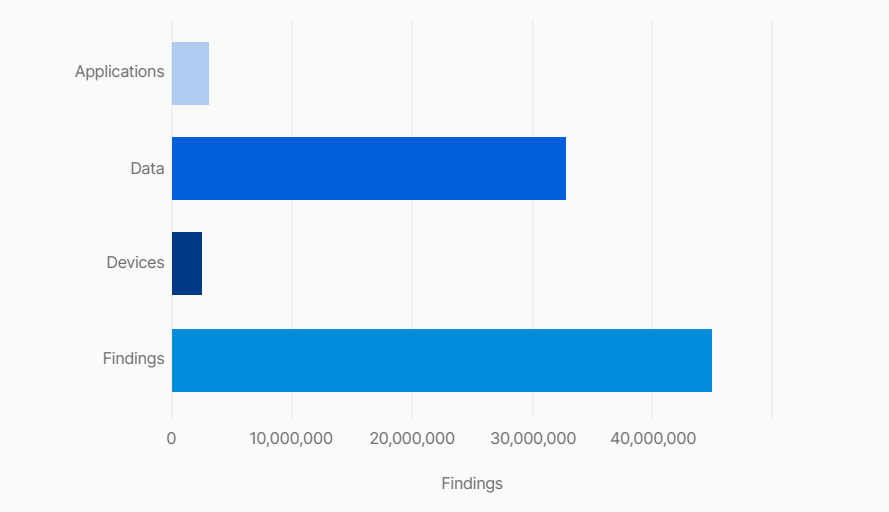

Policy Relationships

Table 5.3: Count and Relative Frequency of Relationships between **POLICY** *and Other Classes*

Type	Number	Total
Relationships between **POLICY** and **APPLICATIONS**	3,810,876	6.56%
Relationships between **POLICY** and **DATA**	40,488,148	69.69%
Relationships between **POLICY** and **DEVICES**	1,270,824	2.19%
Relationships between **POLICY** and **NETWORKS**	1,587,880	2.73%
Relationships between **POLICY** and **POLICY**	7,212,926	12.42%
Relationships between **POLICY** and **USERS**	3,725,182	6.41%

Chart 5.6: Relative Frequency of Relationships between **POLICY** *and Other Classes*

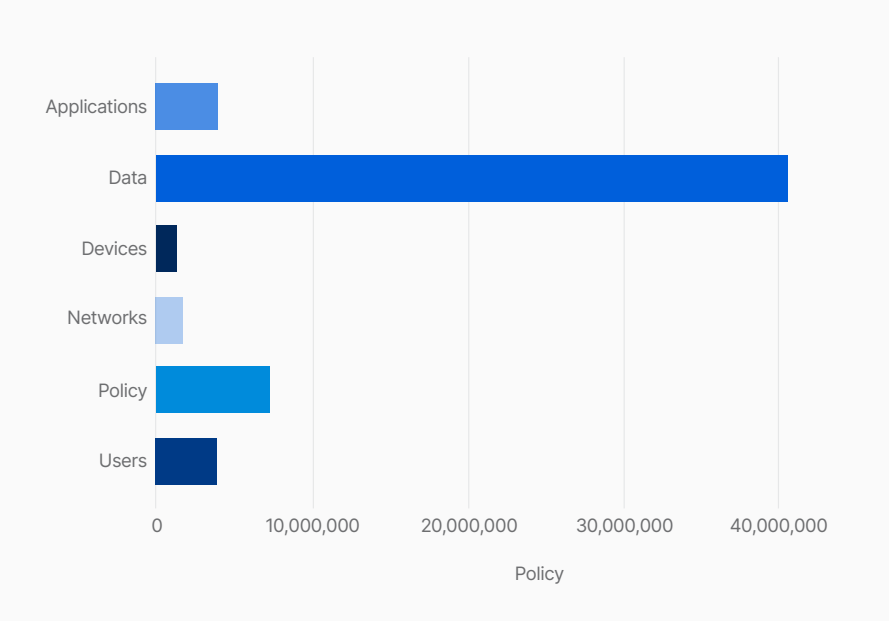

Relationships between **POLICIES** and **ASSETS** are a lovely thing. They indicate that organizations have created guardrails of policies, procedures, and rulesets to protect their sensitive assets. More **POLICIES** indicate better efforts to manage and secure what matters, and more rules to govern the assets with a strong tendency to go rogue (especially humans).

The majority of **POLICIES** exist to govern **DATA**, **USERS**, and **APPLICATIONS**, with a healthy amount of guardrails allocated to **DEVICES** and **NETWORKS**. The only things blatantly missing from this data are **POLICIES** to guide and guard **FINDINGS**, and better understand the types of **FINDINGS** that should create alerts or automated responses.

JupiterOne

User Relationships

USERS relate to almost everything in the enterprise, especially **DATA** and other users.

In the future, we hope to see a greater number of relationships between **USERS** and **POLICY** to illustrate the ever-important role of users who are restricted by policies, accept policies, or bear the responsibility to review a policy at least annually. Sheer number of policies isn't necessarily the best indicator of **POLICY** coverage, so future analysis is needed to understand whether any correlation exists between policy, policy relationships, and security posture.

The relationship between humans and **FINDINGS** is another area that's critically missing and worth further exploration, especially if security teams grow to the point where they can better manage their backlogs of findings.

Type	Number	Total
Relationships between **USERS** and **APPLICATIONS**	394,886	3.60%
Relationships between **USERS** and **DATA**	4,333,933	39.46%
Relationships between **USERS** and **DEVICES**	464,801	4.23%
Relationships between **USERS** and **NETWORKS**	489,037	4.45%
Relationships between **USERS** and **POLICY**	1,276,726	11.63%
Relationships between **USERS** and **USERS**	4,022,872	36.63%

*Table 5.4: Count and Relative Frequency of Relationships between **USERS** and Other Classes*

