

# Sonasoftware

Introducing  
**Sonasoftware SAIBRE**



Sonasoftware AI Build  
& Runtime Ecosystem

## Executive summary

AI underpins the 4th industrial revolution. However, delivering successful AI projects is a challenge due to the growing burden of AI debt. Here, we explain how SAIBRE, or AI ecosystem, can help streamline the process of delivering successful AI applications.

### What is AI debt

AI debt is a measure of how far behind you are with your AI adoption strategy. There are three elements to be aware of. Data AI debt relates to issues with collecting all the relevant data and correctly analyzing it to avoid issues like bias. This can be a real challenge if your data is spread across multiple systems. Technical AI debt is a measure of how brittle your AI solutions are. AI models have a tendency to drift and break over time. This is exacerbated by the tendency to create each AI solution using a one-off custom process. Finally, human AI debt reflects the difficulty in recruiting a complete team with the expertise needed to deliver end-to-end AI solutions.

### What are the constituents of any AI ecosystem?

Every AI solution requires four key building blocks. Firstly, a data platform, which is able to collate and engineer the data into AI datasets. Then this data is used to train one or more machine learning models in the AI platform. Once you have a fully trained model, you need a runtime and strategy to allow it to be deployed into production. Finally, you need to put in place monitoring for data health, model drift, and overall system performance. SAIBRE delivers all these elements, hence our decision to name it the Sonasoftware AI Build and Runtime ecosystem.

### How does SAIBRE deliver end-to-end AI?

SAIBRE enables our team to deliver complete end-to-end AI solutions for companies across many industry verticals. Using SAIBRE, we can work with your team to conduct feasibility studies on your existing data. If these suggest there is a workable AI solution, we can then move on to building a proof of concept for you. This involves using SAIBRE to train and validate a suitable machine learning model. Finally, SAIBRE's light-touch runtime allows us to deploy your new solution to production. Sitting alongside all this is our AI-powered smart monitoring solution, which ensures your new AI application is always performing at its best.

## Introduction

AI underpins the 4th industrial revolution. It is the most disruptive technology since the emergence of the Internet. AI promises to be transformational for businesses right across the board. It can help cut costs, boost productivity, increase efficiency, and give new insights. With AI, companies can adjust prices in real-time to maximize sales. They can predict stock demand to ensure they never overstock or run out of products. They can accurately predict credit risks and ensure they only lend to the right people. This transformative power encourages lots of executives to launch AI initiatives of their own. Yet the majority of AI initiatives will end up failing due to AI debt.

## What is AI debt?

AI debt is a way to understand the negative impact of failing to deliver successful AI transformations. There are three component parts: data AI debt, technical AI debt, and human AI debt.

- **Data AI debt.** As you will see later, AI is completely dependent on data. This is why data scientists are core team members for any AI project. However, data can cause a number of problems which we collectively refer to as data AI debt. These include issues from collecting all your data from the diverse locations it is stored in to adding unwanted bias during data engineering.
- **Technical AI debt.** AI applications can be challenging to build and deploy. All too often, teams approach them as one-off projects, resulting in brittle designs that are hard to deploy and prone to failure. This is in marked contrast to how you would build any other modern application.
- **Human AI debt.** One of the reasons so many companies suffer from AI debt is a lack of appropriately skilled and experienced team members. Delivering a successful AI project requires a mix of skills that is hard to find. You need to build up a team with skills ranging from mathematics and data science to AI engineering and MLOps.

## How do we address AI debt?

Our AI ecosystem, SAIBRE<sup>1</sup> can help you to achieve successful AI transformation. It streamlines the process of delivering end-to-end AI solutions. SAIBRE is designed to be pragmatic and lightweight. We developed it to help our own team deliver end-to-end AI projects faster and more reliably. As such, it prioritizes efficiency, speed,

---

<sup>1</sup> SAIBRE stands for the Sonasoftware AI Build & Runtime Ecosystem.

and simplicity over features and raw power. SAIBRE specifically addresses data and technical AI debt, and working with our team allows you to solve your human AI debt.

