



The Anatomy of an API

2023 EDITION



Table of Content

Introduction	3
Key Findings	6
Methodology	8
API Design	10
API Performance	17
API Security	21
API Market	25
Conclusion	33



Introduction

API usage is exploding. That's just a fact. Everyone can feel it, but more importantly, numerous reports and surveys back that up. Without going into too many details, we'll just highlight a couple of key numbers:

90%

of developers
use APIs ⁽¹⁾

25.1%

API management is growing at
a 25.1% CAGR ⁽²⁾

83%

of all internet traffic
belongs to API ⁽³⁾

Today, APIs are fundamental building blocks of every modern company, business, and product. **The web has enabled businesses to enter the digital world, but APIs are unlocking their full potential.** APIs allow businesses to leverage the power of data they already have: data that might be sitting idle or siloed off from the rest of the world completely. With APIs, that data can be put to work and even dramatically amplified by connecting it to different data sources. That's the true power of APIs: enhancing interoperability, promoting innovation, and unlocking new revenue streams.

Proof points that investing in APIs is a good idea have existed since the 2000s when companies like Salesforce, Amazon, eBay, Google, and Facebook started building the first commercial APIs. Today, those companies have a combined market cap of 4 trillion dollars and are consistently ranked as the top companies in the world. More recent examples include companies like Stripe, Twilio, Zapier, and others pushing the API economy forward while generating massive amounts of shareholder value.

But building good APIs is hard, and building a business around APIs is even more complicated. The perfect analogy to this is what happened to the web recently.

⁽¹⁾ "State of the Developer Nation 19th Edition" Accessed November 27, 2023. <https://www.developernation.net/resources/reports/state-of-the-developer-nation-q3-2020>

⁽²⁾ "API Management Market Size, Industry Share Forecast." Accessed November 27, 2023. <https://www.marketsandmarkets.com/Market-Reports/api-management-market-178266736.html>

⁽³⁾ "State of the Internet Reports - Akamai." Accessed November 27, 2023. <https://www.akamai.com/our-thinking/the-state-of-the-internet>.


As companies rushed to build websites, the world came to an inflection point around 2015 and realized that most websites are poorly developed and slow. Mostly because people weren't concerned about quality, nor was there enough data to help people understand what it means to build a good website.

We think history is repeating itself with APIs and want to help change that. Fundamentally, with this report, we tried to answer one seemingly simple question:

————— **How does an average API look like in 2023?**

By understanding what an API looks like today, we can understand where we are as an industry, where we want to go, and, more importantly, how we get there. It's always hard to take inventory of an ecosystem that has been growing for years, especially without a real standard to lean on. We think surveys are not as efficient because they rely on people and not data. People tend to make the data more subjective than it is - it's simply in our nature. That's why we decided to publish this report based on objective, actual APIs, and data flowing through them. That's where our platform, Treble, comes into play.

Treble is an **end-to-end APIOps** platform that helps businesses build, ship, and scale APIs. It enables customers to efficiently address the top 3 API pain points:



- 01** Visibility and understanding of API data
- 02** Governance and API standardization
- 03** API security

These three pain points cost businesses billions of dollars annually in lost revenue, trust, innovation, and customers.

Businesses integrate Trebble through a simple SDK/Agent available for more than 20 different programming languages, platforms, and gateways. The integration takes less than 3 minutes and, out of the box, gives customers six core products:

API Observability

Captures requests in real-time and generates more than 40 data points for every single one of them

API Governance

Runs automated tests measuring API quality across industry standards and best practices

API Analytics

Provides Google Analytics style insights into your API, its usage and users

API Security

Scans every request for over 20 API-specific security threats

API Documentation

Automatically generates/updates documentation for every endpoint and creates a developer portal with OpenAPI Specification support

A^(P)I Assistant

Provides an AI integration assistant that can generate integration code, tests, or SDKs in any language based on the API documentation



Customers can choose to deploy Trebble using one of the following options:

01

• SaaS

the simplest and fastest way of using Trebble without worrying about infrastructure and scalability

02

• Private cloud

allows enterprise customers to run Trebble and store data within their private cloud on AWS

03

• On-prem

enables enterprise customers to run Trebble and have complete control over data storage and retention on their off-cloud infrastructure.

As part of its SaaS offering, in 2023, Trebble has processed over **5 billion API requests across 9000 different APIs**. We've spent the past two months analyzing anonymized data from a large subset of these APIs to try and answer many long-standing and open-ended API-related questions.



Key findings ---

AI APIs doubled in 2023.

96%

The number of AI-related APIs grew by a staggering 96% compared to 2022., outpacing APIs within most other industries combined.

Most API requests are made mid-week

WED

Wednesday is single-handedly the most popular day in the week when it comes to request volume all year long across all APIs.

Every fifth endpoint on an API is a zombie endpoint

20%

20% of endpoints on an average API don't get a single request for 30 days or more, making them dead weight on the API.

APIs are still very insecure


51%

51% of all requests did not use any form of authentication, and 55% had a medium threat level score.

Most problems with APIs happen on the client-side


4x

Based on HTTP response codes, client-related errors occur 4x more often than server-related errors.




Methodology ---

We used a quantitative methodology to maintain objectivity and obtain the most accurate results in this report. A quantitative methodology is a way of studying things by collecting and analyzing numbers. It uses measurements and statistics to look for patterns and draw conclusions. This method is suitable for being precise and objective, and it helps researchers understand and make predictions about different topics.



This data was anonymized, and no private, secure, or sensitive data was ever included. We used numbers like response size, load times, and how often people used the APIs to understand how well they were working. **Numbers don't play favorites.** They give us precise information to help us make intelligent decisions. Numbers give us accurate measurements, so we know exactly what's going on without confusion. This enables us to make decisions based on factual evidence, not opinions.

In contrast, the qualitative method dives into the “why” behind the numbers. While the quantitative method provides solid evidence for decision-making, the qualitative method adds depth by uncovering stories, motivations, and contextual details that numbers alone may not reveal.



By looking at APIs through the lens of numbers, we're gaining a complete picture of how they work. This way, we can stay on top of trends, make improvements, and confidently navigate the ever-changing world of digital connections.



API Design

Before diving into other vital indicators like performance and security, it's essential to understand how people build their APIs, and for most REST-based APIs, that starts with endpoints. Generally, the more endpoints an API has, the more operations it can do. However, that also comes with a cost - complexity. It's harder to maintain and update an API with a large amount of endpoints than it is a small and focused microservice. Yet again, orchestrating a network of microservices isn't a walk in the park.

Our data shows that an average API has 22 endpoints. The most extensive API we observed had **319 endpoints on a single API**. We also divided endpoints into four different groups based on size.