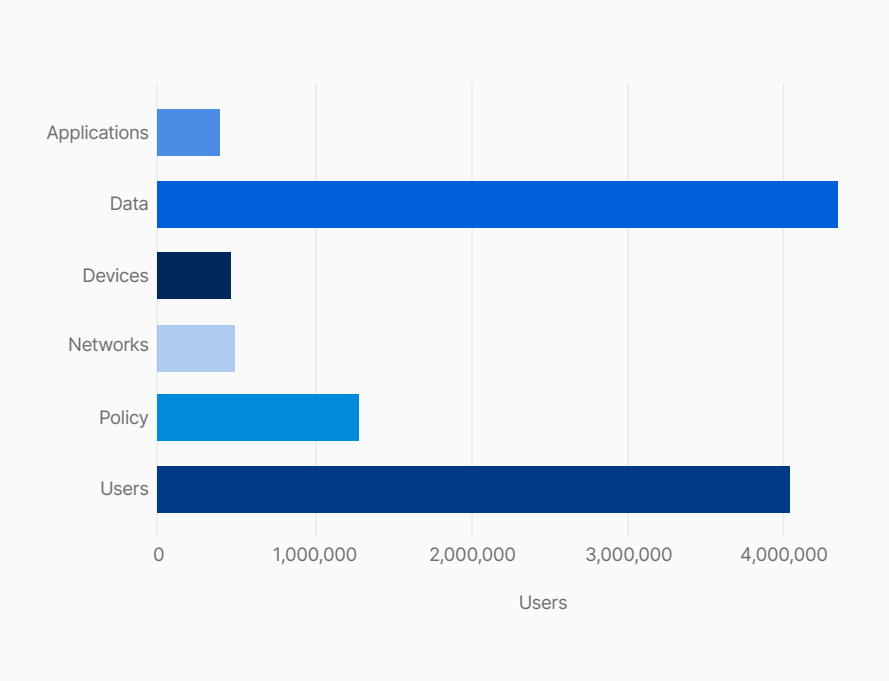

*Chart 5.7: Relative Frequency of Relationships between **USERS** and Other Classes*

JupiterOne

Application Relationships

Type	Number	Total
Relationships between `APPLICATIONS` and `DATA`	1,000,300	5.11%
Relationships between `APPLICATIONS` and `DEVICES`	16,572,918	84.70%
Relationships between `APPLICATIONS` and `USERS`	1,992,467	10.18%

Table 5.5: Count and Relative Frequency of Relationships between `APPLICATIONS` and Other Classes

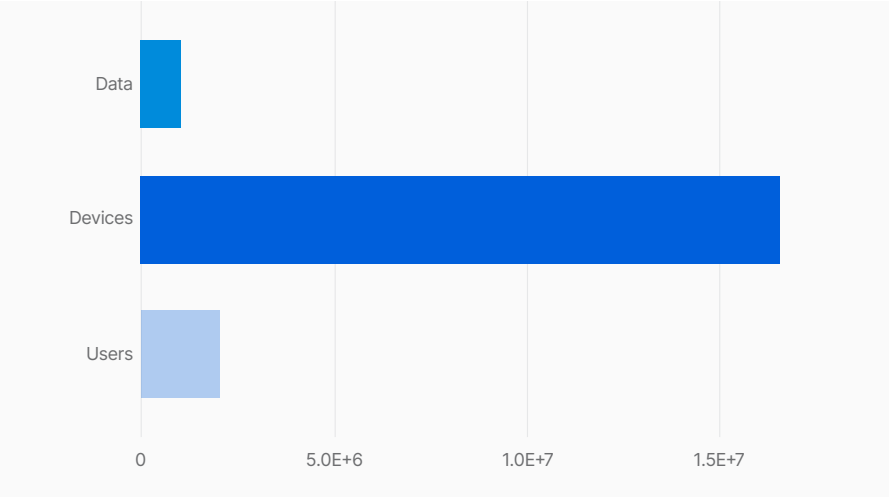

Chart 5.8: Relative Frequency of Relationships between `APPLICATIONS` and Other Classes

JupiterOne

Data Relationships

*Table 5.6: Count and Relative Frequency of Relationships between **DATA** and Other Classes*

Type	Number	Total
Relationships between **DATA** and **APPLICATIONS**	19,258,579	18.40%
Relationships between **DATA** and **DATA**	16,214,785	15.49%
Relationships between **DATA** and **DEVICES**	27,325,185	26.11%
Relationships between **DATA** and **NETWORKS**	12,769,825	12.20%
Relationships between **DATA** and **POLICY**	2,165,472	2.07%
Relationships between **DATA** and **USERS**	26,914,459	25.72%

*Chart 5.9: Relative Frequency of Relationships between **DATA** and Other Classes*

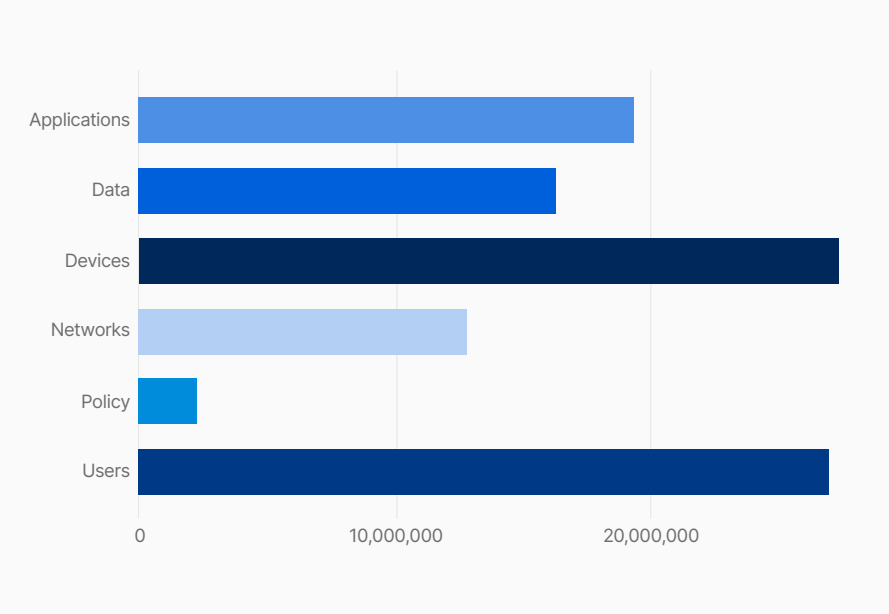

There are really relatively few relationships between **DATA** and **POLICY**, especially in proportion to other **DATA** relationships. This superclass is likely to contain a number of mission-critical assets for many organizations, including records that contain personally identifiable information (PII) or personal health information (PHI). Additionally, we must consider the importance of data retention policy for a business to meet legislative, contractual, and regulatory obligations.

