elastic

Real-world problems require real-time data:

A strategic guide to putting your data to work with search and AI

How IT leaders can improve digital customer experiences, increase operational resilience, and reduce cyber risk by putting existing, untapped data to work in real time with search and AI

Table of contents

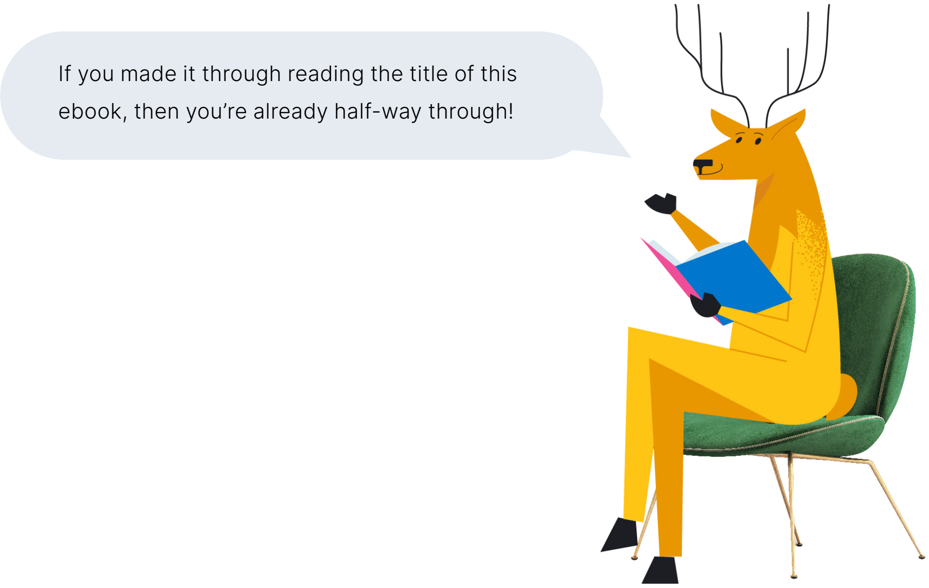
If you made it through reading the title of this ebook, then you’re already half-way through!

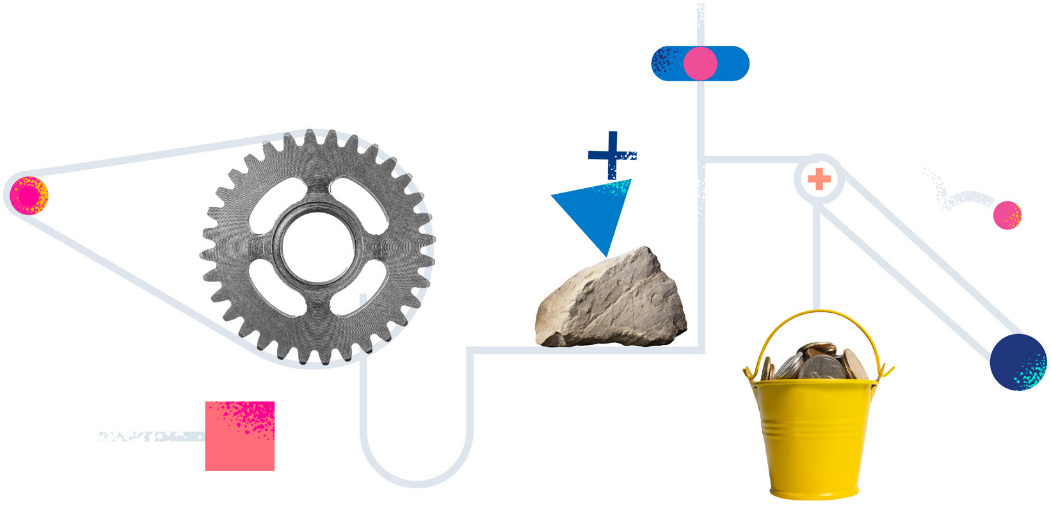

Introduction

As a consumer and an employee, you expect nothing short of seamless, secure experiences across all of the applications, websites, emails, texts, and video calls that you interact with every day. As an IT leader, you're expected to keep all of these underlying systems running and secure to make your customers happy, keep your employees empowered, and meet your business goals.

And that's all coupled with macroeconomic conditions that add pressure to find cost savings without compromising your IT performance and customer experiences. That's no easy task.

With the ongoing need to increase visibility into the performance of critical applications and infrastructure, dial-up cybersecurity, and improve the ability to surface relevant information, what if we told you there's an opportunity to condense your tech stack and save along the way?

Organizations that use real-time data for the right purpose are:[1]

- 8x more likely to grow revenue by 20% or more
- 1.4x more likely to uncover new revenue streams
- 1.6x more likely to create data-driven experiences
- 1.8x more likely to commercialize their data

[1] Forrester, The state of insights-driven business, 2022.