When we look at the distribution among those groups, it looks like this:

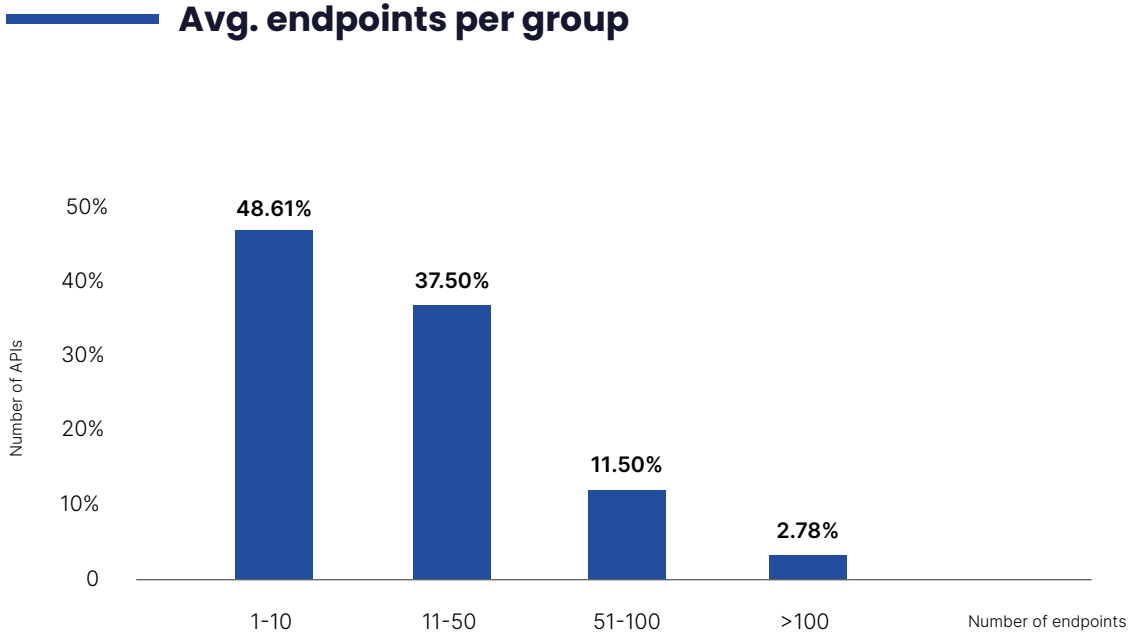


Figure 1. Avg. endpoints per group, source: Trebble; sample data 2023

As you can see, most APIs have 1-10 endpoints, indicating that many companies choose a microservice-oriented architecture over a monolithic approach when building APIs.

Let's go a level deeper and analyze the type of endpoints that an average API uses. HTTP request methods determine endpoint types. Each method has a specifically intended use case, such as storing and getting data.

There are 9 of them: GET, POST, PUT, PATCH, DELETE, OPTIONS, HEAD, TRACE and CONNECT. Looking at their usage across requests, we get the following distribution:

HTTP method usage

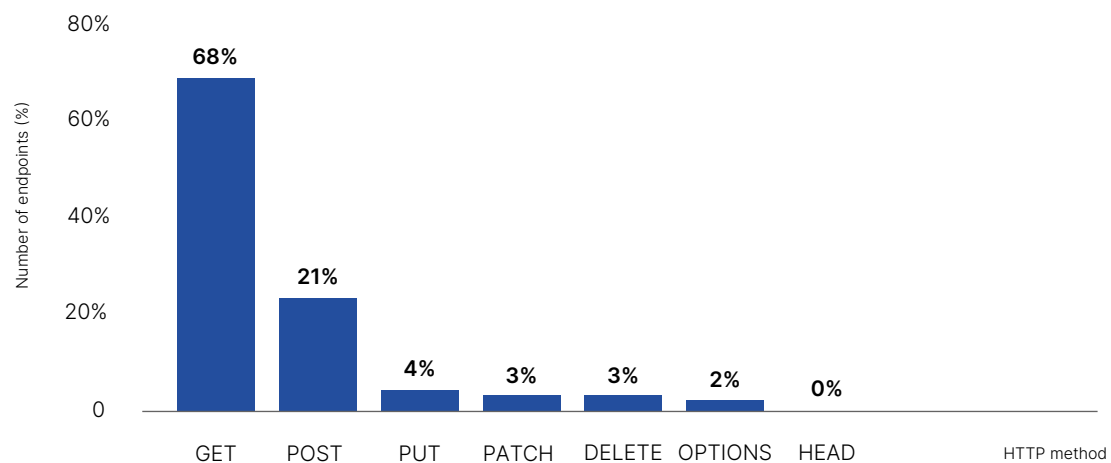


Figure 2. HTTP Method usage, source: Trebble; sample data 2023

Virtually only two methods are predominantly used. **68% of all endpoints are GET, and 21% are POST.** All other endpoints account for less than 3% each. These results mean two things:

- 01** • HTTP methods are underutilized, and engineers generally use the POST method to do the work of PUT, PATCH, and DELETE operations.
- 02** • Most endpoints (and most APIs) are developed primarily for consuming information, as that's the sole purpose of the GET HTTP method.



Shailendra Bade
Engineering Director at
American Express

“API usage is booming - it's not a trend, it's a business essential. Manual management? Not viable. For sustainable growth, observability and governance are key. Automated tools are no longer a choice, but a necessity. Without these, failure isn't just a cliché, it's a reality.”

While discussing endpoints, we should differentiate between active and zombie endpoints. The notion of zombie APIs has existed for a while now, but based on the data we see, we believe that zombie endpoints pose a more significant threat.

Our research shows that 23.5% of all endpoints are zombie endpoints on an average API with 11-50 endpoints or more.

— Zombie vs. Active endpoints




Figure 3. Zombie vs. Active endpoints, source: Trebble; sample data 2023

That means **every fifth endpoint is a zombie endpoint** and has not received a single request in the last 30 days. The danger of zombie endpoints isn't that users are not using them but rather that they are rarely updated and maintained from an engineering perspective. This poses various security threats that often lead to data leaks and attacks.



Bill Doerrfeld
Editor in Chief at Nordic APIs

“As I said in a recent talk, APIs are more relevant than ever, and this report provides some data-backed context into trends we are sensing in the industry concerning design, performance, and security. One thing is it confirmed what I had suspected, that a significant number of endpoints are sitting inactive. Also, hump day is real!”