There are significantly more relationships between **POLICY** and **DATA**, in which policy drives the data, than the inverse which is explored in this section. The relationships where policy is driven by data is low at 2% of total. The resultant mean is 1,705 data-driven policies per organization.

For organizations to appropriately classify, protect, and retain their **DATA** at such scale as is reflected in this research, organizations need to adopt new ways of creating policy based on data. In practice, this could mean classifying data according to specific criteria such as the organization's unique risk appetite, business requirements, and the characteristics that are common to their most sensitive data.

Device Relationships

Table 5.7: Count and Relative Frequency of Relationships between `DEVICES` *and Other Classes*

Type	Number	Total
Relationships between `DEVICES` and `APPLICATIONS`	267,518	68.35%
Relationships between `DEVICES` and `DEVICES`	5,750,857	7.34%
Relationships between `DEVICES` and `NETWORKS`	866,655	16.37%
Relationships between `DEVICES` and `USERS`	287,285	7.94%

Chart 5.10: Relative Frequency of Relationships between `DEVICES` *and Other Classes*

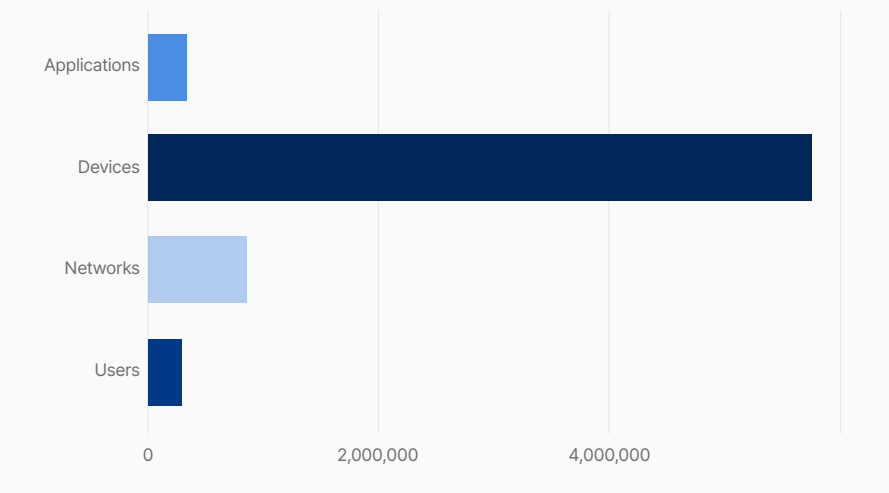

`DEVICES` are among the least-related types of cyber asset. `DEVICES` have understandably few relationships between `APPLICATIONS`, `USERS`, and `NETWORKS`, which could indicate that organizations are relying on a limited number of trusted usage scenarios to monitor against insider abuse.

What is notable, however, is the high number of relationships between devices and the fact that the `DEVICES` superclass is composed of cloud hosts. This is another indicator of the sophisticated, serverless infrastructures being deployed at modern organizations and the architectural shift towards microservices and resiliency. It's a wondrous thing, until you consider the implications of governing the soaring number of host-to-host relationships, which can exist without the oversight of humans.

Indirect Relationships are Surprisingly Important

Device relationships is an accessible example of why indirect relationships matter. A device might only be used by one user. However, if you connect the device back to all the digital identities of that user and their access to SaaS, cloud, and data, the device has dozens of indirect relationships.

Indirect relationships are known as extended traversals among graph theory nerds, but they're important to everyone with an interest in security or cyber assets. Indirect relationships allow practitioners to determine the blast radius, or impact, of an asset if a device or account is compromised.

No spoilers, but you can expect the JupiterOne research team to drop some in-depth analysis into asset relationships in the relatively near future.

Network Relationships

Table 5.8: Count and Relative Frequency of Relationships between `NETWORKS` *and Other Classes*

Type	Number	Total
Relationships between `NETWORKS` and `DATA`	5,191,469	68.35%
Relationships between `NETWORKS` and `DEVICES`	557,632	7.34%
Relationships between `NETWORKS` and `NETWORKS`	1,243,146	16.37%
Relationships between `NETWORKS` and `POLICY`	602,871	7.94%

Chart 5.11: Relative Frequency of Relationships between `NETWORKS` *and Other Classes*

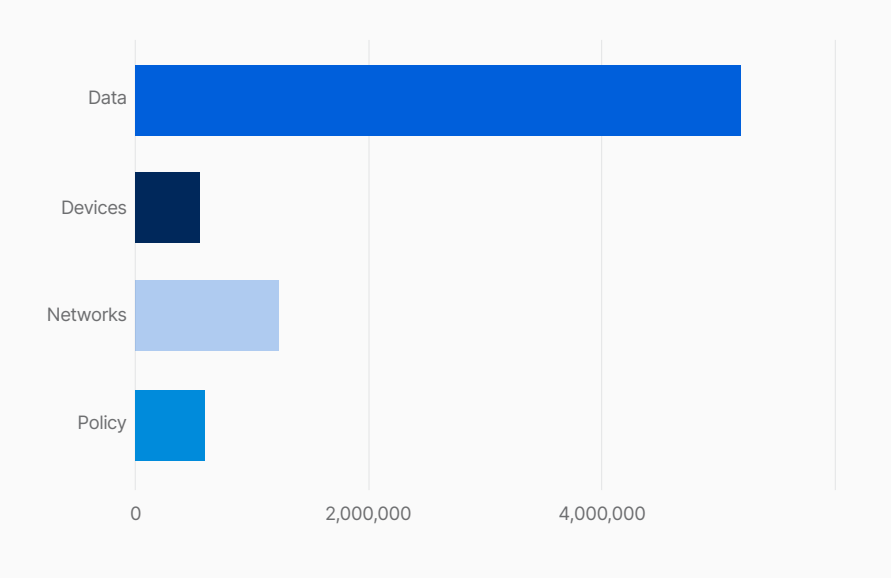

`NETWORK` relationships was a category that held few surprises. The idea that over 68% of network-driven relationships create data is logical when you begin to consider flow data, packet data, and APIs. While there are significantly fewer relationships between `NETWORKS` and `DEVICES` or `POLICY`, there's evidence that network state has influence over `POLICY` and connected networks.

Section 6

Top Queries — What Do Security Teams Care About?