## What do we actually mean by AI?

Artificial intelligence or AI is not a new concept. The term has been used for at least 70 years. But it only recently became mainstream as computers became more powerful. AI can exhibit narrow or general intelligence. Narrow intelligence covers systems that can solve one very specific problem, but can't be used for anything else. General intelligence systems are able to teach themselves to solve problems. When we talk about AI we are almost always referring to the former. Although systems like DeepMind and GPT3 are starting to exhibit true general artificial intelligence.

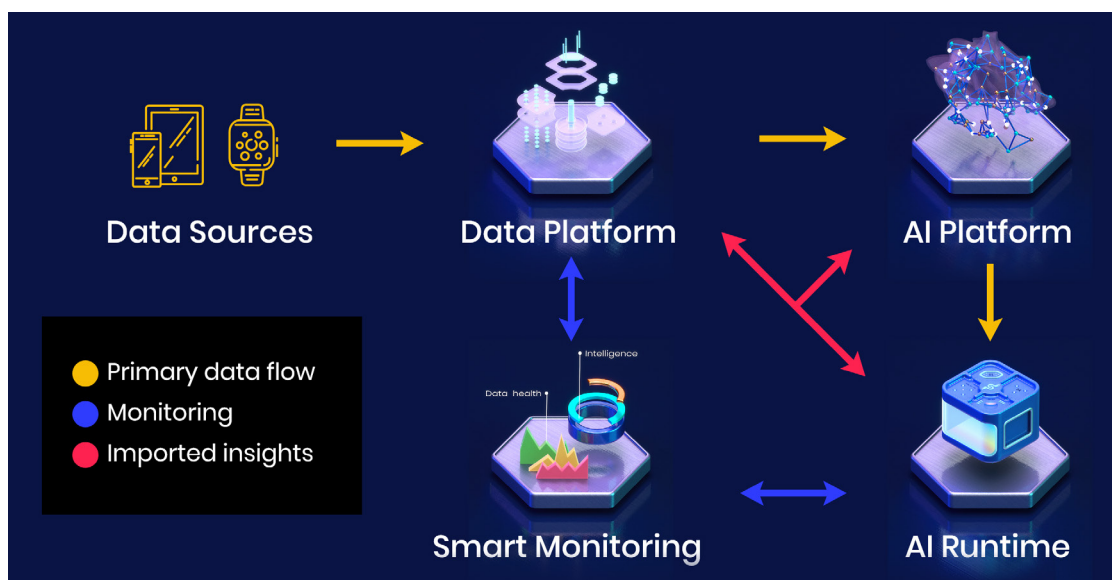
For the rest of this document, we will be talking primarily about narrow artificial intelligence. More precisely, we will be looking at systems that leverage machine learning (ML). That means they rely on approaches that allow a computer to teach itself to solve a particular task. This category covers the overwhelming majority of practical use cases for AI in the business world. In most cases, we rely on **supervised learning** models. These are great when you have data that is well described and where it is obvious what you are looking for. We also use **unsupervised learning** at times—this is useful when you only have the raw data and may lack context or granular understanding of what it shows. Finally, **reinforcement learning** systems work using a system of virtual rewards and punishments. Whenever the model makes a good decision it gets rewarded. If it makes a poor decision it will be punished.

## The constituent elements of any AI solution

Every AI solution is created from a number of basic building blocks as shown below.

### Data sources

AI solutions are completely dependent on data. Without data you won't have any AI. That data can come from many different sources. For most businesses, this will primarily be your historical data records. But other sources include sensors, image data, external data feeds (for instance weather data), and 3rd party data. Data may be discrete or time series, digital or analogue, binary or text based. Whatever form it's in, the more data you can find, the better your solution is likely to be.