Next, we looked at HTTP response codes to understand how the relationship between clients and servers is evolving. These codes indicate the result of a client’s request to the server. They are standardized, universally understood, and come included with every HTTP server. There are 63 different HTTP response codes grouped into five groups:

- 1xx range: Informational responses
- 2xx range: Success responses
- 3xx range: Redirection messages
- 4xx range: Client error responses
- 5xx range: Server error responses

When we look at the data on how these response codes are used, we have the following distribution across groups:

Response code group distribution

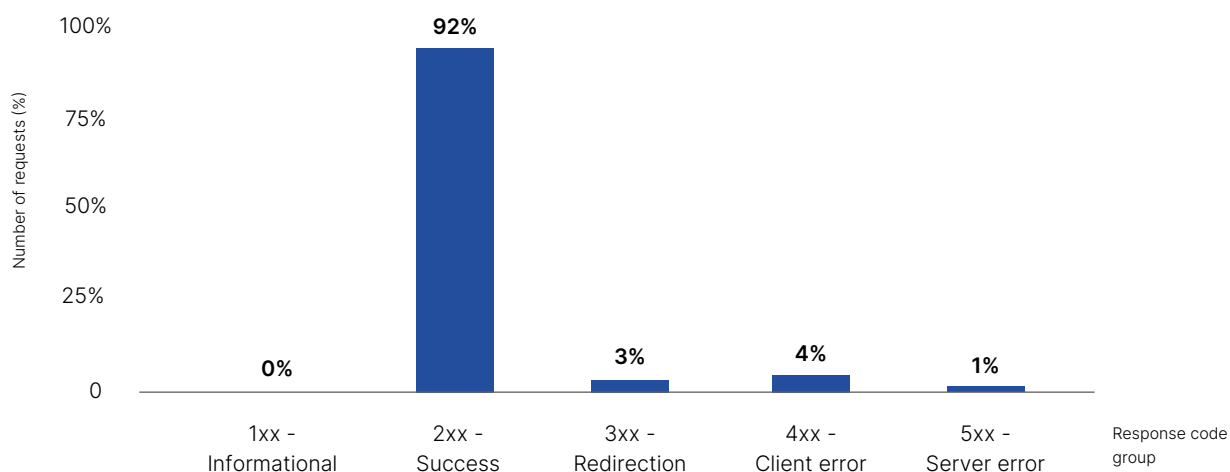


Figure 4. Response code group distribution, source: Treble; sample data 2023

Given that the values of almost all of them are so low, it might make sense to repeat these as raw numbers:

- 1xx Informational: 0%
- 2xx Success: 92%
- 3xx Redirection: 2.62%
- 4xx Client Error: 4.45%
- 5xx Server Error: 0.92%

As evident, **92% of all requests return a response code in the 2xx range**. If we dive deeper into the top codes from that range, we get the following breakdown:

- 200 (OK): 87.27%
- 204 (No-Content): 2.08%
- 201 (Created): 1.83%
- 206 (Partial Content): 0.50%

This outcome is expected, as 200 OK is the most popular HTTP response code of them all.

However, the rest of the codes tell a rather interesting story. Data shows that client errors (4xx range) are 4x more common than server errors (5xx range). Translated to human language, this means that even when problems do occur on APIs, they mainly occur because of problems on the client-side. If we drill down on the top 4xx response code, we get these numbers:

- 401 (Unauthorized): 2.14%
- 404 (Not Found): 1.18%
- 429 (Too Many Requests): 0.80%
- 403 (Forbidden): 0.50%

When clients do make requests to API, they tend to forget to authorize or request resources that no longer exist. Usually, unauthorized requests happen more in the development phase, and missing resources are a production problem.

The final piece of the API design puzzle is the technology engineers use to develop these APIs. We can understand this data by looking at the SDKs Trebble users use. To get started with Trebble, engineering teams need to add the Trebble SDKs or agents that send the data they collect to Trebble. All SDKs are open source on GitHub and take less than 3 minutes to integrate. More importantly, Trebble can easily differentiate between them. Looking at their popularity, we have the following list of TOP 10 used SDKs:

Most popular Trebble SDKs

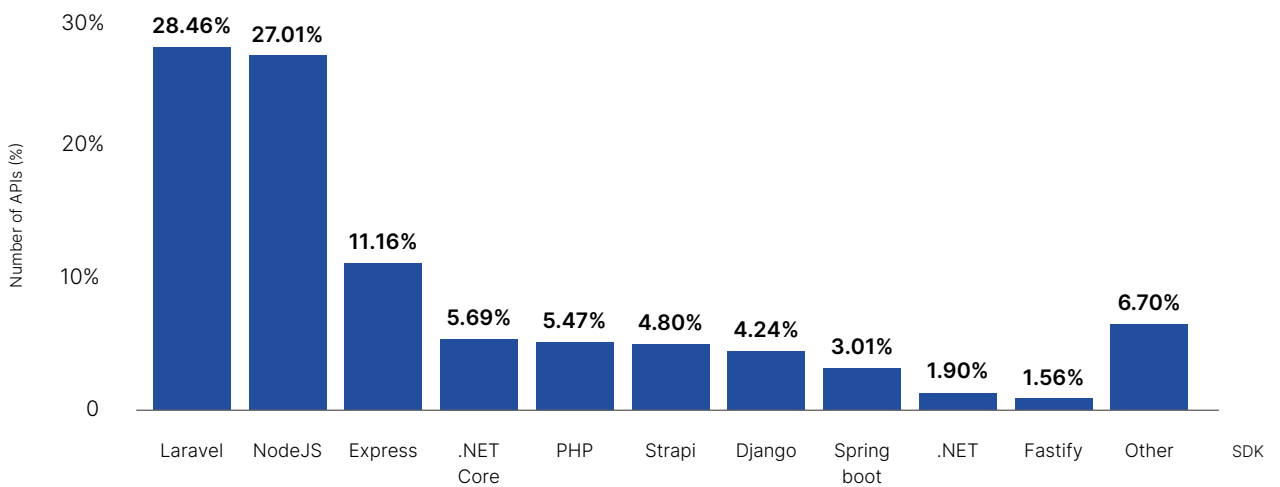


Figure 5. Most popular Trebble SDKs, source: Trebble; sample data 2023

Laravel (framework in PHP) and NodeJS are the two most popular choices for building APIs. This might seem surprising to some, but a recent report from W3Techs⁽⁴⁾ shows that PHP is by far the most used server-side language, with 76.6% of the market share. If we grouped the SDKs based on language, we quickly get to three core technologies:

- Javascript
- PHP
- .NET.

Javascript and PHP were at the top of the list last year, but we are seeing a significant increase in the use of Microsoft technologies like .NET. That shouldn't be surprising because Microsoft has invested heavily in its developer platform, making adoption easier and connecting many other platforms like Azure, Github, and AI tooling into one cohesive strategy.