Query Data

The data included in the analysis of queries is derived from log performance data among users of JupiterOne's Cyber Asset Attack Surface Management (CAASM) over a one-week period of time, September 28-October 5, 2021. This data is largely reflective of queries that users created to generate alerts on negative changes on their environment, as well as queries created as a single-use examination of their asset relationships.

In isolation, this fully-anonymized query data can provide an understanding of what security practitioners care about, since these queries are generally written by individuals in security leadership (CISO), security engineering, or DevOps / DevSecOps roles. When query data is viewed in comparison to the realities of assets and asset relationships from a near-identical period of time, it's possible to create some understanding of blindspots and invisible assets.

The queries included in this report are a combination of searches into asset inventory and relationships, along with alerts created by enterprise security teams. **DEVICES** are the most common target of human-generated cyber asset query. Further analysis of this category reveals efforts to understand the relationship between **DEVICES** and **FINDINGS**, **DEVICES** and **POLICY**, as well as **DEVICES** and **USERS**.

USERS are another common target of human-generated queries and alerts, which typically sought to understand the resources that employees could access, the **POLICY** that governed human behavior, and human relationships **DATA**, **NETWORKS**, or **APPLICATIONS**.

Superclass	Count	Total
(All)	340,575	8.83%
APPLICATIONS	275,985	7.16%
DATA	515,847	13.38%
DEVICES	1,088,813	28.24%
FINDINGS	452,216	11.73%
NETWORKS	242,984	6.30%
POLICIES	326,958	8.48%
USERS	612,497	15.88%

Table 6.1: Composition of Top Queries by Asset Superclass

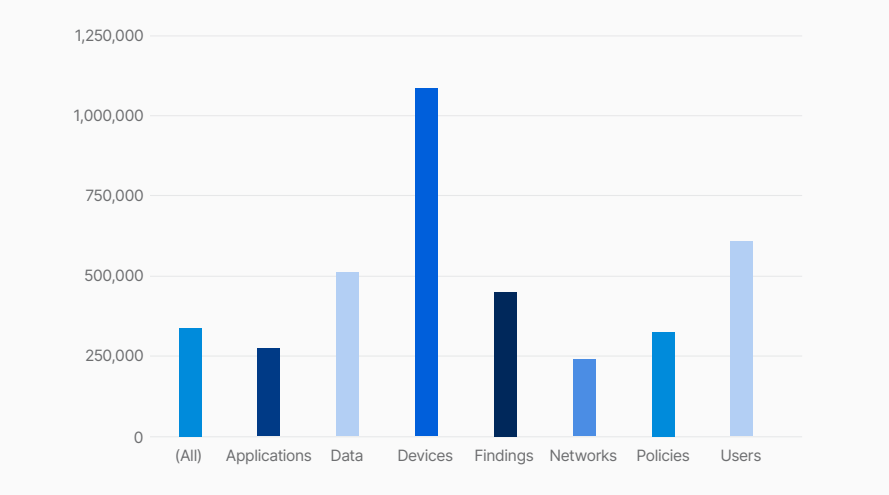

Chart 6.1: Relative Frequency of Top Queries by Asset Superclass

JupiterOne

Top Queries by Relationship Category

Table 6.2: Composition of Relationship Class Included in Queries

Relationship	Number	Total
Allows	143,328	3.70%
Assigned	180,519	4.66%
Connects	21,537	0.56%
Everything Else	289,129	7.46%
Has	164,685	4.25%
Implements	37,393	0.96%
Is	33,963	0.88%
Logs	42,854	1.11%
Monitors	81,404	2.10%
Protects	271,473	7.00%
Relates	2,543,605	65.60%
Trusts	22,380	0.58%
Uses	45,346	1.17%

Chart 6.2: Relative Frequency of Relationship Class Included within Queries

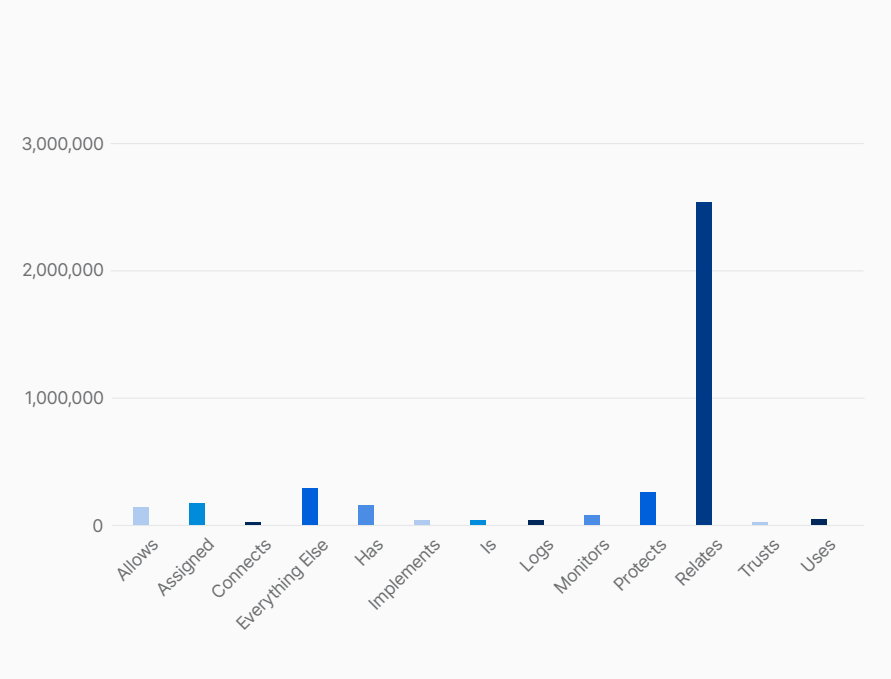

The “everything else” category of queries reflects queried relationships with less than 0.5% frequency, in proportion to total queries. Most common queries within the “everything else” category refer to relationship classes such as performed, denies, implements, and opens.

The incredible diversity of relationship classes is reflective of the enterprise technology stack. Relationship classes are largely defined by technology vendors and less often defined by security teams.

