Deskpro

GUIDE

You need an AI-powered help desk and data sovereignty:

A support leader's guide to getting both

Table of contents

International AI infrastructure puts pressure on data privacy and support team tech stacks

You've been evaluating AI-powered help desk software for your support team, and after issuing an RFP and going through demos and follow-up calls with multiple vendors, you've got a clear top choice. Compliance informs you that your chosen provider doesn't meet your organization's data sovereignty requirements because the AI models they use are accessed via a public cloud with infrastructure outside the sovereign area.

Or you've been using a SaaS help desk with AI features for several years, only to have your Security team tell you to migrate off it because the vendor can't provide adequate protections when your data leaves your region for AI processing.

Both scenarios are becoming common for support leaders in regulated industries and government agencies in regions with strict data sovereignty or privacy requirements, like GDPR in the EU. Data privacy laws are continuously evolving, and, with many countries adding geographic sovereignty restrictions, SaaS solutions may no longer be compliant, especially when factoring in AI features that send data to a public cloud.

Navigating the help desk market when you must maintain data sovereignty can be complex, and it can be challenging to know where to begin. We've created this guide to help you understand:

- How data privacy and sovereignty requirements impact the help desk evaluation process–especially when AI tools are involved
- The different paths you can take to keep your data in your region while still using AI
- What to look for when vetting help desk vendors

Why do data privacy laws and sovereignty requirements matter to support teams?

Organizations in regulated industries and regions with strict data governance laws are increasingly requiring their support teams to move their tech stack to sovereign environments due to concerns about foreign jurisdictional reach. One of the major drivers is the legal conflict between the 2018 US CLOUD Act and the local laws of many countries.

What is the US CLOUD Act?

The US CLOUD Act is a law that allows US law enforcement to obtain a warrant to access data stored by US-based cloud providers, even if that data is located outside the US.

Under the US CLOUD Act, the US government could theoretically compel a US-based cloud provider, such as Amazon, Microsoft, or Google, to turn over data stored in their public cloud, even if that data is stored in cloud infrastructure in another country.

This law notably conflicts with data privacy and operational sovereignty laws in the EU, including GDPR, the EU Data Act, and the EU AI Act. While these laws don't explicitly require data to be stored in the EU, many organizations in highly regulated sectors (including government, finance, and healthcare) still view sovereign environments as the only way to eliminate jurisdictional risk entirely.

While the EU is one of the most prominent examples of a region enacting data privacy and sovereignty requirements, many other countries and regions have implemented data sovereignty frameworks that impact where government agencies and local organizations in regulated sectors can practically store data. For example:

