



— SOLUTION BRIEF

Cloud Native Business Intelligence Products

THE INSIGHTS NEEDED TO DELIVER A RELIABLE PLATFORM

Contents


Product overview and assumptions	1
Potential pain points	3
The proposed solution	6
Expected outcomes	10
Separation of Concerns	12
Contact details	13
FAQ and the answers	

Data-driven movement is gaining traction in many companies around the world. Making decisions based on data is a modern way to conduct business in today's economic climate. In the long run, this puts companies ahead of their competitors, still living in an old era of "decision guessing".


Thus, many different business intelligence products are being built to support decision makers. **But building a successful SaaS in the ML and BI area poses many engineering challenges.** This solution brief helps you understand how utilizing AWS Cloud in Business Intelligence solutions enables rock-solid products that are easy to build and ready to scale up.

Product overview and assumptions


Usually, a product's main goal is to solve its customers' (tenants') problem by analyzing their data and providing valuable insights or suggesting an action. Artificial intelligence and machine learning is often heavily involved in the process of generating these insights.



As a first step, each **tenant's data is imported** into the system or, in some industries, held and accessed directly from their data centers. Various mechanisms are involved to perform import and ETL functions depending on the scale and characteristics of incoming data (real-time, streaming, batch, incremental, big data).



Whenever fresh data is loaded into the system, a set of data engineering tasks (such as feature engineering) is performed. Afterwards, **machine learning algorithms are trained** to provide important insights.



The end result and the actual output for a tenant of the product varies, but usually consists of a set of interactive dashboards that can access imported data and the machine learning results. **All insights are visualized** to allow tenants to better understand their business. It can also generate a PDF report or provide data for other apps via a dedicated, secure API.

The product itself is usually accessed via a dedicated website that embeds these dashboards. Each tenant can have their own set of users mimicking the structure of their organization and create their own group of users with a different set of permissions. Each group can be allowed to perform various actions within the application.



Potential pain points

There are many issues and concerns that might arise at various stages of such projects. Even if business goals are clearly defined and data scientists have their models ready, the purely engineering side of such projects requires a vast amount of expertise.

Beginning with **the lack of a research lab**, your data scientists and business intelligence developers might have no place to conduct their experiments. Depending on the type of the product, they may be limited by the lack of compute power, memory constraints, reproducibility issues, or a lack of experiments or model registry.





A team or company might have **no SaaS experience** either. **Even if the product was successfully delivered to one tenant, this itself is only a milestone on the long road ahead.** The questions along “how do we scale up businesswise? how can we deliver the same business intelligence product to dozens or hundreds of customers at the same

A dedicated tenant portal, a smooth onboarding process and a huge focus on “reusability” within the product are absolutely essential. A change or a new feature for one tenant must be applied to all tenants with as little manual work as possible. It is often unclear how to achieve that in every step of the process.

Technology must inevitably track your business needs and goals closely. **Scaling the platform up in terms of compute power is crucial.** You must ensure that you can rapidly add new tenants as they go. At the same time, you cannot create overengineered behemoths due to their cost. Architecting a solution that costs little to nothing in periods of low usage, but can automatically scale up to thousands of tenants is not an easy task. Serverless services can help, but experience is required to build on top of them. Additionally, ensuring consistency between tenants, handling transparent deployment of new features, structuring data warehouses and lakes, having the ability to rollback and deploy or customize one tenant’s product - all are hard problems to tackle.

Productionalizing ML solutions poses a challenge due to **lack of MLOps workforce on the market**. MLOps requires a specific mix of various skill-sets (DevOps, data engineering, data science), which is found in the wild. Thus, even with an experienced data science team and already trained models, it may be hard to actually deploy them in a scalable manner.

The security aspect is non-trivial as well. Siloing tenants' data is a tremendous effort if the team is unprepared. Additionally, sometimes a tenant will require a dedicated environment for their instance of product. Mitigating “noisy neighbors” issues is also a challenge.