Section 1:
Your business problems are data problems

As an IT leader, you are expected **to optimize enterprise applications and infrastructure for availability and performance.** The average application, comprising of 50 to 100 services with multiple deployments, can generate more than 300 GB of data per hour during an incident or outage.[2] And when IT downtime can cut enterprise profit by 9%,[3] every second of an outage matters. Imagine if a grocery store's online ordering system went down days before a big holiday. If the stress of the holiday season wasn't enough, this would certainly make for some unhappy customers. Unplanned downtime not only leaves a bad taste in customers' mouths (no pun intended) but can easily amount to millions of dollars lost. [4]

These are data problems.

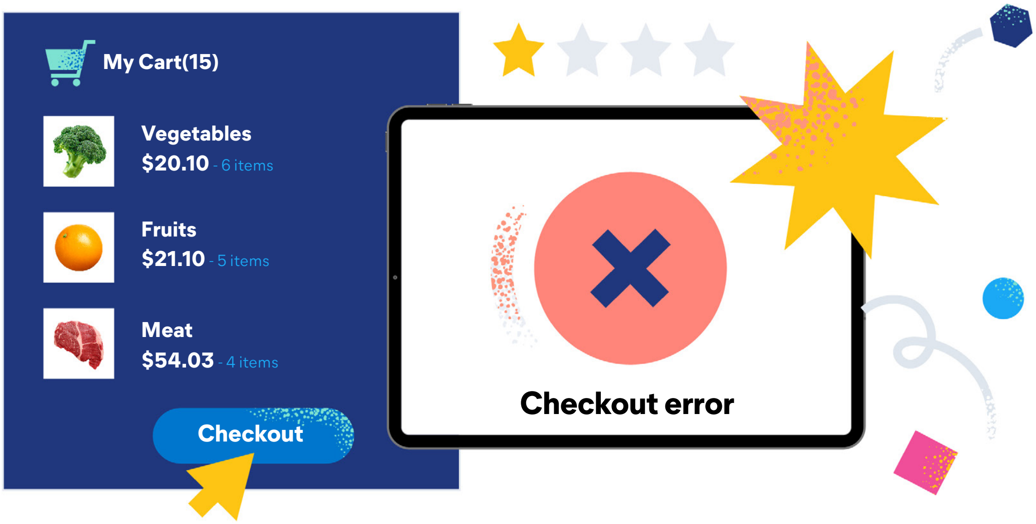

[2] IBM, AIOps: Unleashing the hidden insights in unstructured IT data for better IT operations management, 2021.
[3] CIO.com, IT downtime cuts enterprise profit by 9%, says study, 2024.
[4] IBM, Cost of a Data Breach Report 2023, 2023.

As an IT leader, you're also expected to prevent security threats and detect and resolve incidents quickly when they do occur. The average Fortune 500 enterprise generates more than 10TB of security events per month.[5] You and your team are always sifting through the exponential avalanches of security data, worried about the next threat and the impact it could have on your brand and the disruption it would have on your organization's day-to-day operations.

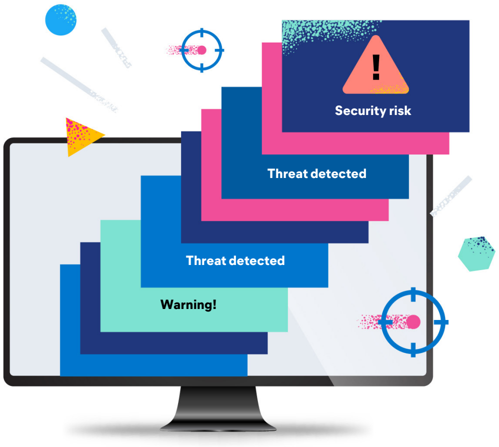

These are data problems.

IT leaders are expected to connect the right people and teams with the right information, at the right time regardless of where the information is, or the format of the data. For example, a customer needs help installing their video doorbell. By interacting with a self-service experience, the customer gets frustrated that they have to enter in their product number to even start the troubleshooting process. The experience is cumbersome and clunky. About 77% of consumers say that offering poor self-service support is worse than not offering any at all since it wastes time.[6]

These are data problems.

[5] AT&T, What to log in a SIEM: SIEM and security logging best practices explained, 2020.
[6] HigherLogic, An Exploratory Research Study: Customer Experience and Customer Self-Support, 2020.

All of these challenges are fundamentally connected to data. In fact, the average enterprise stores more than 71PB of structured and unstructured data on-premises alone, not even including cloud.[7] A mountain of data that could be hiding key insights to business growth, indicators of compromise, and everything in between. This data shouldn't just be stored, but rather put to work.

Top data challenges and opportunities for C-Levels[8]

69% Data quality issues

69% Organizing unstructured data and making use of it

69% Poor real-time data accessibility

67% Inability to efficiently use AI to analyze data

67% Inability to monitor data and use insights in real time

Unfortunately, only 40% of data within organizations is actively being put to work[9] today, which leaves an undesirable amount of data taking up space and costing money to store without adding any value. Using this untapped data will enable you to keep customers happy, keep your systems up and running, and keep your organization protected. But how?