The elements you need for any AI solution

## Data platform

All these different data sources tend to be spread across multiple platforms, with different data types in different formats. These need to be collated into one central location and then processed. Traditionally, AI used multiple elements to achieve this: data lakes for storage of bulk unprocessed data, data hubs to convert that data into a usable format, and data warehouses to store the processed datasets. A data platform combines all these and also includes tools to help with data exploration and engineering.

## AI platform

The core of your AI solution is the AI platform. This is the system where you take your data and use it to train machine learning models. These models can then be refined and validated to ensure they deliver the required effect. There are main families of models that are of interest here

- Forecasting
- Anomaly detection
- Classification

These will be described in more detail later. Of course, there are other types of model, but they tend to be less useful for most business use cases.

## Deployed model

Once you have a trained model, you need to deploy it in your production systems. This means solving some key problems to understand how it will integrate with your existing systems.

- 1 Where will the model actually run? As a service on the AI platform, in one of your own servers? Or deployed to multiple devices?
- 2 How will the model receive the data it needs? This includes understanding any required data processing the model expects.
- 3 What will its outputs look like? How will you make sure the model outputs are actually useful?
- 4 Who will be responsible for deploying the model and have you included them in the design process?

## Monitoring

The final part of any successful AI solution is monitoring. This covers three main areas. Firstly, you need to monitor that the data sources are working properly—a model without reliable data is useless. Secondly, you need to monitor that the AI solution is running correctly. That means good old-fashioned systems administration or DevOps. Finally, you need to monitor the performance of the AI model. Over time, this is likely to become increasingly inaccurate in a phenomenon known as model drift. Once it drifts too far, you will have to retrain it or even replace it.

## Your data is valuable

Every company generates data as a byproduct of doing business. For decades, this data was largely ignored. Occasionally, a company might use analytical techniques to gain some business intelligence insights. But often the data was just archived and forgotten. Almost 20 years ago, this started to change significantly. Companies like Google and Amazon began to realize how their data could be used to drive new business models. In Google's case, this allowed them to shift from being a company that provided search results to a company that consumed data to drive advertising revenue. In turn, that has seen their parent company, Alphabet, expand to become one of the most successful technology companies in history. In Amazon's case, it allowed them to go from being an online book retailer to one of the world's largest multinationals with interests in retailing, cloud computing, and consumer electronics. But none of that would have been possible without data and machine learning.

## Types of AI application

AI can be used to deliver a wide range of benefits to businesses. Many AI solutions

are available as off-the-shelf applications that you can simply attach to your existing systems. However, these are often quite generic and only solve very specific use cases. Support agent chatbots are the perfect example of this. But what can AI achieve if you develop your own solution? Here are some typical use cases for custom AI.

## Anomaly detection

Anomaly detection is about identifying data points or transactions that fall outside the expected normal range. A good example is credit card fraud detection systems. These might flag a transaction because it happens in a different country, or is for a larger sum than usual. However, anomaly detection can be used for many other things.

**Identifying risky investments.** PE and VC firms have to assess the risk of any potential investment they make. This is done during the due diligence process. Anomaly detection helps identify red flags within the financial data and forecasts submitted during this process.

**Predictive maintenance.** Many industries rely on heavy plant or expensive production line robots. If these go wrong, they can halt production and may even cause serious damage if there's a catastrophic failure. Anomaly detection allows you to identify issues before they become failures and thus fix them in time.

## Forecasting

Machine learning is particularly good at identifying patterns in data. These patterns allow an ML model to accurately forecast future events based on past knowledge and the current state. Such models are really useful across many different industry verticals.

**Demand management.** Many utility markets consist of providers and resellers. Resellers hope to be able to sell the commodity at a profit but the price they are charged will depend on the peak demand seen by the provider. If you can predict this peak demand, you can potentially shift it and thus reduce your costs. This is increasingly common in the US electricity market.