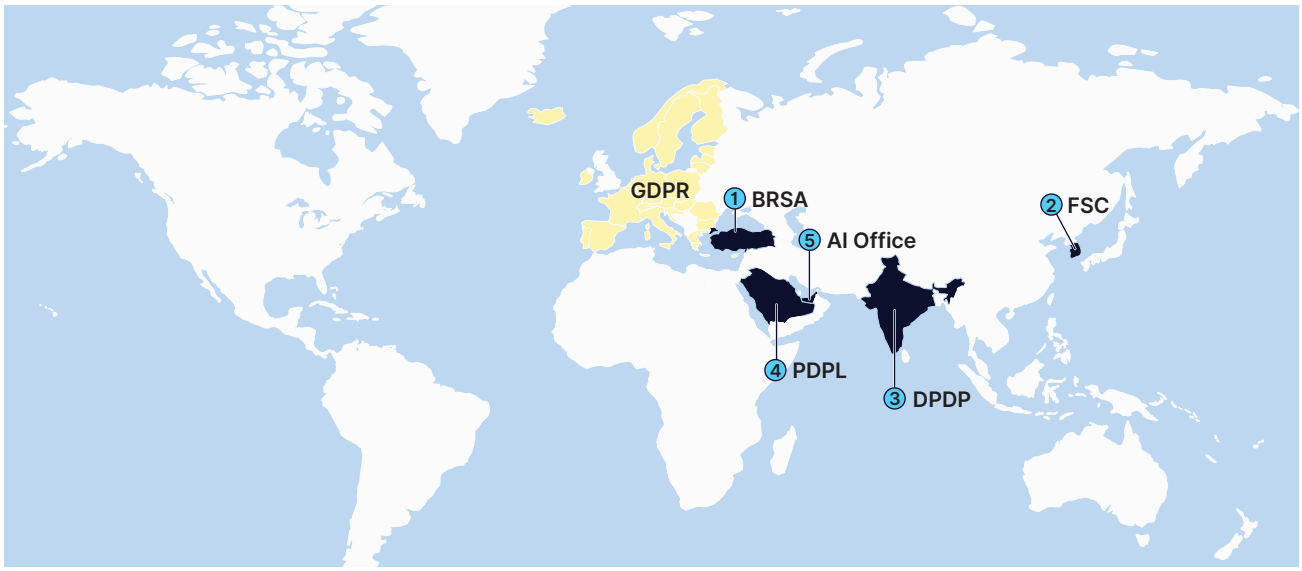

1 **Turkey's Banking Regulation and Supervision Agency (BRSA)**	Has long enforced a "closed-loop" system for banking where all "primary and secondary systems" (production servers and disaster recovery sites) used by banks must be physically located within Turkey.
2 **South Korea's Financial Services Commission (FSC)**	Allows the use of SaaS and Generative AI within internal networks but requires that any AI training or SaaS processing involving personal data must still occur through private cloud infrastructure with servers located in Korea.
3 **India's Digital Personal Data Protection (DPDP)**	This framework, while still in its early stages, establishes a rights-based, consent-driven framework governing how organizations collect, store, and protect personal data. When it comes to AI, the government can designate entities as "Significant Data Fiduciaries" and may specify categories of personal data that must be processed exclusively within India.
4 **Saudi Arabia's Personal Data Protection Law (PDPL)**	Mandates that strategically or nationally sensitive data be stored within the Kingdom. The Kingdom is making significant investments in local AI infrastructure to create compliant local alternatives.
5 **The UAE's AI Office**	Hasn't set hard requirements for data sovereignty, but it provides strategic direction at the federal level and has established incentives for local AI infrastructure that keeps data inside the UAE.

For support teams in regulated industries around the world, data sovereignty requirements are increasingly likely to impact AI usage. If you're not subject to AI data sovereignty requirements now, there's a high chance you will be in the next few years–and implementing AI-powered technology that runs in a sovereign environment now may be the best way to protect your organization from future risk.

How data sovereignty impacts the help desk buying decision

The intent of data sovereignty laws is to keep customer, personal, or enterprise data within certain legal or geographic boundaries, protecting it from foreign jurisdictional reach. That means data must be stored–and in some cases processed–by servers and/or cloud infrastructure in your region, and may only leave those boundaries under specific legal conditions or safeguards. Additionally, the cloud infrastructure must be owned and operated by local vendors to eliminate the risk of a foreign vendor potentially being compelled to turn over the data stored in their cloud.

Unfortunately, most cloud-based SaaS support platforms create a sovereignty problem by default because the software runs on a public cloud operated by one of the hyperscalers: most commonly AWS, Microsoft Azure, or Google Cloud. Data that is processed by the software—such as support tickets, customer or employee PII, and bank account information-–is stored and processed in data centers chosen by the SaaS vendor and operated by the hyperscaler, which may be located outside your region and subject to different laws.

Real-world scenario: Let's say your SaaS help desk provider runs their multi-tenant software on Google Cloud. You are based in Sweden and must comply with GDPR, which requires adequate protections for any EU personal data transferred outside the European Economic Area: protections that may not be in place under a standard SaaS arrangement. Because Google is a US-based company, it is subject to the US CLOUD Act, meaning US law enforcement can compel them to produce data regardless of where it is physically stored. This creates a direct conflict: complying with a CLOUD Act order could require your help desk provider to transfer data in ways that violate GDPR, leaving your organization exposed to regulatory risk.