[7] IDC, Meeting the New Unstructured Storage Requirements for Digitally Transforming Enterprises, 2022.
[8] Elastic, Solving business challenges with data and AI: 5 insights from C-suite leaders, 2024.
[9] BARC, How Much Information Available to Companies Is Used for Decision-Making?, 2024.

> **To address these data challenges, organizations need a way to derive value from data continuously, in real time.**

So, how do you do that?

With **the precision of search and the intelligence of AI**, you can enable users to find the answers that matter from all of your data in real time, at scale. Search technology by design is built to surface the most relevant pieces of information at speed and scale, but search technology is only capable of generating lists of results. On the other hand, generative AI can generate exact answers using its computational power and intelligence, but it has no context of your organization's knowledge.

When you combine the precision of search with the intelligence of AI and your proprietary data, you can find answers to enterprise problems from all data, in real-time, and at scale.

Search combined with the intelligence of AI is able to quickly ingest messy data at enormous scale and enable all types of real-time analytics on top of that data. Even when data is complex, the type of analysis required isn't known ahead of time, or analytics are needed in real time, the combination of search and AI is able to take on the challenge.

Most business problems are actually data problems

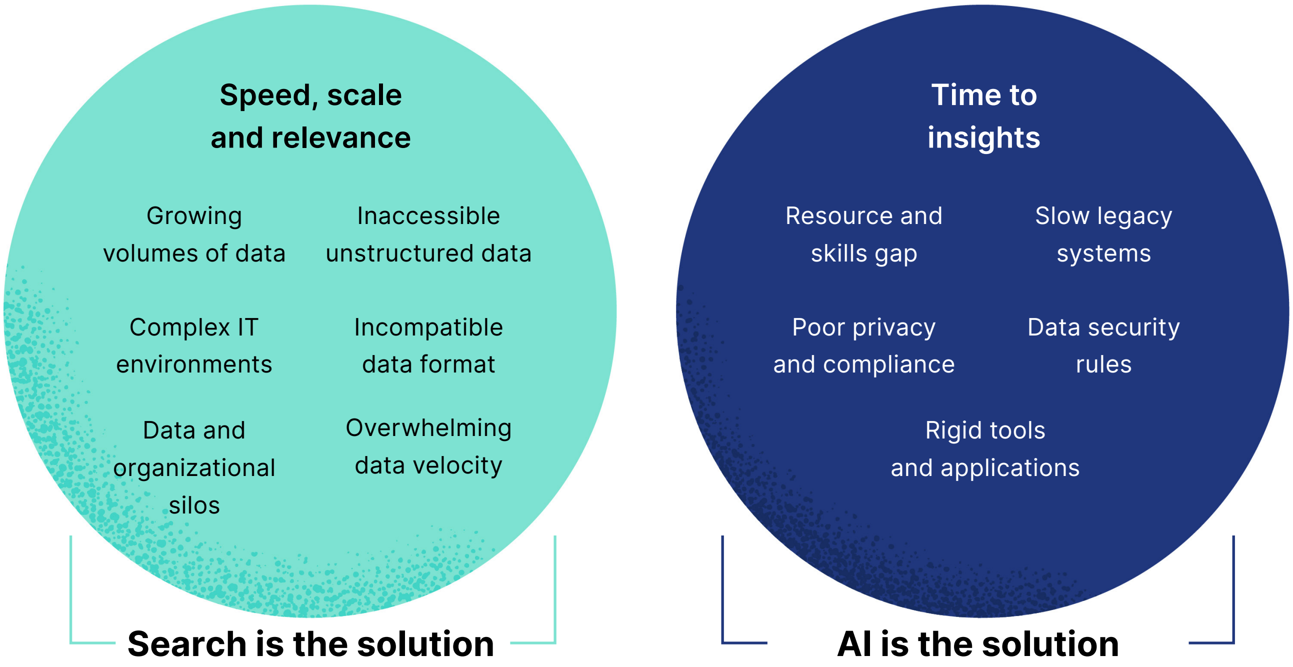

Speed, scale and relevance

Growing volumes of data

Inaccessible unstructured data

Complex IT environments

Incompatible data format

Data and organizational silos

Overwhelming data velocity

Search is the solution

Time to insights

Resource and skills gap

Slow legacy systems

Poor privacy and compliance

Data security rules

Rigid tools and applications

AI is the solution

How do you put search and AI to work? Read on to find out everything you need to know in order to analyze your data, extract insights, and continuously derive value in real time. But first, a bit of insight into the challenges organizations are facing.

Section 2:

Data challenges and business complexities continue to accelerate

Organizations of all sizes are experiencing exponential growth of unstructured data and business complexities are soaring to new heights. You're most likely dealing with at least one of the challenges below.