JupiterOne

Queries vs. Reality: The Questions We Ask and What's Really in Our Asset Inventories

Table 6.3: Composition of Queries vs. Assets vs. Relationships by Class

Superclass	Percent of Queries	Percent of Cyber Assets	Percent of Cyber Asset Relationships
APPLICATIONS	7.16%	5.43%	6.71%
DATA	13.38%	20.48%	35.90%
DEVICES	28.24%	10.99%	2.46%
FINDINGS	11.73%	40.69%	28.63%
NETWORKS	6.30%	7.61%	t2.61%
POLICY	8.48%	2.85%	19.93%
USERS	15.88%	11.96%	3.77%

Chart 6.3: Composition of Queries vs. Assets vs. Relationships by Class

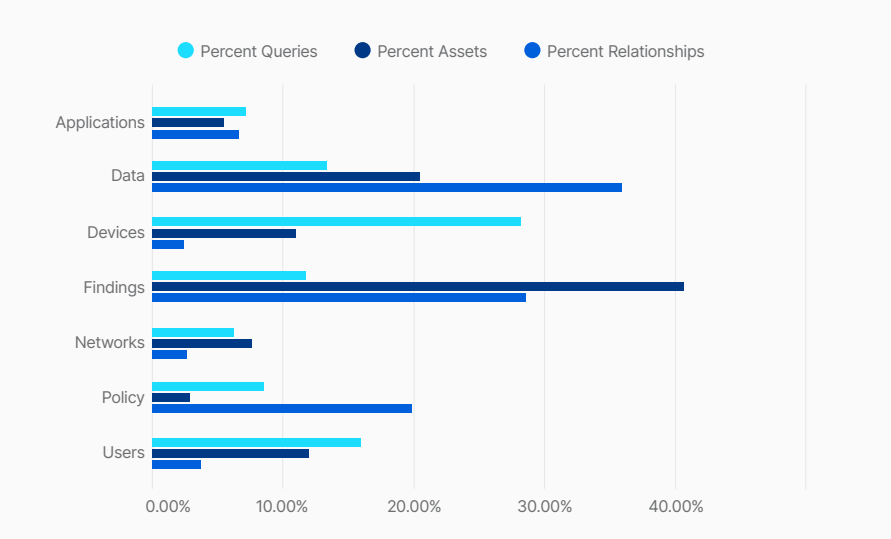

Queries and alerts don't always match the composition of our asset inventories or relationships. In fact, our efforts to understand how queries matched up to reality found significant differences between what security teams care about and what's actually expanding their attack surfaces.

DEVICES received a disproportionate amount of attention from security teams, especially when viewed in conjunction with the actual occurrence of device cyber assets and device relationships.

The percent of queries dedicated to understanding **DEVICES** outpaced it's asset inventory proportion by nearly 18%.

USERS received a lot of attention, but this is likely warranted and reasonable considering the disproportionate impact of humans on enterprise security posture. 16% of queries were interested in **USERS**, compared with a 12% occurrence in the proportion of cyber asset inventories.

The superclasses of cyber assets where queries are significantly fewer than actual occurrence of assets or relationships could indicate blindspots among cybersecurity teams. **DATA** and **FINDINGS** are both a significantly lower percentage of queries than actual proportional occurrence. **NETWORK** query proportion was also overtaken by actual occurrence.

JupiterOne

Traversals

In order to understand traversals and their significance to our analysis, we must return to graph theory. Each edge, or relationship, that is included in a query is counted as a traversal for the purpose of this analysis.

- For example, a query that sought to identify all **`USERS`** or all **`POLICIES`** would have zero traversals.
- A query to determine whether a vulnerability scanner finding was a CVE would include one traversal.
- A query to determine whether a cloud host had any vulnerability scanner findings and whether these findings were a CVE would include two traversals.

Traversal analysis reveals a lot about the security practitioner's ability to use asset relationships to identify blast radius. Security teams understand that virtually no assets exist in total isolation and that the most actionable **`FINDINGS`** are the ones that assess the unique, toxic combinations of misconfigurations and overprovisioned access.

Table 6.4: Count and Relative Frequency of Query Traversals

Number of Traversals	Query Count	Total
0	1,045,561	60.25%
1	544,521	31.38%
2	81,808	4.71%
3	53,188	3.06%
4	1,470	0.08%
5	8,368	0.48%
6	479	0.03%
7	48	0.00%
9	41	0.00%
14	7	0.00%

JupiterOne

Traversals

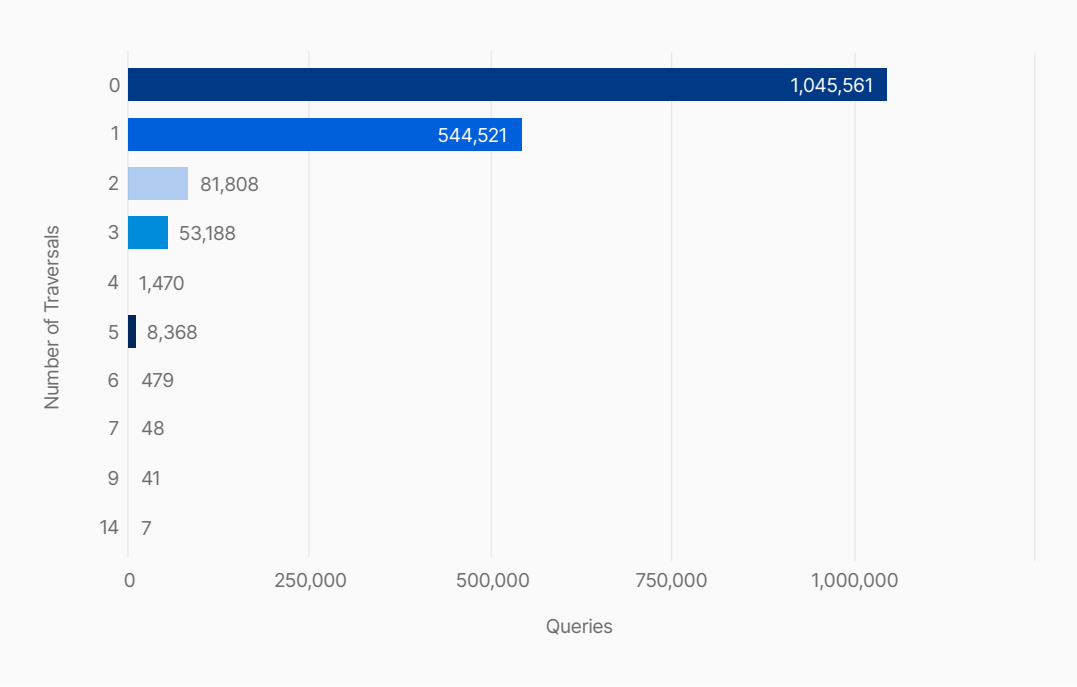

Chart 6.4: Relative Frequency of Query Traversals

The relative rarity of queries and alerts with more than two traversals also reveals that not all security practitioners are taking full advantage of the relationship-based graph model for attack surface management. Queries with 3 or more traversals are collectively a low 3.7% of total, despite the potential of such highly-specific queries.