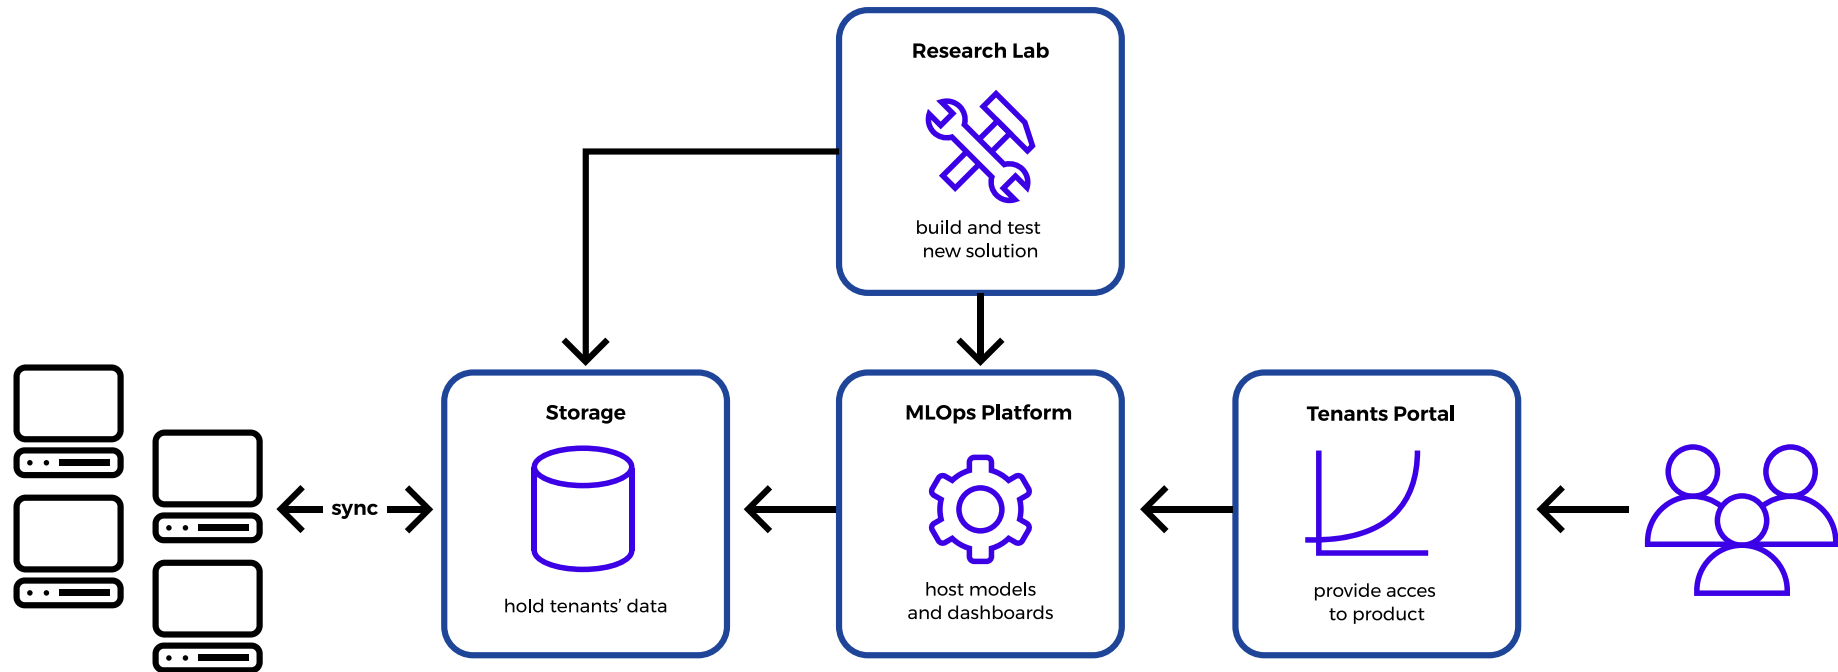


Additionally, even if solutions to these problems were identified, they usually mean that a cloud migration process has to be done. With **no prior cloud experience**, this poses a significant challenge and sometimes existential threat to the product. If you plan on using multiple SaaS data products instead of one cloud, you will end up in a **fragmented, hardly manageable engineering environment**.