Customer and employee expectations are higher than ever. In fact, 84% of customers expect[10] and hope for brands to adopt new digital solutions to deliver products and services to them. And when it comes to employees, they need their technology stack to work for them and integrate disparate systems across the organization in order to feel empowered to do their best work. IT teams must monitor all systems to ensure problems are diagnosed and fixed quickly, which leads to a huge amount of data that must be monitored.

Security attacks are on the rise with a 20% increase in attacks from 2022 to 2023,[11] and have been further exacerbated by AI.[12] IT needs to partner closely with security teams to ensure security is unified across the business and built into employee education. Attempted breaches and attacks targeted at individual employees are only expected to increase.

[10] Appnovation, The Digital Consumer Trends Report, 2021.

[11] Harvard Business Review, Why data breaches spiked in 2023, 2024.

[12] FBI, FBI Warns of Increasing Threat of Cyber Criminals Utilizing Artificial Intelligence, 2024

A shift in emphasis to providing business value. The days of IT simply being a cost center are in the past. IT teams are shifting from just delivering technologies to optimizing them.[13] IT is now expected to maximize the organization's technology investments to improve operational efficiencies and, ultimately, generate tangible value. All this to say, IT teams are expected to substantially help generate revenue.

Digital transformation initiatives continue to influence IT strategies as organizations move their processes to more digitized environments and workflows. With the goal of aligning growing data in real time across environments and platforms, organizations have a lot more data to manage, process, and extract insight from to transform their business and meet employee and consumer expectations.

With these complexities, come business goes you need to meet:

1. The need to elevate customer and employee experiences to keep customers satisfied. This relies on seamless information discovery, holistic visibility into IT systems, and strong data protection.
2. The need to improve operational resilience to keep systems reliably up and running, which depends on resilient data storage and comprehensive insights into IT systems and security events data.
3. The need to reduce security risks to keep systems secure and sensitive data protected, which requires a secure data store and comprehensive monitoring of security events and IT systems.

[13] CIO.com, Key IT initiatives reshape the CIO agenda, 2023.