When AI enters the picture: The bigger sovereignty challenge

Beyond the sovereignty challenge created by SaaS applications in general, there's also the AI piece of the puzzle. Even if you're using a help desk that runs on a regional cloud, operated by a local subsidiary of one of the hyperscalers, its AI features may not be sovereign-compliant. Help desk vendors typically power their AI features using large language models (LLMs) hosted on public cloud infrastructure, often in the US or other location that is separate from the regional infrastructure hosting the help desk itself. So even if the help desk stores and processes data in-region, its AI features may be routing personal data outside your sovereign boundaries for processing.

Real-world scenario: Imagine your organization is based in Spain, and your help desk provider runs an instance of their SaaS platform in the Spanish region, enabling you to stay compliant up until now. However, the help desk provider now offers AI features, such as a chatbot and suggested responses for agents, built on OpenAI's GPT-5. The OpenAI GPT-5 infrastructure is not located in the Spanish region, meaning you can't use the AI features in your help desk platform without violating data sovereignty requirements.

The cost of getting this wrong

So what does all of this mean to you in practical terms for support leaders? If you implement a help desk that doesn't meet your organization's sovereign requirements, potential consequences could include:

- Fines and legal violations
- Reputational damage
- Time and resources wasted evaluating software you can't use
- Delayed rollout of your help desk, creating service disruptions
- Inability to use AI features to improve support operations

The best thing you can do is to keep data sovereignty requirements front and center from the beginning of your help desk search. Partner with security, legal, and compliance team members to make sure you understand the requirements your software and its AI features are subject to and determine which help desk vendors and deployment models meet your needs.

In the next section, we'll walk you through the four primary paths to maintaining data sovereignty with a help desk:

- Hyperscaler sovereign cloud
- Virtual private cloud (VPC)
- Regional private cloud
- Fully private (on-premise or colocation) deployment

4 paths to deploying your help desk in a sovereign environment

Hyperscaler sovereign cloud

With this model, a hyperscaler operates the sovereign cloud data centers within a sovereign region, so data stays in your country or region. For example, the Microsoft Sovereign Cloud adds a governance layer to the Microsoft public cloud and private cloud environments, allowing highly regulated organizations to use Azure services while maintaining control of data residency.

This allows countries and organizations to potentially meet sovereign standards quickly while benefiting from the expertise and innovation of the hyperscaler. It can also give organizations access to secure AI through managed services like Amazon Bedrock, the Azure OpenAI Service, or Google Vertex.

However, it doesn't always fully solve the data protection problem that's driving many regulated organizations to look to sovereign clouds in the first place: the hyperscaler has legal and operational control over the cloud infrastructure, and, in most cases, they are still subject to their country's laws around data retrieval.

One exception: the AWS EU Sovereign Cloud was specifically built to avoid these challenges. It operates under a parent company that is locally controlled in the EU, led by EU citizens, and subject to local laws.

Pros	Cons
✓ Scalability and disaster recovery supported by the hyperscaler	× Risk of vendor lock-in (dependence on provider's proprietary AI and software ecosystems)
✓ Fast access to the latest AI and other hyperscaler innovations	× Hyperscaler maintains operational and legal control over physical infrastructure
✓ Potentially faster to deploy than other sovereign options (if a hyperscaler already offers a sovereign cloud in your region)	× Hyperscaler may be subject to their country's laws around data retrieval
✓ Lower upfront cost than fully private deployment	
✓ Access to specialized architectures like Microsoft's Sovereign Landing Zones, providing automated guardrails for compliance	

Virtual Private Cloud (VPC)

A virtual private cloud is a logically isolated network environment hosted within a hyperscaler's broader public cloud infrastructure. Unlike a standard multi-tenant SaaS deployment, a single-tenant deployment in a VPC gives your organization a dedicated slice of the hyperscaler's infrastructure, with controls over where your data is stored and who can access it. When you're contracting for a VPC, you can specify that you want it to be located in your sovereign region. This makes it a step up from a standard SaaS arrangement, and for some organizations, it may be sufficient to meet their compliance requirements.