Lastly, **some technological decisions made early, directly impact your future roadmap**. Are you sure that all the picks are suited for you, and have you put enough thought into each one? A wrong choice might mean months or even years of refactoring. Is your visualization tool SaaS-friendly? Can you quickly onboard new tenants? How can you reuse your dashboards or algorithms between tenants? There are hundreds of such questions waiting to be answered.

The proposed solution

The typical product naturally splits in four engineering areas:



Main responsibilities of the Storage area are:

- Holding tenants' data in a centralized manner
- Managing and maintaining import and pre/post processing jobs
- Ensuring strong data security

The Research Lab area is focused on:

- Building new prototype features
- Visualizing insights on dashboard in a user-friendly way

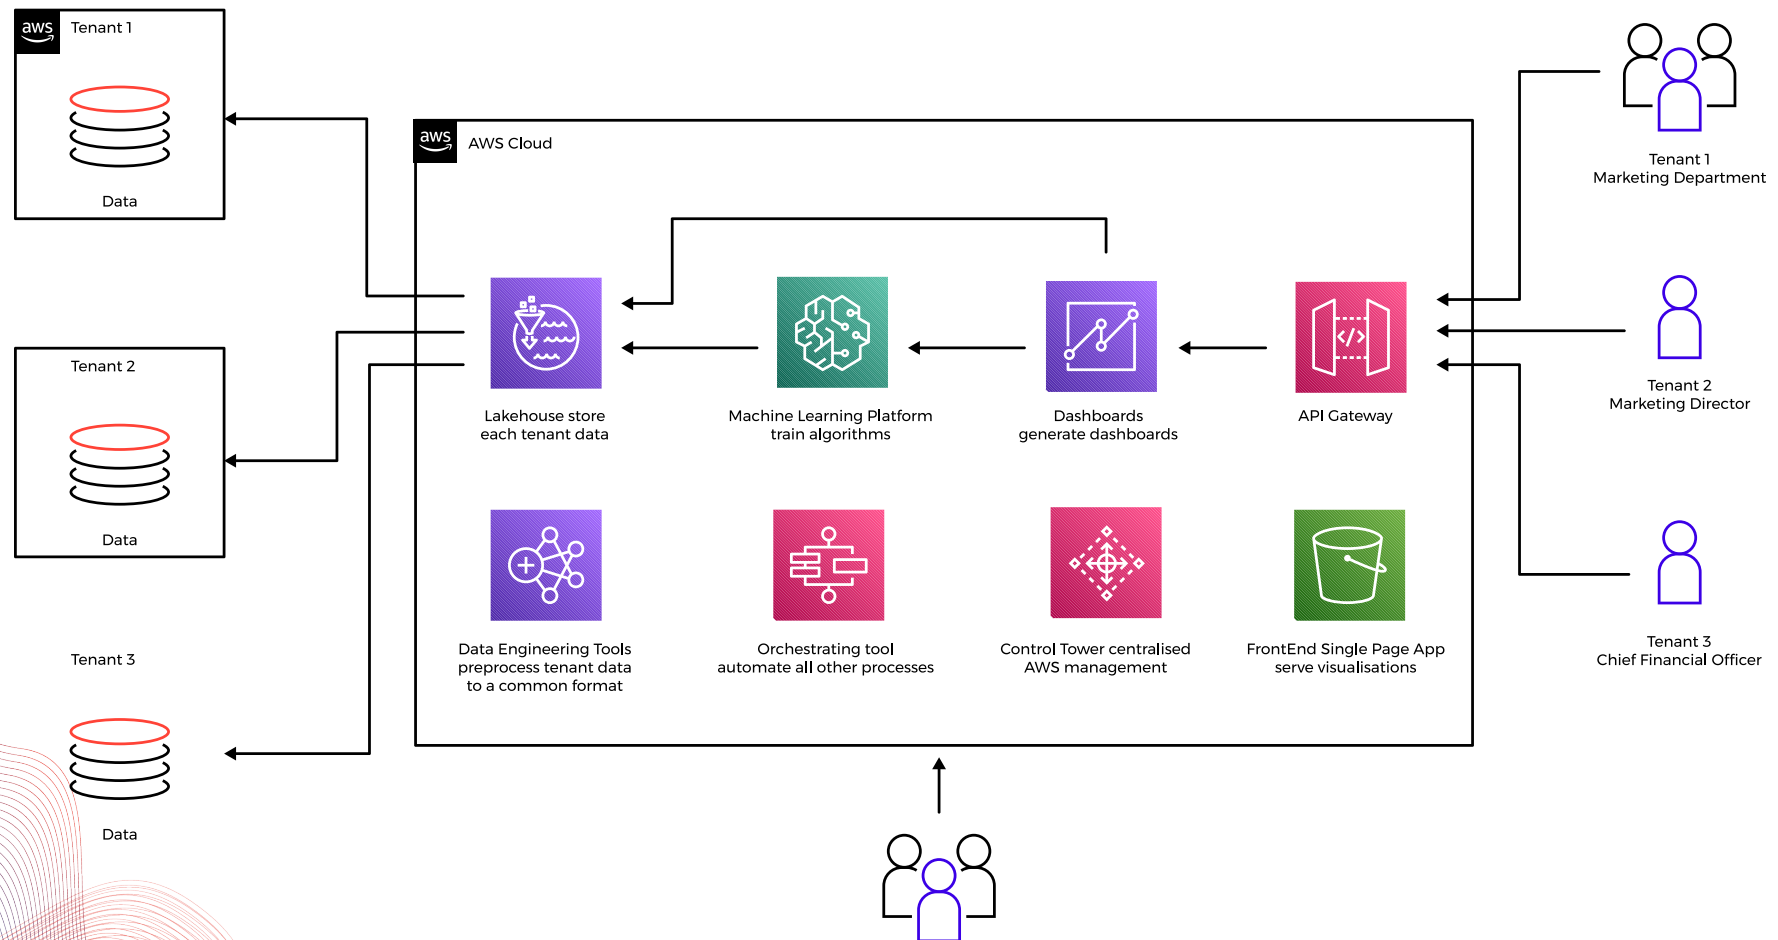
The MLOPs Platform area is here to support:

- Transforming the prototypes from Research Lab into fully-fledged features of the product
- Handling automation, resilience, observability, deployment, rollbacks etc. of machine learning models and data visualization dashboards
- Serving necessary artifacts to Tenants Portal

The goal of the Tenants Portal area is:

- Handling all tenant onboarding processes (registration, customer support, payments, invoices)
- Enabling tenants to hook their data sources into our system (via plugins or using dedicated assistance)
- Allow tenants to create their own users and user groups within their account
- Serve tenants their dashboards and machine learning models

With that in mind, we can proceed with mapping the above idea onto AWS infrastructure. While the birds-eye view presented above shows the “final architecture a few years from now on” and could suggest a waterfall methodology. In fact, every part of it can be incrementally created in each sprint iteration using agile methods. Through a series of meetings, building PoC and MVP all the way up to the final product, we solve all headaches for you.






AWS Lake Formation will serve as a lakehouse / data lake solution storing data. Infinitely scalable, low maintenance and tight governance is what best characterizes this service.



AWS Glue, AWS EMR or Amazon Kinesis will provide a data engineering toolkit, to handle all data related operations and automate import and ETL processes. These services can handle any amount of data, be it in a batch or streaming manner.



Amazon SageMaker will be the bread and butter of the research lab and MLOps platform. This fully managed set of services will handle all the machine learning algorithms across their lifecycle. From prototyping and exploratory data analysis through model training to deployment and model maintenance.



Your dashboards could be created in Amazon QuickSight, a fully managed dashboarding solution in AWS. If you prefer to stick with your existing solution (like Tableau, Power BI, Looker or Qlik), this can be easily adapted.