Traditional tools ultimately cause more problems than they solve

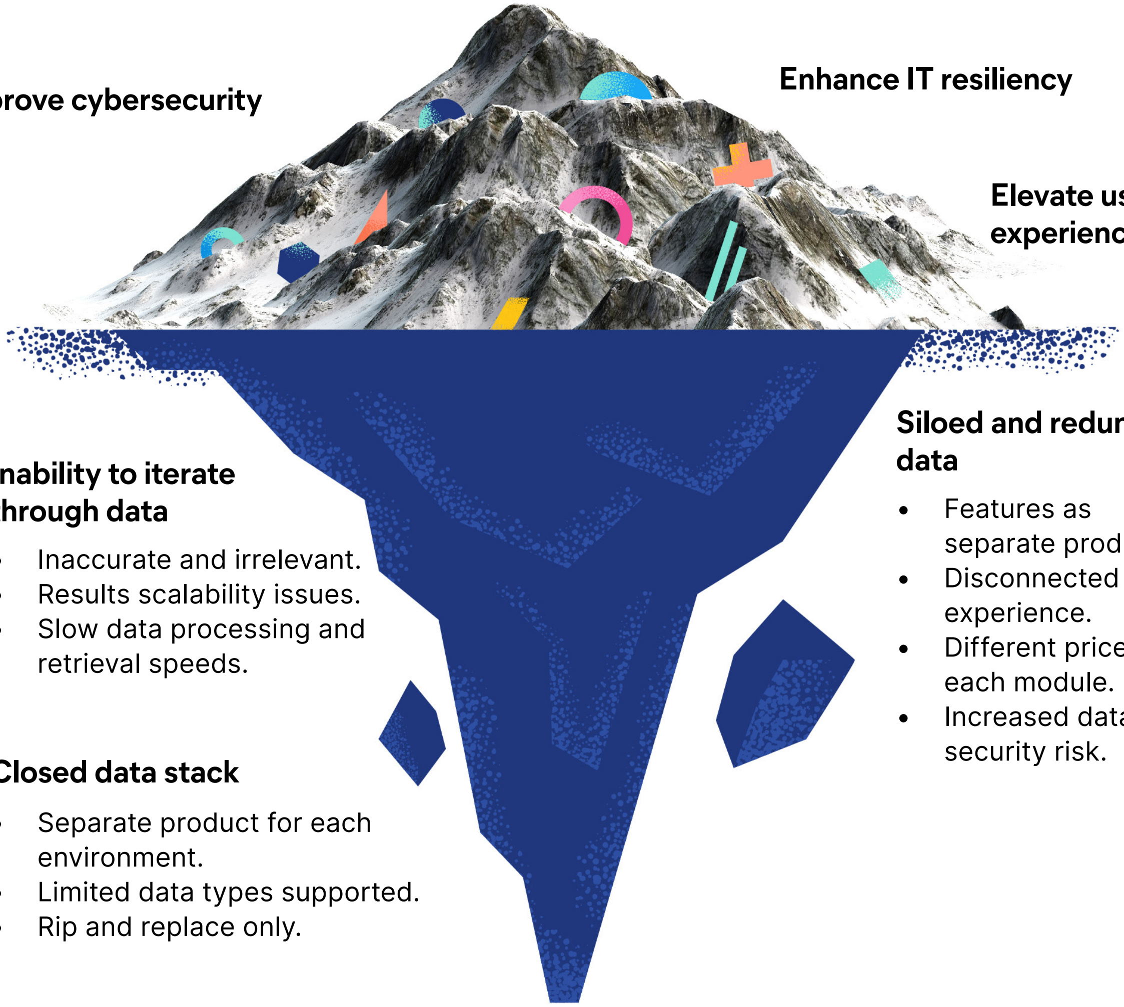

Improve cybersecurity

Enhance IT resiliency

Elevate user experiences

Inability to iterate through data

- Inaccurate and irrelevant.
- Results scalability issues.
- Slow data processing and retrieval speeds.

Closed data stack

- Separate product for each environment.
- Limited data types supported.
- Rip and replace only.

Siloed and redundant data

- Features as separate products.
- Disconnected user experience.
- Different price for each module.
- Increased data security risk.

Your giant mountain of data leads to a mountain of tools tasked with managing and making use of all that data. This leaves you with an even bigger problem that is hidden from plain sight.

The traditional tools and methods you deploy to solve business problems lead to:

1. **An inability to iterate through data.** Slow iteration speeds can lead to inaccurate and irrelevant results and scalability issues. This impacts the ability to iterate through your data in real time. And when you can't iterate, your data is useless.

2. **A closed data stack.** With different products for each environment, there are limited data types that can be supported by each product within each environment. That means that you won't be able to correlate data across those environments; you'll never be able to get a holistic view of all of your data. When you want to update your product, you will need to rip and replace the entire system.
3. **Siloed and redundant data.** When you have all of these different products that each perform a few tasks, you're getting fractured user experiences for your data teams, increased costs with different pricing for each module, and ultimately, siloed and redundant data. All of this leads to increased security risk, unplanned downtime, and poor customer experiences. You then have a sizable data tech stack that doesn't allow for your data to be used across multiple environments.

Demands escalated by ongoing data challenges

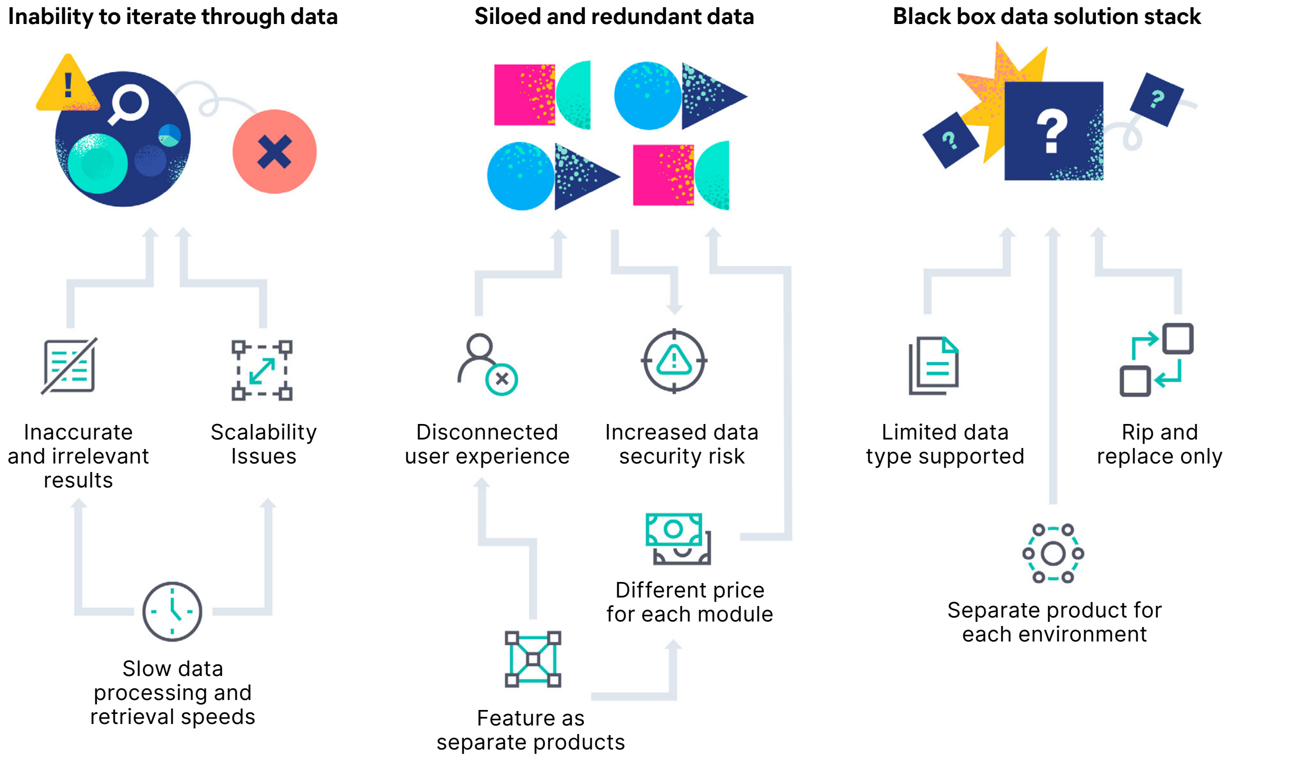
Inability to iterate through data
!
Inaccurate and irrelevant results
Scalability Issues
Slow data processing and retrieval speeds
Siloed and redundant data
Disconnected user experience
Increased data security risk
Different price for each module
Feature as separate products
Black box data solution stack
?
?
?
Limited data type supported
Rip and replace only
Separate product for each environment