**Stock management.** Most retail businesses operate on very tight margins. They need to ensure they have sufficient stock to meet demand, but they must avoid overstocking. Firstly, warehousing costs money and secondly, the excess stock may never be shifted if it is seasonal. Forecasting models allow you to accurately predict exactly what stock to buy and when.

## Classification

One key family of ML models allow you to classify data and thus draw inferences. Such models enable several key use cases.

**Lead scoring.** Every sales or marketing professional knows that not all leads are equal. Some leads will string you along for months and never sign while others will become your best customer. Classification systems can use your past history of leads and deals to predict which new leads are hot, which need nurturing, and which are a lost cause from the start.

**Credit risk.** Credit institutions have to assess the risk of a loan default in order to price it correctly. Often, this is done heuristically, for instance assuming a certain percentage of loans will default. But an AI model can accurately classify how risky a potential loan is based on knowledge of past defaults and the current credit climate.

## How do you unlock your data's potential?


So, how do you actually unlock this hidden potential in your data? The answer is, through good data science. Data science is not a new field, but it became big news with the advent of AI. Without it, you can't create a successful AI project, despite what some people say. But unfortunately, legacy approaches to data science still abound. These lead to three pitfalls that trap inexperienced data science teams.

- Introducing bias
- Issues with data processing
- Lack of end-to-end vision

Before you even contemplate trying to create an AI solution, you need to address these data science problems.

### Avoiding bias

The first stage of any AI project is data engineering. This means taking all the available data and converting it to clean data ready to be processed by an AI model. This almost always needs some input from subject matter experts. However, it is essential not to let their preconceptions lead to bias. Humans are incredibly good at seeing patterns and correlations. So good, that we often see patterns that aren't really there (pareidolia), or are purely coincidental.



Many people will see an image of a face in this rock formation

We often find that companies rely heavily on their past experience when it comes to data engineering. But the pace of change is so fast nowadays that data features quickly lose or gain significance. Here at Sonasoftware, we ask our clients to paint the big picture for us. This lets us understand what data there is and what it might show. However, we then analyze the data in-depth to ensure it supports and suppositions the client may have.

## Improving data processing

Having completed the initial data engineering, we then start the Exploratory Data Analysis. Data is king for AI, and in this phase, we try to determine whether there is sufficient data to create robust AI models. This requires both skill and patience from our data scientists. They need to resist the temptation to move forward until they are sure the data fidelity has been established.

The adage here is garbage in, garbage out. In other words, if you feed garbage data into your AI model, you just get garbage out. This is true however sophisticated your AI platform is. And it is one of the biggest reasons why so many AI projects fail. We often see companies trying to create AI solutions on the basis of insufficient or inaccurate data. By contrast, if you rely on Sonasoftware, we will ensure the data exploration is done correctly. Usually, we find the right data and can continue to the next phase. If you don't have the right data, we will give you advice on how to start collecting the data you need.

## **AI is an end-to-end problem**

One of the biggest mistakes companies make is to lose sight of the need to actually deploy their new AI solution into production. AI models are trained and validated on complex infrastructure. Actually deploying them in production isn't easy. You will be reliant on your engineering team. But engineering teams are typically cautious when they hear about great new tools that will transform the company. That's because they end up having to do all the deployment and maintenance. This can be painful and time-consuming. Then there's the perennial issue of over-promising and under-delivering.

### **The wrong metrics**

Firstly, we often see people choosing the wrong metrics to evaluate the success of their AI model. In turn, this leads to a model that may not be useful in production. Effectively, they focus on creating a perfect model, rather than solving an actual business problem. Successful AI initiatives establish what metrics really matter to the business. Only then can you validate if your solution is delivering good value and ROI.