However, it's important to understand what a VPC does and doesn't solve. Because a VPC is still hosted on a hyperscaler's public cloud infrastructure, the hyperscaler retains operational and legal control over the underlying hardware—meaning they remain subject to their home country's laws around data access and retrieval. A VPC deployment with AWS, for example, doesn't eliminate CLOUD Act exposure the way an AWS EU Sovereign Cloud deployment can. If your organization operates in a region with strict data sovereignty requirements, or in a highly regulated industry, a VPC alone may not be enough to satisfy your legal and compliance obligations—particularly once AI features enter the picture, since AI processing may route data outside the VPC boundary entirely.

Pros	Cons
✓ Lower cost and faster to deploy than other sovereign options	× Hyperscaler remains subject to their home country's laws, creating jurisdictional risk
✓ Available in a wide range of geographic regions	× Logical isolation only: infrastructure is not physically dedicated or regionally governed
✓ Scales easily with organizational needs	× Vendor lock-in risk
✓ Shared Responsibility Model for security, with the hyperscaler managing cloud security only	× AI services accessed through the VPC may route data outside the VPC boundary
✓ Compatible with managed AI services offered by the hyperscaler	

Regional private cloud

An alternative approach that is adopted by some organizations, including some financial institutions and government agencies, is to deploy their tech stack in a private cloud operated by a regional vendor in their sovereign territory.

This option lets organizations tailor their cloud usage to their compliance requirements because they can choose where they want to physically store their data. However, this model can be difficult to scale due to limited on-premise resources, and it requires organizations to rely on the private cloud service provider for maintenance and ongoing innovation. It also requires the private cloud vendors to make the significant investment of building their own AI infrastructure to provide AI foundation model services within the sovereign environment.

Pros	Cons
✓ Regional vendor operates under the laws of the sovereign territory	× Difficult to scale due to limited regional server capacity
✓ High degree of customization for specific industry regulations	× Lower R&D budget and level of cybersecurity expertise compared to hyperscalers
✓ Regional vendor handles hardware upkeep and maintenance	× Requires vendor to build and manage AI infrastructure

Fully private (on-premise or colocation) deployment

For support teams in the most highly regulated industries, there's a third path to maintaining data sovereignty–deploying your support software on premise in:

1. A data center that your organization owns and operates, OR
2. A colocation center in which your organization rents servers and infrastructure

Due to the complexity and cost of managing data centers and their infrastructure, this will likely only be your path if your organization has its own data center(s) or colocation facility already. However, if that's the case, a fully private, on-premise deployment can be the best path to maintaining full control of your data.

To make this option work, you'll need to look for support software vendors that offer fully self-hosted deployment options. If you plan to use AI features in your software, you'll also need to determine if you can connect AI models that your security and compliance teams have already vetted. You may find your organization requires you to use an AI foundation model that can be deployed fully on-premise.

One increasingly popular option for larger enterprises is to adopt an Enterprise Private Cloud, where a hyperscaler will manage a "mini-cloud" with hardware and services running directly in your own data center. Examples include AWS Outposts, Azure Local, and Google Distributed Cloud (GDC).

Pros	Cons
✓ Organization maintains full control of data and (in most cases) hardware	× Highest cost and complexity
✓ Potential for air-gapped (fully offline) environment for the most sensitive security use cases	× Limited and slower scalability
✓ Minimal third-party risk	× Organization is responsible for hardware maintenance and security
✓ Can be optimized for high-performance local AI models	× Requires significant internal IT resources to maintain

The next decision: AI implementation

If your support team is planning to use any of the AI features your help desk vendor offers, you'll need to think about their approach to AI in addition to the deployment models they support. Essentially, there are two paths you can take with AI: use an AI model your vendor provides or connect to a model your organization is already using.