Section 3:

Why you need to combat these challenges

We already addressed how data volume will continue to exponentially increase. And with that increase in data, we'll see an increase in the complexity of security threats, requiring more monitoring. As the use of SaaS solutions increases, so will the number of cloud providers you use, and so will the number of partners using those solutions and providers. With the increased economic uncertainty, organizations will find it more important than ever to have a tech stack that consolidates solutions and saves money.

> **480 EB of data produced daily by 2025**
> **(1 EB = 1000 PB = 1,000,000 TB)17[14]**

Organizations that use real-time data for the right purpose[15] are:

8x
more likely to grow revenue by 20% or more

1.4x
more likely to uncover new revenue streams

1.6x
more likely to create data-driven experiences

1.8x
more likely to commercialize their data

[14] IDC, Worldwide IDC Global DataSphere Forecast, 2022-2026, 2022.

[15] Forrester, The State Of The Insights-Driven Business, 2022

Section 4:

Combat these challenges with the precision of search and the intelligence of AI

The way for you to make use of all this data is through a solution that combines the precision of search with the intelligence AI that is delivered on a single platform with a flexible, scalable architecture.

When it comes to your data, search and AI help you:

1. **Capture all of your data**
   The first step in making use of all this data is bringing it all together. From the cloud to on-prem and machine-generated data to consumer-generated data, all of your data needs to be brought together — very quickly and at scale — into a system where you can derive value from it.

   Every error message, security event, text log line, and time stamp needs to be captured from all of your data sources. Terabytes of this type of data are being created every day. Sure, one single log could mean nothing, but it could be the indicator of the latest ransomware attack or an indicator that your system is about to reach max capacity.

   > **Before any business critical event happened in the past, there were data points that signaled there would be a crash or a security breach.**

   You need a line of sight into everything happening in your infrastructure. Search makes it possible.

2. **Make your data searchable and analyzable**

   Once you've captured all of that data, you now need to make it searchable and analyzable in real time . Don't be content with just storing your data. You need to put it to work for you.

   Think about your unstructured data, like a log file. You're not sure which elements of that log file will be of value to you. By asking freeform questions through search and iterating your search terms, you can quickly get the actual information you're looking for.

   Humans aren't the only ones that can make use of your searchable and analyzable data. Machines can as well. When you have a million log lines to go through, you need AI running analysis across the data to find the insights you want.

3. **Make your data explorable in real time**

   Make your data visual, explorable, and easy to interpret with curated workflows to help you make sense of your data and take action. You may currently have disconnected workflows and solutions that are drawing from over 400 data sources[16] and using over 300 applications[17] to manage it all. This impacts the time it takes to gain insights from all that data and ultimately, the time it takes to act on it.

   Imagine one of your security analysts has siloed data sources that are only accessible through different tools. They would have to go through each tool and tie the context together manually. That wastes their time and yours — leaving you vulnerable to security risks.

[16] Matillion, Matillion and IDG Survey: Data Growth is Real, and 3 Other Key Findings, 2022.

[17] Businesswire, Less than Half of Company SaaS Applications Are Regularly Used by Employees, 2021.

By exploring the data seamlessly within one unified platform with the intelligence of AI, you can iteratively refine the information and apply relevance to the data insights in real time. This enables you to derive value from all of your data continuously in real time to address broader business challenges, opportunities, and priorities, like improving customer experience, improving resiliency, and mitigating security risks.

The precision of search and the intelligence of AI give your enterprise:

Speed: Just like Google, type out your query and press enter to get answers in natural language from all of your data sources. Not seeing what you're looking for? Just enter another query and get those results just as fast.

Scalability: As your data grows (and it will grow!), the combination of search and AI allows you to seamlessly meet your needs at any scale, with no hardware-driven limitations.