Secondly we often come across naive over-optimism. Too much effort is focused on the upside, sometimes forgetting the risk that comes with probabilistic outcomes. All too often, data scientists search for accuracy without paying attention to reality. This is when AI teams start promising "100% improvement in efficiency" or "halving the costs." Many AI initiatives have died due to a loss of confidence and a perception that the project didn't deliver on its promises.

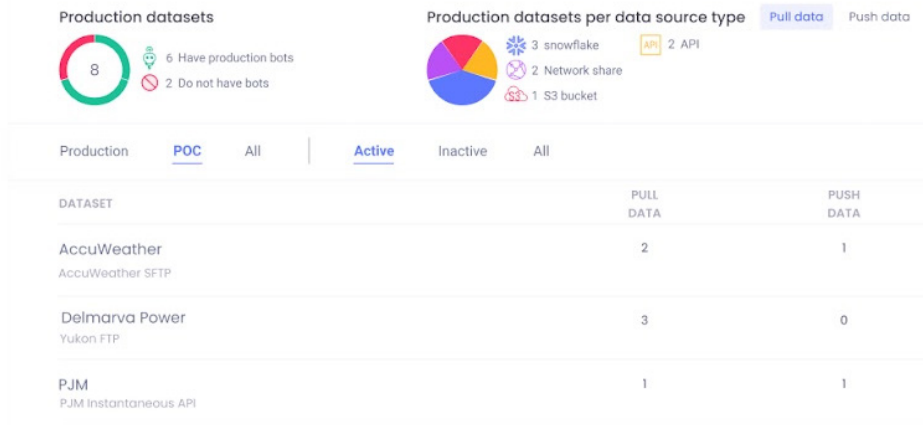
## **SAIBRE—a complete AI ecosystem**

We developed SAIBRE to solve the problems we have identified working on AI initiatives with clients. SAIBRE consists of four modules that all work seamlessly together. These deliver all the elements required for an AI ecosystem except for the raw data sources.

### **The SAIBRE data platform**

As we explained above, you need to put at least as much effort into data exploration as you do in creating the ML models. SAIBRE includes a data platform that helps streamline this process. It allows you to import data in multiple formats, or to link data from an API or public 3rd party data source. It provides a suite of tools for data


exploration and engineering. These let you clean the data and transform it into a usable dataset for machine learning. Importantly, this also makes it much easier to reuse data between different projects.



SAIBRE makes it easy to build, maintain, and reuse datasets

## SAIBRE AI platform

Once the data platform has done its work the dataset is fed into the AI platform. This takes a modular approach to building ML models. The data can be further engineered and refined, or combined with other datasets. Models can be added from a growing library of templates, or it allows the addition of custom code. The model can then be trained and validated. It's even easy to run model competition where several models are trained in parallel and the best one is selected according to the metrics that matter.




SAIBRE allows complex AI models to be built from simple functional blocks

## SAIBRE runtime engine

From the outset, SAIBRE has been designed to simplify deployment. The runtime engine is lightweight and efficient and can even run on many edge devices. Models can be deployed with just a few clicks. You just need to specify which endpoint they will

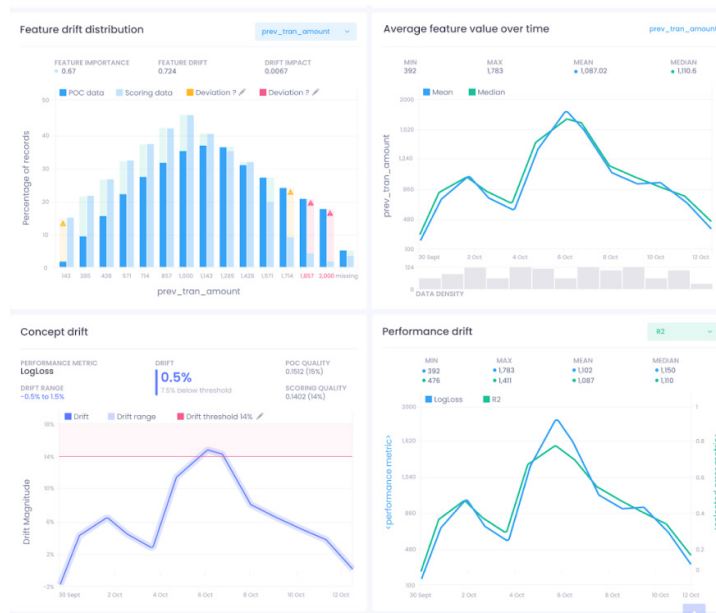
connect to, where data should be sent, and what the model is called. The resulting AI application can then be run in Sonasoftware's infrastructure, or you can receive an installable container to run on your own server.