Using more complex queries allows security practitioners to narrow results by criteria associated with the most critical assets, such as asset class, age, policy, owner, excessive entitlements, or outdated versions.

What it Means for Security

The fact 91% of queries and alerts include one relationship or fewer reveals an industry that is trying to understand the most basic characteristics of their cyber assets. This is likely a natural byproduct of the scale of cyber asset inventories and the advent of policy-as-code. More than ever before, security practitioners are removed from the asset lifecycle until resources are deployed to production.

There is no easy solution to the realities of modern attack surface management, especially if you're trying to draw from traditional approaches to security. The only solutions are greater collaboration between security and DevOps to create more guardrails along the cloud asset lifecycle, and more industry recognition of why cyber asset relationships are so vital to manage contemporary attack surfaces.

Section 7

In Conclusion

In Conclusion

JupiterOne

Interesting Points

We've done a deep dive into understanding cyber asset classification and how assets are related to one another. From our dataset this year, we've been able to grasp where potential cyber asset risks are and the liabilities associated with various assets.

In conclusion, we want to surface and summarize the most interesting points from our research and analysis:

1

Automated asset inventory is mandatory

Automation is a necessity for security teams to manage a vast, fast-changing attack surface. An added bonus is for teams to build or create systems where all of their cyber asset metadata can flow into one view to empower complete visibility and accelerate actions. Modern assets are created through automation, so any part of asset management that is done manually is doomed to fail.

2

Digital transformation has truly changed everything

The adoption of cloud services, resilient architectures, and agile development life cycles has created a cyber asset mix where traditional IT infrastructures are a tiny percentage of assets. We need to rethink everything, including our commitment to the skills pipeline, policy, and collective definition of best practices as an industry.

3

Understanding cyber asset relationships is crucial to effectively manage risk

Assets alone reveal little actionable intelligence—especially compared to relationship-based analysis—to understand relationships between assets, findings, users, and policies. Relationships between assets drive security context and increased security context is how security programs improve.

4

`FINDINGS` merit a deeper look to manage vulnerabilities

Since `FINDINGS` account for more than 41% of enterprise assets and attributes, this affects how security teams are resourced and manage vulnerabilities. Security practitioners could experience alert fatigue or miss out on true positives unless the industry increases the amount of talent and technology to help triage, investigate, remediate, and tune findings.

5

Queries are often misaligned with actual asset inventories

The queries and alerts created by security practitioners do not reflect the actual composition of modern cyber asset inventories. Security practitioners are focused on `USERS` and `DEVICES` while directing relatively little attention toward `DATA` and `FINDINGS`. The relatively low number of relationships in most queries and alerts also indicates that security practitioners are working to understand what they have, not to understand which assets are most business-critical or vulnerable.

Appendix A: Methodology

Ultimately, the SCAR intends to be a rigorous resource by security practitioners and researchers, for other practitioners. This report was created with an appropriate level of rigor and attention-to-detail that the security community should demand.

Approach to Error and Corrections

We acknowledge that the authors and reviewers of this report are human, and accordingly, imperfect. Despite significant review, it is entirely possible that mistakes could happen due to human error. Any errors discovered will be promptly corrected in the report, and noted in the Appendices in the future in a section labeled "Appendix C: Corrections."

We believe that peer review is crucial, which is why we are providing SCAR resources such as a partial data set, graphs, and our entire cyber asset relationship graph data model on Github for readers. We encourage you to check our analyses or find your own stories within the data for your own research and reporting, while providing credit to JupiterOne as your data source.

*Will be available by May 1, 2022

Disclaimer

No claims are made that the research in this report is completely representative of all organizations worldwide. Claiming to have a perfect understanding of cyber assets at organizations of all sizes and industries would be unscientific and unreasonable. Instead, we simply claim to have analyzed a large data sample.

We believe that many of the findings are appropriate for general application to other organizations, especially cloud-native organizations, but we do not claim total global representation or a complete absence of bias. Readers of this report (and all other reports) are encouraged to be objective and critical of methodology applied.

While we have exerted many possible controls against bias such as rigorous review of this research by our peers and transparency into our data, bias nearly always exists and it's healthy to acknowledge this fact.

Acknowledgment of Selection Bias

The data sample analyzed for SCAR is sourced from the organizations who use JupiterOne's Cyber Asset and Attack Surface Management (CAASM) product. Accordingly, selection bias is possible based on the organizations who find value in JupiterOne's product.

Our customers are generally cloud-native, which means the organizations analyzed for the SCAR could have a lower number of legacy systems and on-premises-based deployments than the true, global mean. This is due to the fact that CAASM products such as JupiterOne are generally designed for the asset and attack surface management requirements of cloud-native architectures.

Notably, the data in the SCAR sample could be further limited based on the customer's chosen integrations with JupiterOne's CAASM product since not all customers integrate all of their systems. Integrations typically include all IaaS and PaaS products, but not necessarily all SaaS, especially not SaaS that fall outside the administration of the security team (e.g., our customers may not always integrate their CRM or marketing asset management SaaS with JupiterOne). Also, not all customers choose to create integrations between JupiterOne and homegrown systems, or integrations with legacy systems.

Some customer data on findings may be omitted as well, since JupiterOne offers customers the option to filter Qualys integration findings for severity medium or higher. Other sources of findings, such as AWS Guardduty or Inspector, are not similarly filtered. Future editions of the SCAR will attempt to explore the distribution of findings by severity, to better understand the proportion of critical, high, medium, low, and informational findings.

Appendix A: Methodology

In short, the data reflected here is likely impacted to some extent by selection bias, and the result is likely a more cloud-native mixture of cyber assets than the reality of cyber assets at organizations worldwide. We hope to better understand, articulate, and control for possible selection bias in future iterations of the SCAR.