Relevance: The context of a search will be different between a security analyst conducting a search versus an SRE versus a customer. Context matters. Search and AI provide relevant, contextual results.

Iterative exploration: All of these aspects combined gives you the opportunity to iteratively explore and analyze all of your data. You are able to slice and dice in different ways by searching different terms

How your solution goes about putting your data to work in real time matters. Also, how that solution is delivered matters. Your search and AI solution needs to be simple, flexible, and work across all of your environments without you needing to move your data from where it is.

Section 5:

What to look for in a search and AI platform

When comparing solutions, we've compiled the many features you should consider.

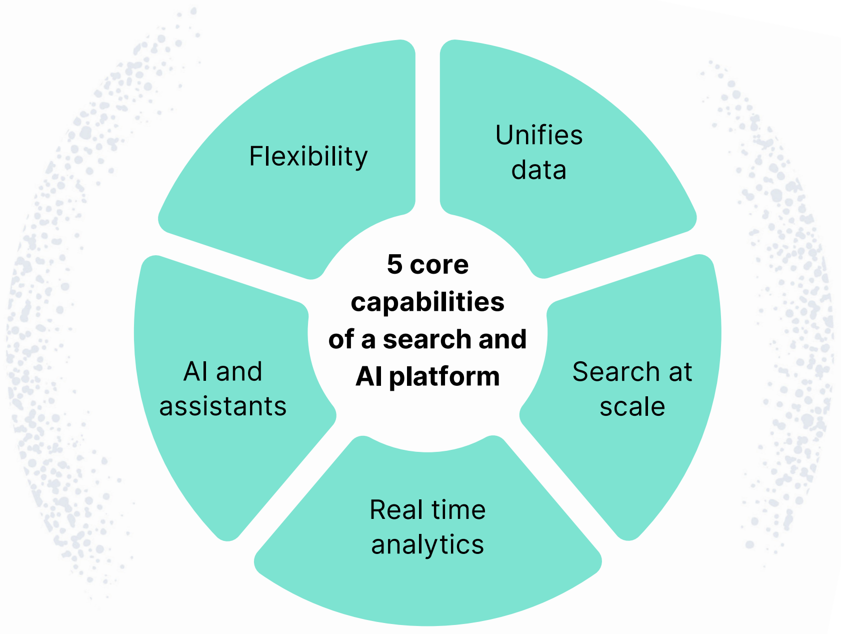

1. Unifies your data on a single platform

A solution that combines the relevance of search with the intelligence of AI that is delivered on one platform provides simplicity through an end-to-end experience with your data lifecycle. From ingest to insights, you're using one data store that uses lifecycle management, search capabilities, access rights, and machine learning. And all of this needs to encompass data from the back end to the front end of your systems.

With a single, unified platform, you get:

☑ **A unified user experience** across all of your solutions and data stores. With one, single-user experience, IT teams don't need to relearn a new tool each time a new solution is deployed.

☑ **Uniform resource-based pricing** means you only pay for the resources you consume, independent of your use case. This is essential as you scale and add more users (and, of course, more data), knowing you're only paying for the additional resources you consume.

☑ **A single data store** reduces data redundancy by storing all of your data across different solutions. Since observability data can be identical to security data, there's no need to store that twice and waste resources. With a single data store, you decrease licensing and storage, hardware, and infrastructure costs.

☑ **A common schema** allows you to bring data together that may not reside in the same place, doubling the value of your data. For example, when it comes to user behavior monitoring, you're monitoring customers to understand their buying behavior. But this data is also useful when it comes to security. You can monitor the same data to look for anomalies and patterns. Are these humans interacting with your application? Or bots? Bringing together all these solutions and all of your data allows you to unlock these insights that you may have missed if they weren't correlated.

2. Search at scale

You need a solution that provides relevant results at scale, regardless of data volume. Your solution should have the ability to surface the right information to the right people at the right time by inheriting the speed, scale, and relevance of a search engine. Regardless of the size and diversity of the data pool—for example, you could have petabytes of different data types including geolocation, IP addresses, PDFs, etc—your solution should be able to search for anything at scale.
When you can search at scale, you get:

☑ **IT teams who can minimize downtime** and accelerate root cause analysis by quickly surfacing relevant telemetry data.

☑ **Employees and customers who can find answers fast** across applications, websites, and everything in between. This leads to better customer experiences, increased customer trust, and improved employee satisfaction.

☑ **Security teams who can prevent, detect, and respond** to security threats at scale through real-time analysis of security events and telemetry data.

2. Real-time analytics

You can't wait days for a query to run. You need answers in the moment. The speed of information retrieval is imperative so you can make use of your data in real time.