⁽⁴⁾ "Usage statistics of server-side programming languages for websites." https://w3techs.com/technologies/overview/programming_language. Accessed 30 Nov. 2023.



API Performance

Avg. end points load times




Figure 6. Avg. endpoint load times, source: Trebble; sample data 2023

The performance of APIs can only be described as a story of extremes, but overall, the performance has improved based on the data we have from 2022. On one side, **54.7% of endpoints had an average load time of 150 ms or below**, which is good. On the other spectrum, 26.8% of endpoints had an average load time of 500 ms, which is bad. The mid-tier combined makes 18.5% of endpoints with a load time from 150 ms to 500 ms, which isn't bad but isn't perfect either.

Knowing that most of the endpoints on an average API are GET endpoints, we wanted to understand how different operations impact load times. It turns out that **GET endpoints are 2x faster than POST endpoints**. This somewhat makes sense because storing and updating data on the server requires more database operations and checks, which slow down responses from the server. Also, GET requests can be cached from the server side, which API Gateways often do by default.

When we think about API performance, it's also essential to understand the number of code-related errors that occur on the API. Because Trebble SDKs sit on top of the API code base, they can capture exceptions that happen at runtime. When we examine all the requests' data, we can see that **code-based errors occur 4% of the time**. This might seem small, but if your API gets 5M requests per month, 200K of them had a code-based error on them, and the API didn't do what it was supposed to do. That leads to customer frustration, loss of revenue, and trust.




Figure 7. Request error ratio, source: Trebble; sample data 2023

Another data point we can look at is the Trebble API Score. Trebble actively measures the API quality across three categories on every single request. Those three categories are security, performance, and overall quality. Performance and quality include checks like load time, response sizes, latency, the ratio of code-based errors, caching, and similar checks. The maximum score an API can get is 100, and the **average score across all APIs on Trebble is 50**. That means there is a lot of room for improvement across all three categories, especially performance and quality. Diving deeper into the score data, we wanted to understand how the API score differs from language to language.

Avg. API score per SDK

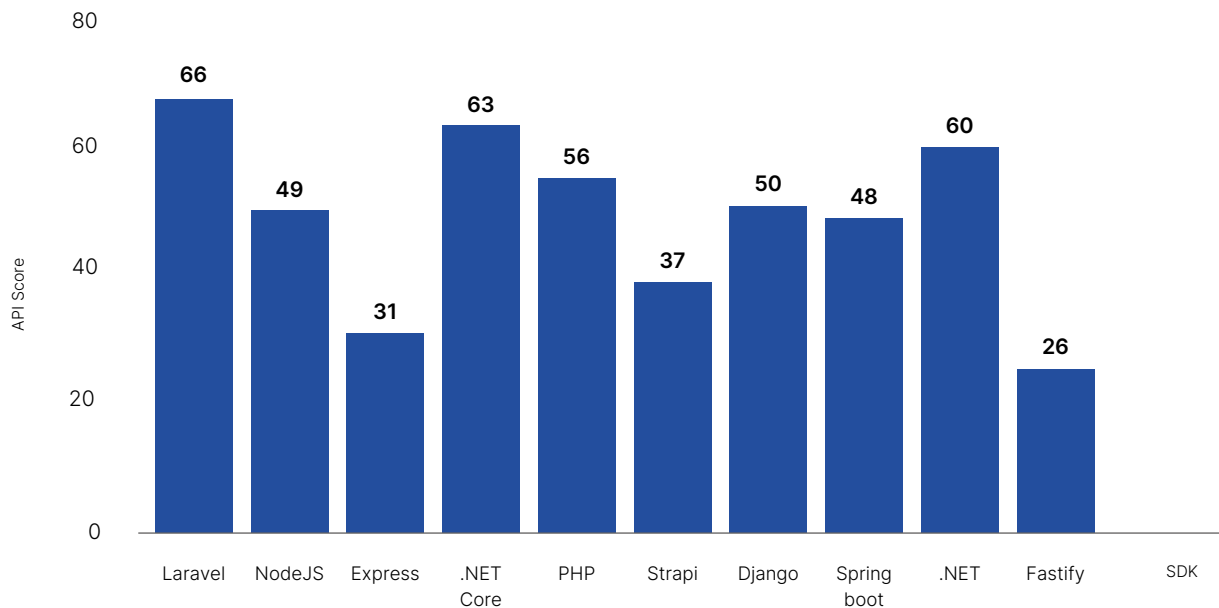


Figure 8. Avg. API score per SDK, source: Trebble; sample data 2023

The data above shows that PHP and .NET-based APIs have an overall higher API score than other languages and frameworks. Laravel is leading the way with an average score of 66, followed by .NET Core with 63 and .NET with 60. This makes sense because building great APIs requires skill and practice. These two language groups have been around since the dawn of the internet. They are primarily server-side orientated, and in most cases, frameworks like Laravel and .NET Core impose a specific set of design and architectural best practices that push the quality forward.



Gulshan Jubaed Prince

Product Manager at Annex Cloud

“In the realm of software engineering, crafting exceptional APIs is a daily mission. The backbone of this mission? Building high-performing and high-quality APIs. Freeing engineers from mundane work and improving developer experience is the need of the hour across industries and engineering teams. It’s this synergy between meticulous monitoring and engineering finesse that will define our industry — where innovation meets reliability in every API built.”



API Security ---

Security is, by its nature, a pretty sensitive topic. Especially on APIs because, in most cases, security starts with design. If you want a secure API that doesn't expose data and checks many of the best practices, you have to design and build it that way. There's no way around that. No magic button, framework, or AI that can help you with that - for now. We've also observed that people tend to overcomplicate security and forget to do the simple things that matter the most.

One of those simple things is authentication. Our data shows that **51% of all requests don't have any form of authentication**. API authentication is the most basic form of API security, where each client gets a unique key that identifies them when making requests. That allows API owners to control how, when, and in what capacity they can access the API. Not using authentication severely degrades API security at a fundamental level.



**Manjunath Iyer
Basaralu Srinivasa**

Engineering Leader at Atlassian

"APIs, despite their long-standing presence, exhibit varying maturity levels. Even within prominent organizations - standards, security, and quality can fall short. To propel our collective progress, we need tools empowering developers worldwide to stay vigilant about standards. My emphasis lies on enhancing API security and documentation standards — a crucial stride for a more robust and reliable ecosystem from day one."