SAIBRE lets you run multiple different AI models in production

## Smart monitoring


The fourth element in the ecosystem is our AI-powered smart monitoring solution. AI solutions often suffer from a key issue called model drift. There are various forms of drift, but all result in the model becoming less accurate over time. A simple example of drift is the impact of unexpected events. Imagine you run a hotel business. In 2020, 90% of your business vanished in the pandemic. Suddenly, all your forecast models for room bookings became obsolete. This is an extreme example, but many factors have a similar (albeit lesser) impact, such as economic climate, weather, and fashion trends. Smart monitoring uses AI to identify when your models are starting to become less accurate. It also includes tools to monitor the quality and reliability of your data feeds. For instance, you may rely on data coming from an external feed that goes wrong and starts to just return a default value.



Smart monitoring can identify AI model drift before it causes problems

# What does a Sonasoft AI project look like?

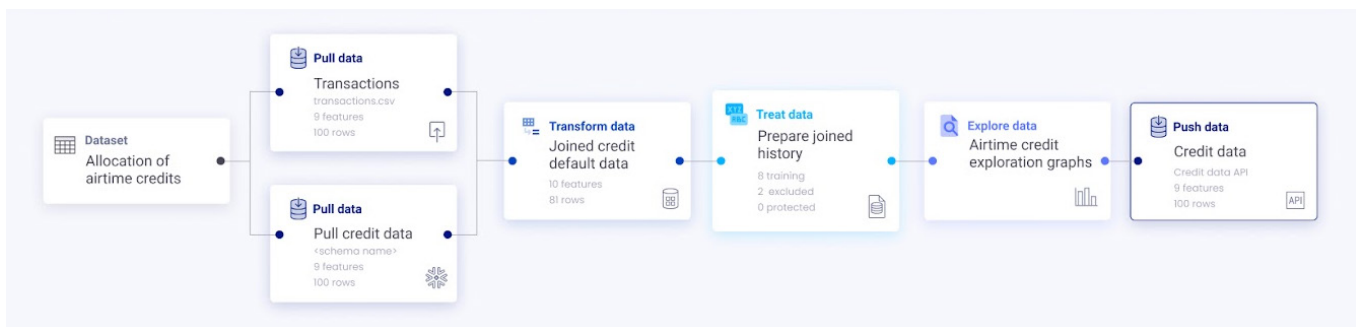
We are firm believers in the transformative power of AI. However, we are also very pragmatic. AI is not magical, even if sometimes people think it is. Our aim is to help clients to deliver practical and impactful AI solutions that actually work. To achieve this, we have developed a 3-step process driven by SAIBRE.



The Sonasoft 3-step process for delivering AI solutions

## Step 1. Data feasibility study

When you first come to us looking for an AI application we will offer you a data feasibility study. Our data scientists and AI engineers will work alongside your team to locate all the relevant data and import it into SAIBRE’s data platform. Then they will collaborate with your subject matter experts to understand what the data is showing. During this process, our team will have several conversations with you to understand what you are looking for and to give you an idea of what is feasible with your data. It may be that you came with a specific idea for what you want the AI application to achieve. Hopefully, the data will be suitable. If it isn’t, we will either suggest alternative applications that are feasible, or will explain how you can get the additional data needed. This whole process generally takes 4-6 weeks to complete.