With a solution that focuses on real-time analytics, you get:

- ☑ **Distributed computing / parallel processing** that allows for the processing of large volumes of data in parallel. When a query is executed, your solution should be able to distribute the query across multiple nodes and shards, each handling a portion of the data. From this, you can quickly get results, even from complex queries and large datasets.
- ☑ **Pre-indexed data** which means the data is indexed as it's ingested, creating optimized data structures that allow for rapid querying. Rather than scanning through raw data at the time of the query, your solution should leverage pre-built indexes that can quickly locate and retrieve relevant information. This provides users with a seamless experience when searching.
- ☑ **No data dehydration,** which eliminates the need for rehydration, meaning no need to reload or reconstruct data from a slower storage medium—like a traditional database and data lake—during a query. This ensures that data retrieval is quick and efficient.

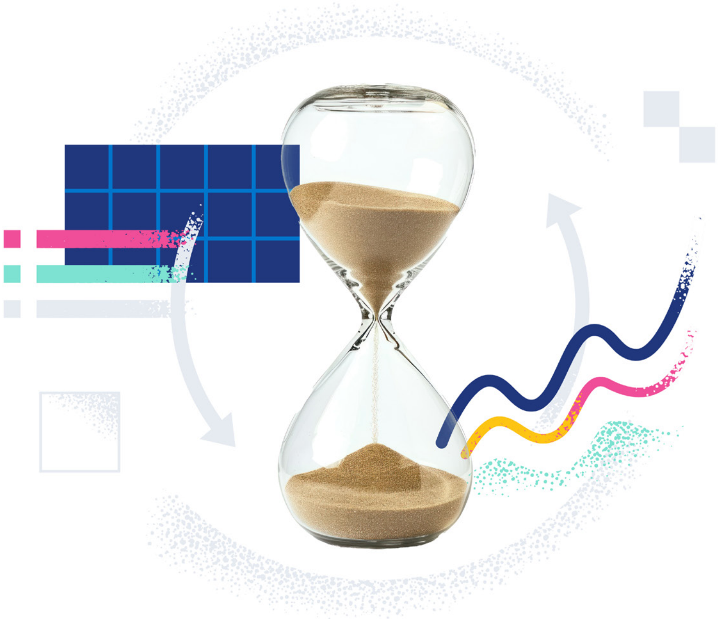

3. AI assistants

AI assistants should be built into your chosen solution. This makes it seamless when you want to ask questions of your data and get natural language responses. By connecting your private data, whether it's business or operational, a knowledge base, or a case history, with an embedded AI assistant, you can elevate conversational AI experiences through retrieval augmented generation (RAG). This will bridge knowledge silos and accelerate problem resolution.

With an AI assistant, you can help:

- ☑ **SREs interpret log messages and errors,** by providing suggestions for optimal code efficiency, writing reports, and even helping identify and execute a runbook. Using an AI assistant, SREs can enable faster problem resolution, improve collaboration, and unlock knowledge silos, allowing teams to focus on building better software
- ☑ **Boost your cybersecurity operations** and make SIEM migrations easier with generative AI. The AI assistant should help with alert investigation, incident response, and query generation or conversion using natural language.
- ☑ **Close the cybersecurity labor gap,** which has risen to a record high of just under four million, by automating repetitive and time-consuming tasks, freeing up analyst time to focus on proactively hunting threats. It also bolsters more junior analysts who no longer need specific domain knowledge to perform certain business-critical tasks. Empowered by access to technical knowledge and capabilities through generative AI, a wider range of professionals are suddenly able to take on cybersecurity roles.

4. Flexible architecture

A flexible architecture provides extensibility, allowing you to scale and grow seamlessly and rapidly. You can ingest everything and bring it all together, which means all of your data, images, documents, logs, metrics, geo locations, IP addresses, and everything in between is in one place.

With a flexible architecture, you get:

- ☑ **A multi-purpose data store.** No matter if it's your security team storing security events or your dev team storing application traces and profiling data or your business team storing product information, a flexible architecture enables you to store it all.
- ☑ **The ability to deploy anywhere.** From cloud, on-prem, hybrid, and Kubernetes, you need a solution that supports everything, everywhere.
- ☑ **Integration with existing tools** so you don't have to overhaul your legacy systems (you've spent money and time on those and you shouldn't need to rip it all out to replace it). A flexible architecture will allow you to run side by side with what you already have and upgrade those legacy systems when it works best for you.
- ☑ **The ability to fill feature gaps** to fit the solution to your exact needs and extend freely with an open code base. If you do see feature gaps, and you have the resources, make sure the solution you chose allows you to take the code and customize it how you see fit. You shouldn't need to rely on the roadmap of the solution vendor.

Section 6:

Unlock your data beneath the surface

The precision of search combined with the intelligence of AI can transform your big, messy data problems into business results by helping you continuously derive value from petabytes of data — in real time. And you don't have to break the (IT) bank because you only need one, flexible platform that combines the precision of search and the intelligence of AI to do it.

See what else you can do with a platform that combines the precision of search and AI:

Mitigate security risk

Improve operational resilience

Enhance costumer experiences