The SCAR Process

The collection and conversion method consisted of the direct recording of cyber asset inventory data for customers of the JupiterOne CAASM solution. A point-in-time snapshot of customer data was captured and recorded into a data lake for querying from a single data capture on October 14, 2021.

Significant and reasonable efforts were made to protect customer anonymity and avoid exposure of critical data to our data science and analyst teams, by ensuring that only sanitized data was included in the data used for analysis. No data that reveals customer sensitive information was included, with "critical" data made unavailable via access controls. JupiterOne's approach to classifying, managing, and protecting the confidentiality of customer data is described in the following documents:

The JupiterOne Data Model
JupiterOne Data Management Policy
JupiterOne Data Protection

All contributors to the SCAR report were required to accept these policies. While IAM controls barred critical data from exposure to analysts, analysts were further instructed to be conservative with what data was used in this report. As such, we may have omitted the analysis of some data that probably could have yielded more interesting insights.

Source data was aggregated into a report by querying the data lake, and then transferred into a spreadsheet for greater consistency and control over granular analysis by asset class. Several analysts worked collaboratively to ensure consistency of analysis efforts, including clustering and grouping of sub-classes within each super class.

JupiterOne's data model allows a single cyber asset to fit multiple asset classes. Similarly, a single cyber asset may enter JupiterOne's graph multiple times due to redundancy within integrations. For example, a laptop may be reflected in both a device management integration and an endpoint detection and response integration. Care was taken to deduplicate all assets with multiple classifications and ensure they were only included once in the final analysis. Several data analysts and security practitioners worked collaboratively to review all groupings of assets with multiple classes, and categorize them into superclasses once according to best fit.

After the creation of data tables and basic graphs, the SCAR analysis was subject to interpretation by the authors, who shared a common lens as long-term security practitioners. The resulting graphs, tables, and written analysis were subjected to several weeks of rigorous peer review by analytics experts, data scientists, security practitioners, and engineers to ensure the data and interpretations were fair.

Appendix A: Methodology

Firmographics

The SCAR analysis looks at organizations of all sizes. Future editions of the SCAR hope to provide additional analysis of assets and attributes within each firmographic category or analysis of how cyber assets vary by company size and industry. Such analysis is only possible or pragmatic with statistically significant sample sizes within each firmographic category.

Firmographic data was not available for all of the organizations included in this report due to a number of factors, including the fact JupiterOne offers a robust freeware version of it's product and company privacy preferences. Accordingly, firmographic data was available for a total of 96 out of 1270 organizations in our analysis, or 7.6% of total.

Company size and industry data was available for 96 of the 1270 organizations included in this research.

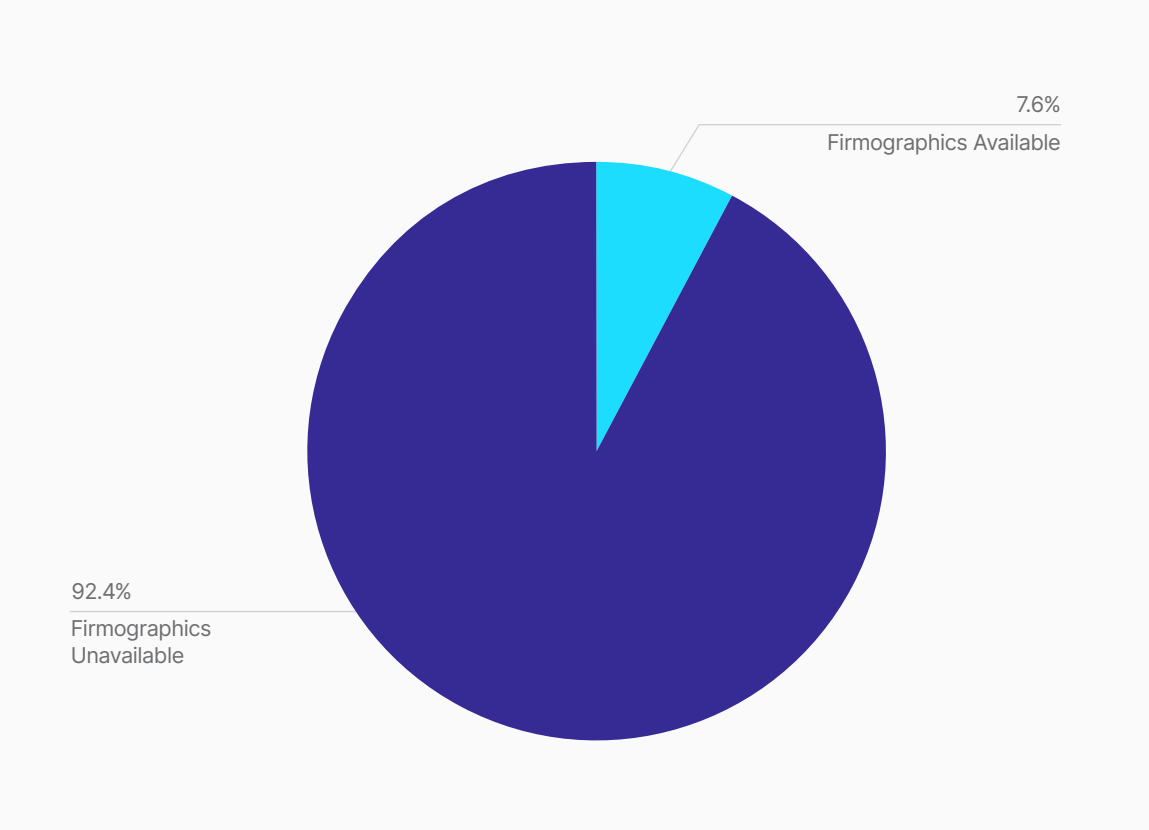

Chart A.1: Relative Frequency of Included Organizations with Firmographic Characteristics and without Firmographic Data

Appendix A: Methodology

Chart A.2: Relative Frequency of Organizational Size for Included Organizations with Available Firmographic Data