Vendor-provided AI

With this approach, you use the LLMs your vendor has selected to power their AI features. Popular open-source LLMs include:

- OpenAI's ChatGPT
- Anthropic's Claude
- Google's Gemini

If your vendor uses AI models hosted in a public cloud, this approach may not meet your sovereign requirements. This is because data may need to leave your sovereign boundaries for AI processing, even if the help desk itself is hosted in your region. If that's the case, you may find that your organization won't allow you to use the AI features in your help desk–causing you to miss out on efficiency gains and improvements to the support agent experience.

Bring-your-own AI

With a BYO-AI approach, you choose the models you want to power your help desk's AI features. Your help desk vendor connects the model via an API key you supply from the model provider.

Providing your own AI lets you select models your security team has already vetted. And, crucially for support teams with sovereign requirements, it allows you to use AI models that can be deployed in your private environment or a sovereign cloud. Examples of LLMs with private and sovereign deployment options include:

- Meta's LLaMA
- DeepSeek
- Cohere
- Mistral AI
- Qwen

BYO-AI isn't an option offered by every help desk provider, so it's important to ask vendors about this up front if you know it's non-negotiable for your organization.

What to look for in a sovereign-ready help desk vendor

As you start evaluating sovereign-ready help desk vendors, there are three crucial areas to keep in mind: deployment flexibility, AI sovereignty, and infrastructure management.

Deployment flexibility

You'll need to ensure your chosen provider offers the deployment model that works for your organizations. Many of the major SaaS help desk providers are cloud-only, or offer a hyperscaler sovereign cloud deployment but only in certain regions, meaning you may not be able to use them in your organization.

Questions to ask:

- What regions or countries do you currently support for sovereign deployments?
- What is your roadmap for expanding to new regions?
- Where is my organization's data stored and processed?
- Can you deploy your help desk platform in our environment

AI sovereignty

If your chosen provider offers a deployment model that works for you, do they also offer AI sovereignty? If not, you may end up being blocked from using any of their AI features.

Questions to ask:

- Can we choose our AI foundation model, or are we locked into your chosen provider?
- If we bring our own AI model, can we deploy it fully on-premise or in our private environment?
- If we're using your AI model, does AI processing happen within the same sovereign boundary as our data storage?
- Do your AI features log or retain prompt and response data, and if so, where is that data stored?

Infrastructure management

You'll need to determine who is responsible for infrastructure management and, if you're choosing a self-hosted deployment, whether the vendor offers professional services to assist with this deployment.

Questions to ask:

- Who manages the data storage infrastructure?
- Can you or your partners manage a self-hosted version of the platform for me?
- Do you help set up self-hosted instances of your platform?
- What does your standard update and patch process look like for self-hosted deployments?

A practical comparison: Which leading help desk vendors meet your data sovereignty needs?

There are a lot of help desk platforms on the market, but how many of them actually meet your data sovereignty requirements? The table below breaks down the deployment models and BYO-AI offerings of some of the top providers. While this isn't an exhaustive list, we hope it's a useful starting place.

Key

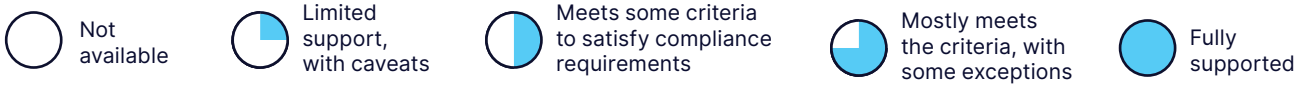

Vendor	Hyperscaler Sovereign Cloud	VPC	Regional Private Cloud	On-Premise Deployment	Enterprise Private Cloud	Bring-Your Own-AI Model
Zendesk	1					
Freshworks (Freshdesk)	1					
Salesforce		2				3
Jira Service Manager (Cloud)	4					5
Jira Service Manager (Data Center)[6]						
ServiceNow				7		
ManageEngine						8
Deskpro						

1 Limited options to choose the region where "data at rest" is stored.

2 Salesforce Private Connect creates a private connection between Salesforce and your AWS VPC, and Mulesoft (a Salesforce subsidiary) offers a true VPC.