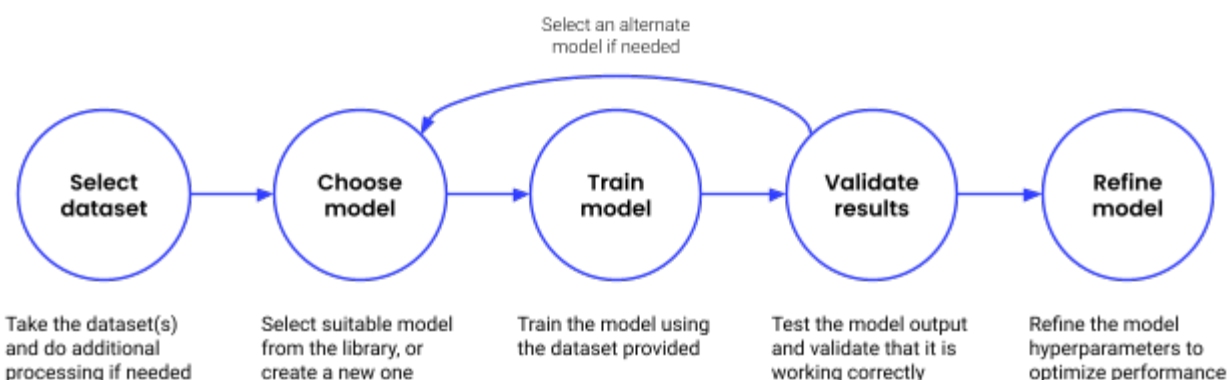


A key part of the feasibility study is creating datasets

## Step 2. Proof of concept

The next step is to sign a contract for us to build a proof of concept (POC). As soon as we have the signed contract, our team will start building the actual AI application. They will take the cleaned dataset from the feasibility study and use it to train a suitable ML

model. This is actually quite a complex process, since the team may need to try several models before they find one that is robust enough.



SAIBRE allows you to construct complete ML pipelines in order to create trained models

SAIBRE also allows us to conduct a process known as model competition. Here, we train several models in parallel and choose the best of these to take further. This speeds up the process significantly. However, the whole process will still take 3-4 months in most cases. This is to allow us time to ensure the model will perform as well as it possibly can and to check that it is robust enough. Once we have a fully functional model we will demonstrate how it will behave when deployed as an AI application.

## Step 3: Deployment to production

The final step is to deploy the model using the SAIBRE runtime environment. There are two options here. Some companies choose a complete standalone install. In that case, you will be given a container that hosts the complete model, the runtime, and all related config files. Alternatively, you can choose a managed service. We will deploy the model for you and monitor it with our smart monitoring system.

### Deployment options for SAIBRE AI applications



SAIBRE AI models can be deployed in three types of environment

As shown above, SAIBRE gives you flexible deployment options. We will advise you which approach works best for your deployment scenario. In all cases, deployment is designed to be near zero-effort for your team.

## Conclusions

SAIBRE is a complete AI ecosystem that is optimized for delivering end-to-end AI solutions. It is a core part of our “zero-effort AI” ethos and helps us provide high quality and reliable AI solutions across many industry verticals. Thanks to SAIBRE, you can benefit from real AI transformation and eliminate your AI debt.

**Zero-effort creation.** We take away the pain of creating fully functional AI applications. Our team brings decades of experience and will ensure that your AI application delivers real value.

**Zero-effort deployment.** Deploying AI applications can be hard. Typically, this is left entirely to your own engineering team. SAIBRE allows you to deploy your new application far more easily.

**Zero-effort maintenance.** SAIBRE smart monitoring will ensure that your AI application is working properly. Where the system detects a problem it will allow you to retrain the model with just a few mouse clicks.

If you want to experience the power of SAIBRE, reach out to us today.