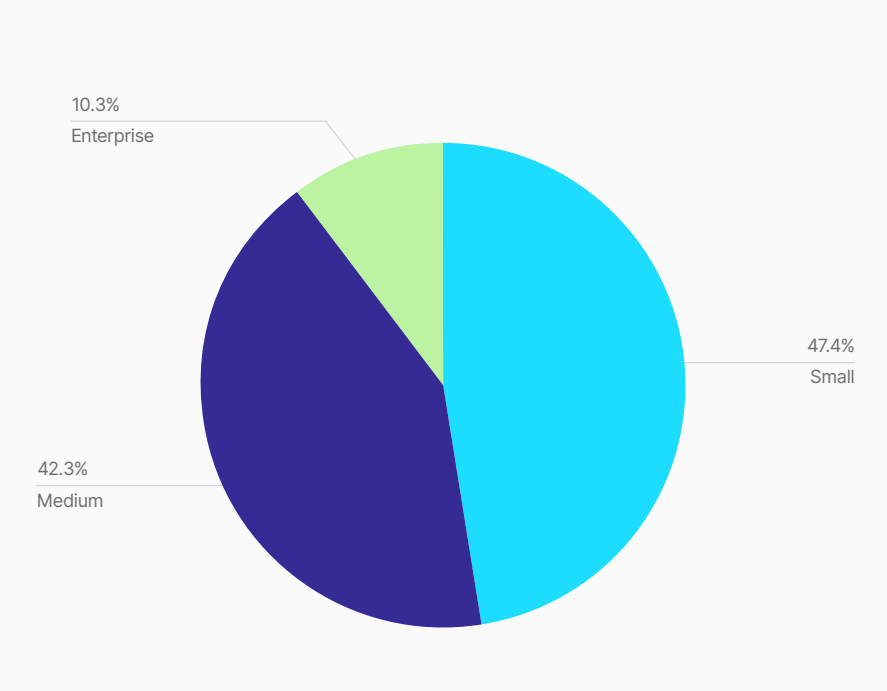

Chart A.3: Relative Frequency of Industries for Included Organizations with Available Firmographic Data

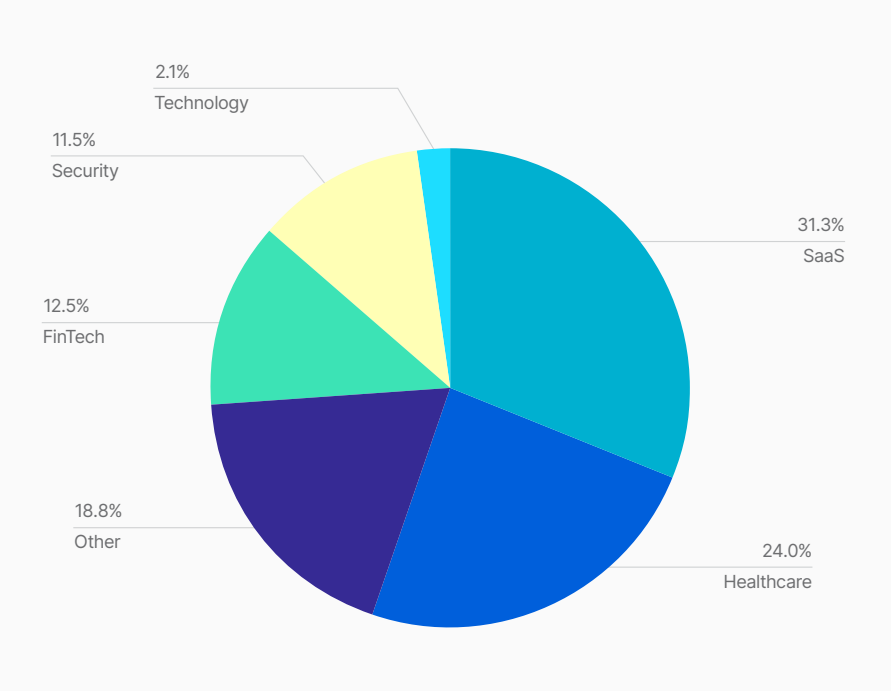

Firmographics by Company Size

Among the 96 organizations with firmographic characteristics, 10% were enterprises, and 42% were medium-sized or mid-market organizations. 47% of the organizations represented in this research with firmographic characteristics were small organizations.

Firmographics by Industry

Among the 96 organizations with firmographic characteristics, the most common industry was software-as-a-service (SaaS) at nearly 30% of total.

Other prominently represented industries include healthcare (24%), Security (11.5%) and FinTech (12.5%).

Appendices

Appendix A: Acknowledgments

The authors would like to express their sincere gratitude to the many people who contributed to this research or assisted in the review process. This report is only possible with the efforts of many people and we are grateful for their support.

SCAR Contributors

Significant contributors to this report include the individuals responsible for the design, editing, promotions, and ideation:

- Ashleigh Lee, Senior Marketing Manager at JupiterOne
- Mark Miller, Senior Director of Content Marketing at JupiterOne
- Melissa Pereira, Director of Communications at JupiterOne
- Kevin Swartz, Director of Digital Marketing & Demand Generation at JupiterOne
- Tanvi Tapadia, Marketing Associate at JupiterOne
- Chum Wongrassamee, Head of Design at JupiterOne
- Dave Moy, Senior Graphic Designer at JupiterOne
- Valerie Zargarpur, Senior Marketing Manager at JupiterOne

SCAR Reviewers

Also, we would like to acknowledge the following people for their review, feedback, and encouragement:

- Chasen Bettinger, Senior Security Automation Engineer at JupiterOne
- Phil Gates-Idem, Chief Architect at JupiterOne
- Tyler Shields, CMO at JupiterOne
- Erkang Zheng, CEO at JupiterOne
- Anne Nielsen, Head of Product at JupiterOne

SCAR Readers (That's You!)

This report highlights the incredible dedication of the worldwide community of security practitioners. We are grateful to be part of such an exceptional, dedicated community. Our goal is to help our security peers advocate for necessary resources and recognition, and move toward a healthier and more secure future for all of us.

Thank you for reading this report, citing it, and sharing your feedback on how we can improve future editions of this research. We are listening, receptive to your thoughts, and you are encouraged to reach us at scar@jupiterone.com.

Cite the SCAR

You are permitted to use statistics, figures, and other information from this report, provided that you cite the source as JupiterOne 2022 State of Cyber Assets Report (or 2022 JupiterOne SCAR) and do not modify the content in any way. Exact quotes are permitted. If you would like to link to the report, we ask that you link to jupiterone.com/scar.

SCAR

JupiterOne.com