3 Not all Salesforce AI features work with their BYO-AI model.

4 Atlassian Isolated Cloud enables customers to choose their region for data hosting but is still managed by Atlassian (Australian-based company).

5 Option to use Atlassian-hosted LLMs only.

6 Jira began sunsetting all on-premise products in 2026.

7 While ServiceNow does technically maintain a self-hosted software license agreement, they typically only grant these licenses to a very small subset of customers.

8 Private AI is only available with Zia, ManageEngine/Zoho's proprietary model.

Deskpro Private: Built for regulated environments

Most help desk platforms take a cloud-first approach. Those that do have sovereign deployment offerings often require you to make trade-offs. They might offer a sovereign cloud setup through a hyperscaler, but not one that is able to avoid the data governance challenges posed by the US CLOUD Act. You might be able to deploy in a hyperscaler sovereign cloud that's compliant for your organization but not bring your own AI model, preventing you from using the help desk's AI features. Or a provider may have private cloud partnerships, but not in your region.

With Deskpro Private, you don't have to make those tradeoffs.

Deskpro Private is an AI-powered help desk platform designed for organizations that require full control over security, data, and deployment, without sacrificing the benefits of a modern customer support platform. We built it to meet the needs of highly regulated industries and organizations that have strict data privacy and sovereignty requirements, including:

- Government agencies
- Healthcare organizations
- Financial services firms
- Aerospace and defense contractors
- Multinational enterprises

Deskpro Private lets you choose the sovereign-ready deployment model that's best for your organization: in a hyperscaler sovereign cloud, VPC, regional private cloud, or on-premise in your data center or a colocation facility.

You can also connect your AI model of choice to ensure you're using an LLM that processes data in your sovereign region and avoids the risk of foreign jurisdictional reach. Here's how it works at a high level:

1. Deploy Deskpro in your private environment.
2. Connect to LLMs such as GPT-5, Claude, Gemini, Mistral, LLaMa, or your own private models.
3. Ground Deskpro's AI features in your organization's data and knowledge, including help center articles, product documents, and resolved tickets. This enables our AI tools to generate context-aware, organization-specific outputs, and sensitive data never leaves your security perimeter.
4. Securely automate workflows and support operations with AI. Capabilities include AI-suggested responses for agents, intelligent ticket triage, ticket summarization, help center article generation, and chat.

With Deskpro Private, you get the benefits of AI and modern help desk capabilities without compromising sovereign compliance. Your agents get more time to focus on complex issues while automating repetitive tasks. The employees, clients, partners, and customers that rely on your support get faster, more accurate responses. And your organization enhances its support operations while maintaining complete data control.

Learn more at **deskpro.com/private**
When you're ready, book a demo: **deskpro.com/demo**

Deskpro

About Deskpro

Deskpro is the developer of AI-driven help desk software that empowers organizations worldwide to deliver exceptional customer and employee experiences that translate directly into tangible ROI, through improved customer and agent retention. Deskpro is the only help desk platform with enterprise-grade security, compliance, and data privacy, enabling all organizations to adopt and deploy AI, even in the most regulated industries and sovereign environments. Deskpro is available as a cloud service, and can also be deployed in VPCs, private clouds, sovereign clouds, and on-premise. Trusted by leading enterprises worldwide in banking, technology, financial services, healthcare systems, aerospace and defense, as well as government agencies, Deskpro is led by experts in enterprise software, customer experience, and customer support.

Learn more, start a free trial, or book a customized demo at **deskpro.com**.