Authenticated vs. Unauthenticated requests




Figure 9. Authenticated vs. Unauthenticated requests, source: Treble; sample data 2023

Another security data point we can analyze is the use of Hypertext Transfer Protocol Secure or HTTPS. It uses TLS (Transport Layer Security) to encrypt data sent from the client-side to the server-side and prevents man-in-the-middle attacks. It's easy to use and, in most cases, free as part of default offerings by major API Gateway players and DNS providers. HTTPS usage is better than Authentication but is still not perfect.

HTTP vs. HTTPS

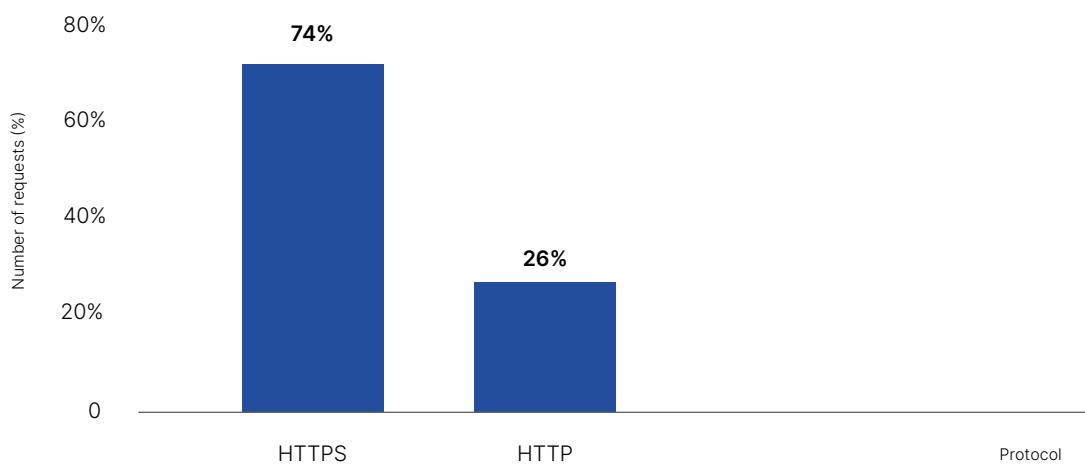



Figure 10. HTTP vs. HTTPS, source: Treble; sample data 2023

74% of all requests were made over HTTPS, while 26% used the HTTP protocol. There is no objective reason why this shouldn't be a much higher number in favor of HTTPS as, as we've mentioned, it's free, has virtually no performance impact, and has clear and immediate benefits.

Another metric that Treble considers when looking at security is the threat level score. Treble runs more than 15 automated checks on every request - specific to security. These checks include, but are not limited to: SQL injection tests in the request payload, authorization usage, HTTPS usage, IP reputation, and many security design best practices. For every request, it gives a **threat level score of either low, medium, or high**, depending on the importance of failed checks. The distribution of those across all requests looks like this:



Security threat level breakdown




Figure 11. Security threat level breakdown, source: Trebble; sample data 2023

55% of requests have a medium threat level score, followed by a low threat level, which accounts for 44% of requests. The medium threat level score mainly implies design-level security issues like not using authentication, using IDs over UUIDs, and exposing a lot of security headers. A high threat level score, which affects 1% of requests, implies serious security threats like SQL injection attacks and other hacking attempts for known bad actors.



API Market

In every market, it is important to understand the players, who they are, what they do, and where they come from. The API market is different, so we wanted to understand which industries are actually building the APIs we use and love.




Figure 12. Top 10 industries, source: Trebble; sample data 2023

If we break down the graph and extract the top 3 industries, we can see the following numbers:

- Information technology & services: 45.17%
- Higher education: 5.84%
- Financial services: 5.13%

It's not a surprise that the IT industry is a leader in building APIs because, in most cases, this includes all tech companies, products, and services we use and love daily. Higher education is a surprise but understandable given that a lot of transformation is happening in the ed tech industry. Similar to financial services, this industry is probably the longest-standing industry when it comes to APIs and probably the one that APIs have impacted the most.



Sanjay Jain

Chief Technology and Product Officer at Freecharge and Digital Business And Transformation, AXIS Bank

“As far as the Financial Services sector is concerned, you can’t think of driving your business without APIs. APIs are the backbone of tapping more partnership opportunities and creating new channels for increasing your user base. More broadly, we are in the digital transformation era, with most organizations going through some form of large-scale shift with the technologies they leverage. APIs are at the center of these changes. Adopting becomes a real challenge if these APIs are not architected well. There has to be a dedicated team and observability platform in place to look after the performance of APIs.”

Given the popularity of AI nowadays, we wanted to understand how that translates to APIs. To determine that, we looked for APIs from the AI industry and those that use a .ai domain name. We found a 2x increase in AI-related APIs from 2022 to 2023.

Growth of AI related APIs




Figure 13. Growth of AI-related APIs, source: Trebble; sample data 2022, 2023

We also wanted to understand if there are any differences between the quality of the APIs across industries. We correlated the Treble API Scores and all the industries to get the following breakdown:

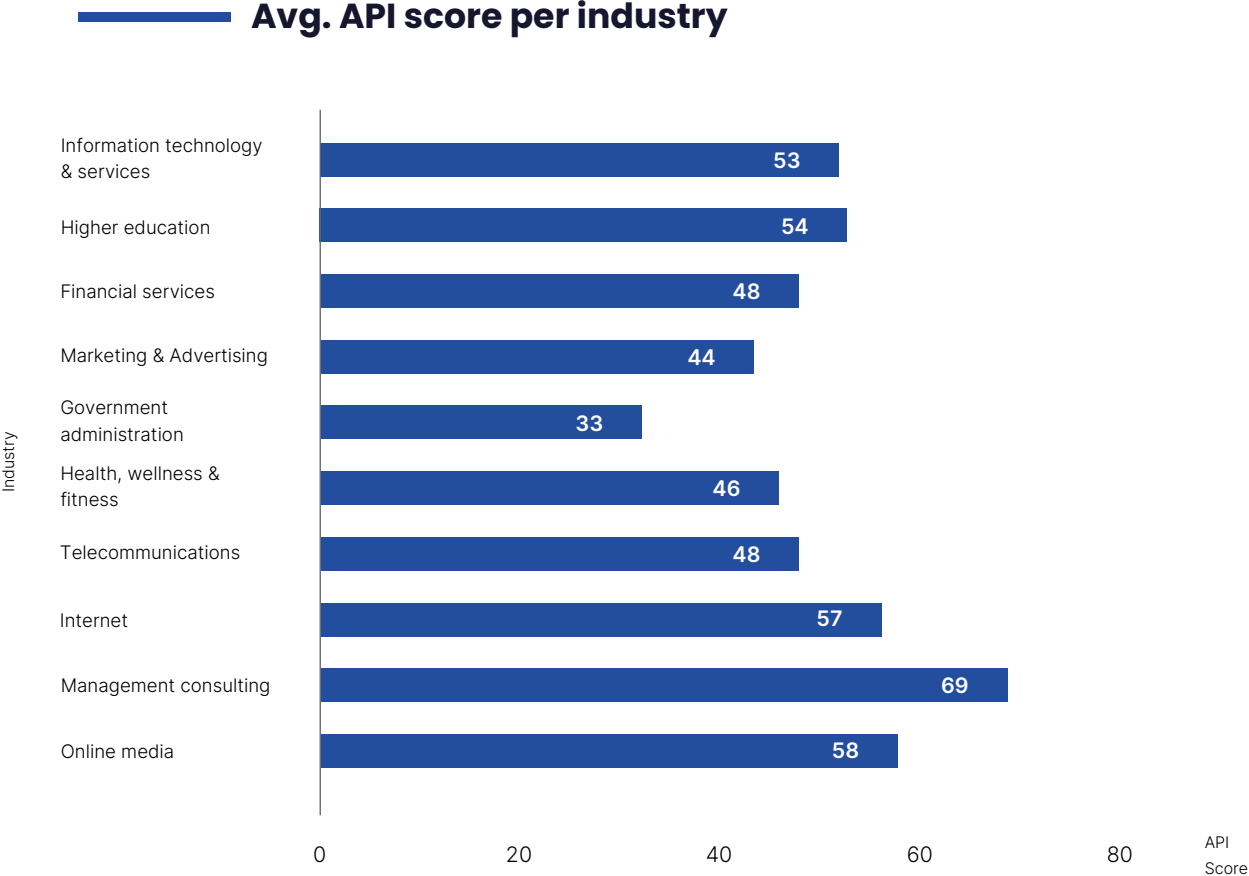


Figure 14. Avg. API score per industry, source: Treble; sample data 2023

In this case, the clear winner is Management consulting, with an **average API Score of 69**. Other contenders include:

- Online media: 58
- Internet: 57
- Higher education: 54

Almost all other industries have a similar API Score ranging from mid-forties to mid-fifties. The lowest score of 33 belongs to Government administration. This makes sense as the government sector is usually slow to adopt new technologies and is just entering the digital transformation age.



Bruno Pedro

author, "Building an API Product."

"It's clear to me that APIs aren't exclusive to the technology sector. Recent surveys show developers represent less than half of all API consumers. This report corroborates the same trend happening with API producers. Banking, education, and healthcare employ an estimated 8% of the global labor. That's a total addressable market of about 275 million, mostly untapped. Compare it to the overcrowded 70 million technology TAM, and it's evident where you should invest next."

Next, let's look at where most requests originate from, grouped by country:

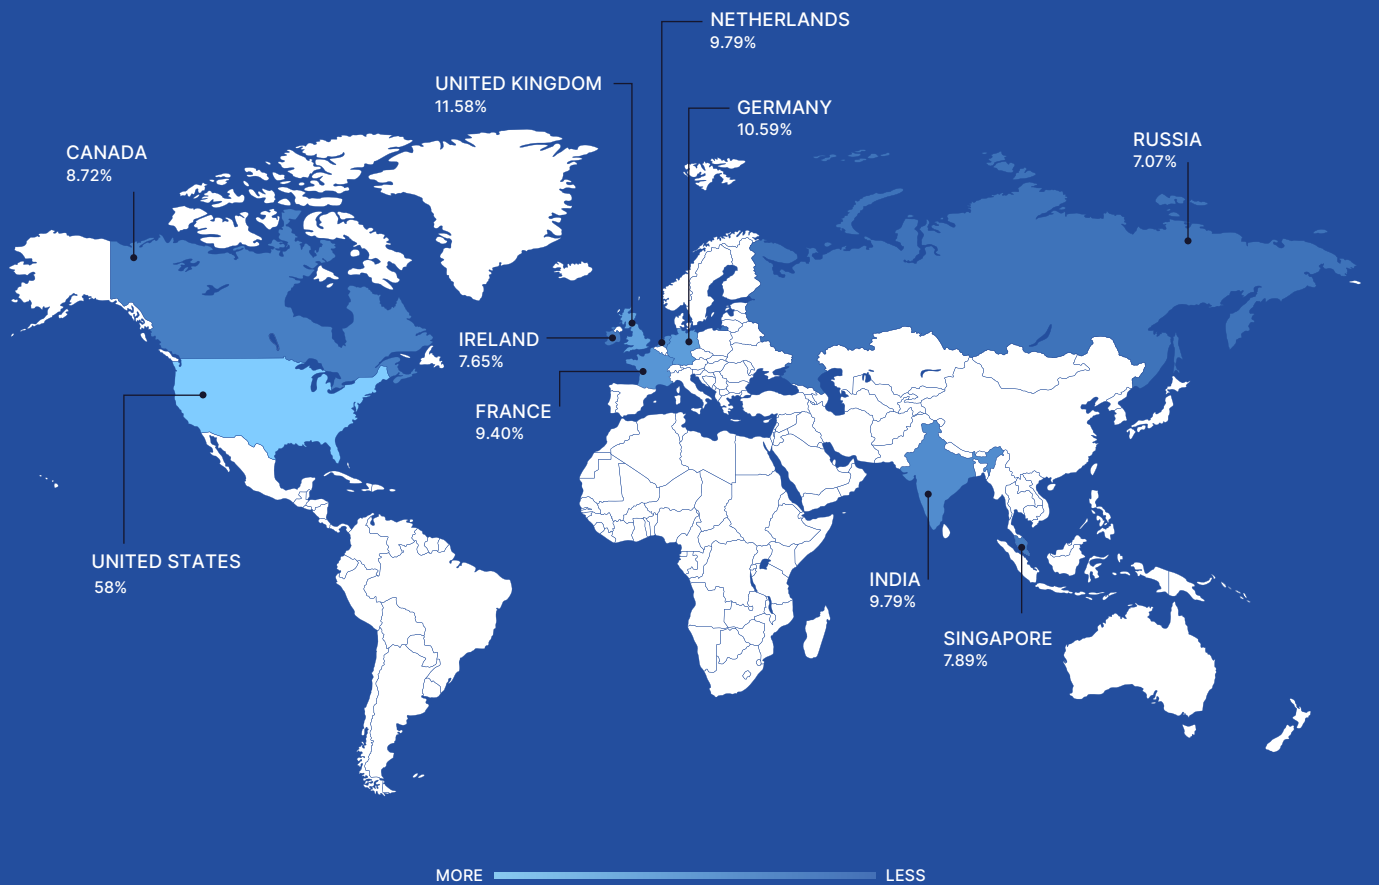


Figure 15. Top 10 countries based on request volume, source: Trebble; sample data 2023

A clear leader by almost double is the United States, followed by the United Kingdom, Germany, Netherlands, and India.