Automation, orchestration and observability will be provided by **Amazon EventBridge, AWS Step Functions and Amazon CloudWatch** respectively.





Finally, the Tenant Portal will be powered by **AWS Lambda and Amazon API Gateway.** Serverless, event-driven solutions will enable us to deliver the project way quicker, while still retaining predictable pricing and great performance.

Chaos Gears, as an APN AWS Partner on Advanced level, can help you obtain additional funding via programs like **AWS Activate** that can get you up to **\$100,000 in AWS credits**. Additionally, other startup-focused programmes in AWS can increase that amount, to provide your research team with necessary training via workshops and training

Expected outcomes

Most importantly, **your product will be able to handle thousands of tenants and rapidly add new ones**. Scaling up business-wise (in terms of processes) and technology-wise (in terms of compute power) will be a breeze. All SaaS & MLOps best practices are kept in place.



Tenants **data will be securely imported and stored in a semi-centralized manner in a lake house solution**, allowing your data scientists to quickly create new machine learning models. Your business intelligence developers will also enjoy a low latency storage solution, perfectly suited for dynamic, interactive dashboards.

Additionally, your **data scientists will enjoy flexible compute and other comprehensive offerings of the research lab created in Amazon SageMaker**. All models and dashboards will be reused between tenants, allowing all of them to enjoy a new feature set every couple of weeks. A/B testing and customization will also be an option.

The **tenant portal will allow your customers to access their product**, handle payment details, choose their plan and manage their users, groups and permissions.

Most parts of the platform will be delivered with **cloud-native serverless solutions**. That itself has enormous benefits - no cost in periods of low usage, predictable pricing model, low maintenance and operational overheads as well as built-in autoscaling. You can go from having one tenant to having hundreds with little effort.

Additionally, we will be able to deliver it way faster when compared to using traditional, machine-based services. The event-driven nature of the solution will ensure that little human involvement is required while operating such a platform.




Separation of Concerns

The separation of concerns and responsibilities should be clear and understood by both sides.

We bring a comprehensive cloud experience in building SaaS and ML products on AWS. While we can consult regarding what your product actually does, we know that only you understand your business best. Thus, we want to focus on enabling your data scientists and business intelligence developers to create the heart of your business and product, **while we build an automated platform on top of their work.**

You extend your research team and focus on bringing value to your customers, while we do all the hard, cloud engineering parts for you. Your product is transformed into a large-scale SaaS Business Intelligence Cloud



That means, our team consists of data engineers, MLOps engineers, DevOps engineers, solutions architects, cloud software engineers and delivery managers. Yours is mostly data scientists, business intelligence developers, business analysts, product owners, marketing specialists, customer success managers and new business developers.

You handle the business & research. We handle cloud engineering.



**Tomasz Dudek,
Solutions Architect at Chaos Gears**

Specializes in Machine Learning, Serverless and SaaS.
Tomasz evangelizes best practices for SaaS that he gathered during
delivering various, complex SaaS projects

Discuss your SaaS issues or this solution with me

tom@chaosgears.com

[BOOK A CALL NOW](#)



FAQ and the answers

Q: Do you have a demo of such a platform running on AWS?

A: Absolutely! Schedule a call with us to get in touch and see that system in action.

Q: How much does the infrastructure of a platform like this one cost?

A: A ballpark figure for one mid-sized tenant (<1TB of data, a couple hundred training hours per month, dozens of dashboard users) is around \$200-1,000 per month. The actual cost varies, depending on a few factors:

How much compute power is required to train machine learning models and how often are they trained?

How much data is stored and how often is it processed?

How often the data is queried? What are the query patterns?

Are your queries complex or do they scan only a tiny subset of the data?

How many users does your tenant have?

Can your tenants share the infrastructure or should they be siloed?

Q: Are there any constraints on technologies? My team uses a custom framework in seldomly used programming language to build train models. And we build on Tableau already!