Going a step further and breaking that down to a more granular level of cities, we see the following:

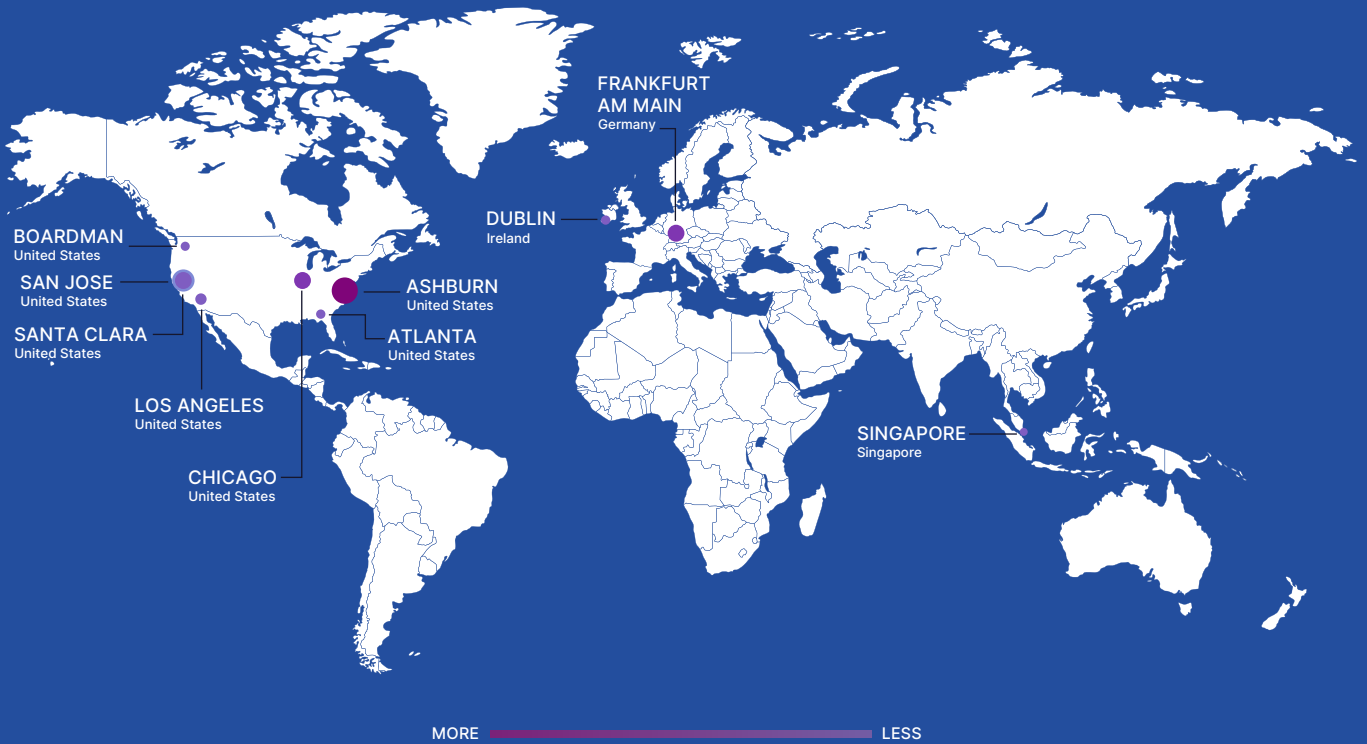


Figure 16. Top 10 cities based on request volume, source: Treble; sample data 2023

Based on this, the capital of APIs is Ashburn, Virginia. You might wonder why a city on the east coast of the United States with a population of 44K makes the most API requests globally. It's simple: Ashburn is the hometown of Amazon's AWS data center for the US East Coast. **This means that most API requests are not made directly through end clients but from various server-side back-ends.** This perfectly matches the microservices trend where many requests hop between different APIs from the client-side to the end location.

We can confirm that by looking at another data set about end-user devices. Treble can automatically differentiate between devices like Desktops, iOS, and Android. Based on that, we get the following breakdown:

Client distribution per device




Figure 17. Client distribution per device, source: Treble; sample data 2023

54.8% of requests originate from desktop-based devices. This includes API requests between different microservices or from one server instance to another. And as we know, there's a lot of them. The battle for second place is a close call, but Android wins over iOS with a 0.3% margin. Apple has a considerable market share in some countries like the USA, but Android has a more significant global market presence.

The final two data points are about time. We've looked at data that helps us understand when API requests are made. The first data set shows us which day of the week is the most popular for API requests all year long.

Request volume per day of the week




Figure 18. Request volume per day of the week, source: Trebble; sample data 2023

The weekly breakdown shows that mid-week is the most popular time for API requests. The winner is Wednesday, with 15.44% of all requests happening that day. The next most popular day is Thursday, and then Tuesday. The lowest traffic day for APIs is Sunday by far. Since the number of requests tends to decrease over the weekend, it would imply that APIs are heavily used in day-to-day work and are still quite business-focused.

Request volume per quarter




Figure 19. Request volume per quarter, source: Trebble; sample data 2023

By extracting the data from 2022, we can compare not just trends but also the growth of APIs. So far, the biggest quarter is Q3 of both years, as well as the quarter where we see the most significant jump in volume. One important caveat is that at the time of publishing this report, there are almost 32 more days until Q4 2023 ends. Yet Q4 of 2023 is already bigger volume-wise than Q4 of 2022.

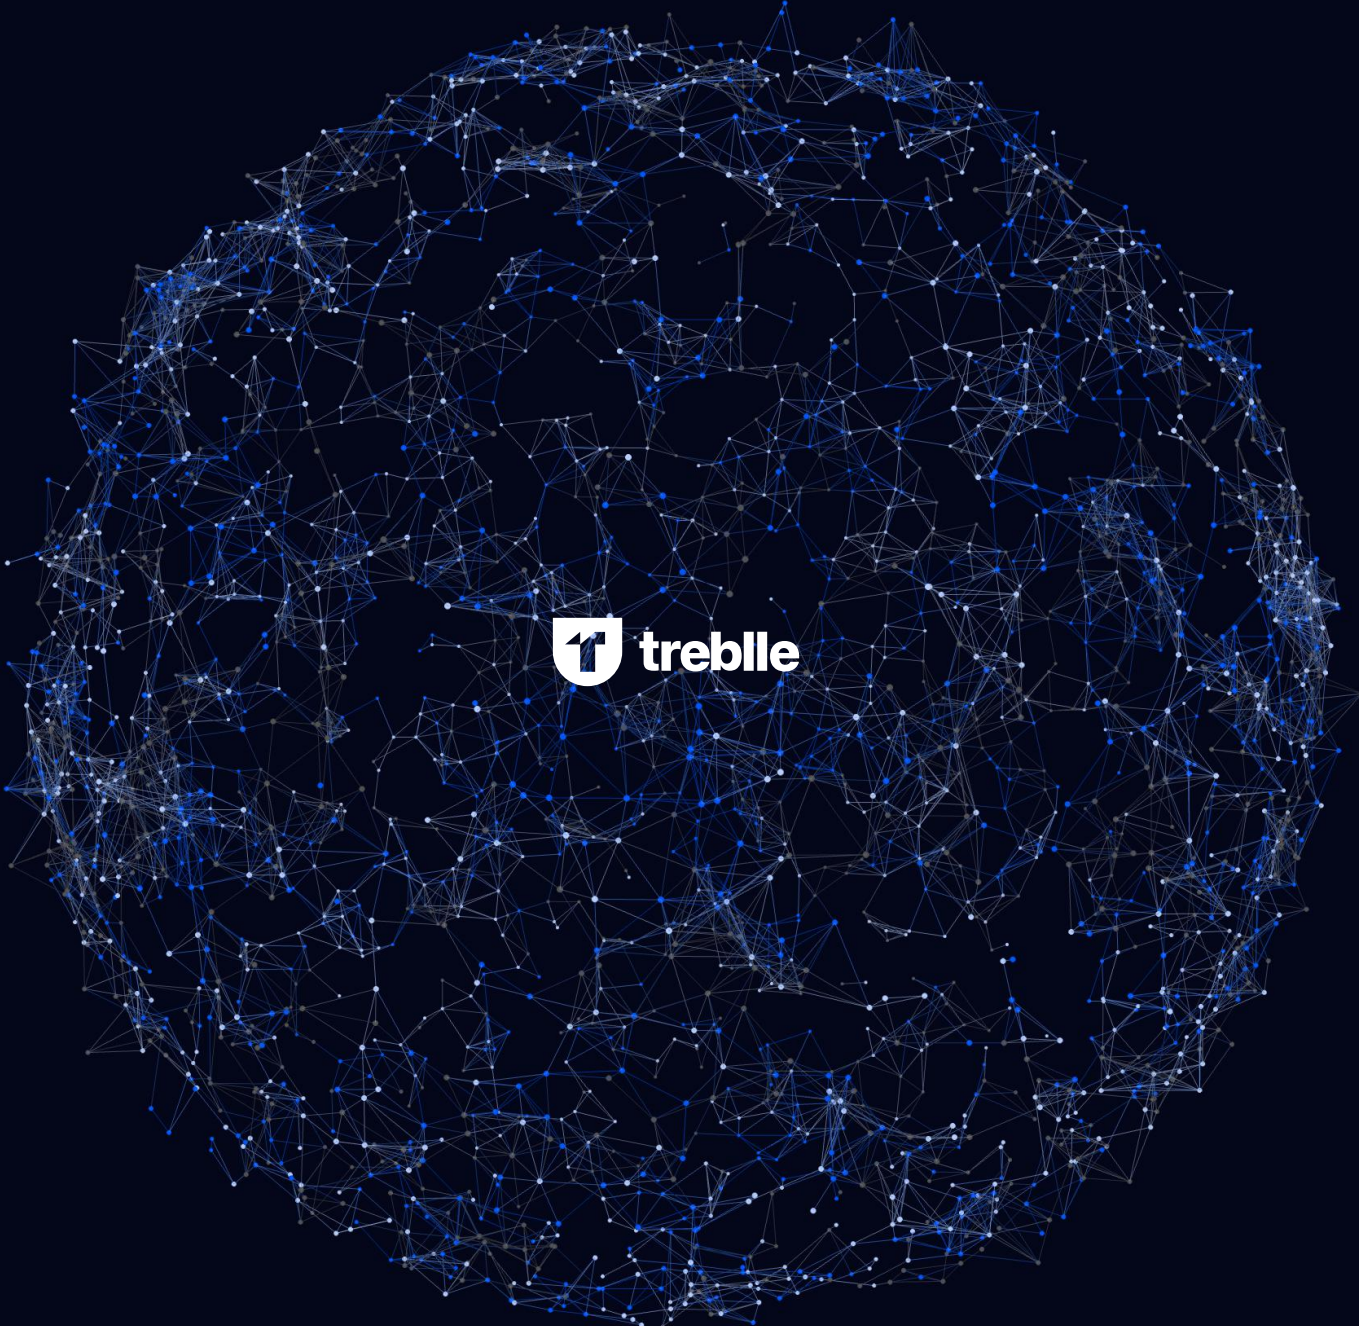
Conclusion

APIs are growing fast. Not just in the sheer volume of requests but also in the number of new APIs spawning up, new industries joining the API economy, and the amount of shareholder value generated. In November of 2023, Stripe, the API payments platform, processed **\$18.6B in a single weekend**.⁽⁵⁾ This puts APIs at the front and center of big business, especially when you factor in the high ROI that APIs offer. It's far cheaper and more efficient to build an API, connect it to multiple clients, or even better yet have others build on top of it. Not only that, but you're also de-risking your tech stack and making it more future-proof. So far, clients have changed in many forms and shapes throughout the years, from computers to mobile phones, watches, and wearables. On the other hand, one thing has largely stayed the same: how data is exchanged - through an API.

Besides growing fast, APIs are becoming increasingly complex, from moving money, controlling supercolliders, and testing F1 cars to facilitating every AI interaction. The stakes for APIs have never been higher as they handle more complex and challenging tasks.

Going forward, the fundamental challenge for businesses will be understanding and democratizing access to API data. Only when you have those two can you enable your team to build better and more secure APIs that bring in more customers and business. These are complex problems to solve and require adequate tooling, just like anything else. With all that in mind, we expect API Observability and Governance tooling to play a significant role in the future of APIs. That's exactly what we at Treble are building: an end-to-end APIOps platform that helps organizations build and ship quality APIs faster.

⁽⁵⁾ "Stripe processed record \$18.6bn over Black Friday weekend." <https://www.siliconrepublic.com/business/stripe-black-friday-weekend-sales-payments-volume-tracker>. Accessed 1 Dec. 2023.



 **treble**