A: Your data scientists may use any language and framework they need. The same applies to your dashboarding technologies. This is just a 100% AWS proposal due to awesome integration mechanisms between various AWS services, but many parts can be swapped for competitive technologies.

Q: Will you provide us with AWS training? We have never used that cloud before.

A: Yes. Your research team will receive a comprehensive workshop in the form of SageMaker Immersion Days to let them fully understand what SageMaker offers. We can also provide more introductory general AWS webinars.

Q: Is a centralized data lake mandatory? My tenants are worried about their data getting leaked to their competitors (which may also be using my product).

A: The architecture presented here is simplified. Each tenant may have a completely separate AWS account with their data. Additionally, for your most security oriented customers, we may design the system such that their data stays entirely on their on-premises.

Q: Will you eventually replace my own data science research team with yours?

A: As mentioned in Separation of Concerns, we absolutely don't aim to replace your core business team. The point is to cooperate together to build a cloud-native, MLOps-oriented SaaS platform that allows you to handle thousands of tenants automatically. We handle the MLOps, Data Engineering and SaaS parts, you handle the actual business and dashboards.

Q: Do you have SaaS expertise? We have never built a product for multiple customers before.

A: Yes. We have deployed multiple SaaS applications, both Business Intelligence related and not. We follow all the best AWS SaaS Factory practices with our AWS certified team. We can aid both in engineering and business parts of your products.

Q: Why AWS? Can't you build the system in another public cloud?

A: Indeed, you can build a similar solution elsewhere. However, building on top of the unquestionable leader of cloud computing ensures you always have access to the newest solutions. Additionally, SageMaker is the most developed ML platform on the market.



Q: Could we build the system on-premises using open-source technologies like MLflow, Hudi, Kubeflow, Airflow, Kubernetes, Seldon and Superset?

A: Yes. These technologies are alternatives to SageMaker, Lake Formation, Glue or QuickSight. However, the effort required to deploy, operate and maintain them is significantly higher than when you run serverless AWS services instead. Suddenly, your DevOps team might have to quadruple in size and eat all your budget - all that while delivering little business value to your customers.

Q: How does your team take care of the knowledge sharing to my team? I'd like to constantly level up the expertise among our team members.

A: We showcase every piece of the solutions in a form of demo meetings and document it in a tool of your choice. We can also provide training or more thorough elaboration of it if you wish. The solution is developed through Infrastructure as Code or Configuration as Code tools and committed to VCS like git, so you can browse it yourself and perform frequent code reviews.

Q: What size of the Operations Team is needed to maintain the aforementioned platform? How do you handle that if needed?

A: Thanks to the serverless nature of the system, the dedicated SRE/Ops team is generally not required in the MVP phase. When the system grows, we can provide 24/7 support with our in-house SRE team or help you build your own. Ops team is usually a magnitude smaller than what the alternative non-serverless solution would require.

Q: How do you monitor costs per tenant?

A: If we go with Silo architecture (every tenant has their own AWS account) you can see the exact cost per tenant. If we share the infrastructure, we can provide some estimations using AWS Cost Allocation Tags, which will show you the approximate percentage of

Q: What kind of either tools or technologies are being used to keep the overall level of the security high?

A: We follow all the best AWS Security practices. Among others - we use multiple AWS accounts for workloads, ensure least-privilege authorization mechanisms via AWS IAM and AWS SSO, set up alerting tools such as AWS GuardDuty to notify us of any breaches, collect every AWS action performed in a separate account, configure AWS Budgets to monitor the costs, block and limit unused AWS services and regions via SCP, encrypt data in transit or at rest.

Q: What does the cooperation flow between companies look like?

A: We treat each other as if we were working in the same team and at the same company. That means, we work on the same VCS repository and AWS accounts, share knowledge through Slack, Teams, Google Meet or similar tools and plan the roadmap together. We also frequently meet at Agile meetings.

Q: How much time is needed to set up the platform?

A: A functional PoC is available in a matter of days. A minimum viable product (MVP) with your first tenant onboarded usually takes us up to 2-3 months to build.



Have a questions?

We can help you!

✉ hello@chaosgears.com