Transcend

Building an In-house Privacy Request System

Tactical, expert guidance on the processes, people, and steps involved in **architecting a future-proof in-house privacy system.**

Transcend

What we'll cover in this guide:

01 Introduction

Thanks to regulatory requirements created by Europe's GDPR and California's CCPA, many companies are investing in automated backend workflows to efficiently automate, process, and respond to data access and erasure requests as well as opt-outs.

Why? Whether your organization receives 1, 100, or 1,000 privacy requests every month, manual efforts require resources and must be understood in the context of team priorities, human-error failure points, and hidden costs.

So it's no surprise that organizations with enough engineering resources and available staff are eager to build automated workflows that reduce potential errors and avoid disrupting roadmaps.

In this guide, we provide a breakdown of the essential elements to build an automated privacy request workflow, with advice from our experts who build these systems for a variety of multinational companies, from fintech to consumer lifestyle to developer platforms.

But there's a crucial step before you pick up tools, and that's knowing the intricacies of what's required. To aid you, we've engaged experts on privacy request infrastructure to develop **six key questions you should have answers for before you start** to guide your cross-functional conversations and help you understand the true scope of the project.

> **Doing this well can provide a reliable and flexible platform to support future changes in laws or your technology stack.**

But be warned, it requires the ability and visibility to effectively anticipate and manage complexities in your data environment. If you choose to build internally vs. leveraging a trusted privacy engineering partner, prepare and plan for complexity. On top of that, the variation and fragmentation among different privacy laws mean your infrastructure needs to remain flexible and nimble for the long term.

02 The 6 questions to answer before you start

Like gathering ingredients before trying a new recipe, you need to first understand which information, people, and tools you'll need to be successful.

Our experts have years of experience implementing these systems for our customers, so here are their six key questions you should be able to answer, and your cross-functional colleagues should be in alignment on before you start building.

QUESTION 1

What kind of data subject(s) do you have?

Different types of data subjects can be involved in various processing activities and often require different workflows for data rights accessibility. Common categories of data subjects include *users*: people who engage with your organization or product by logging in to a dedicated account or portal, and *non-users*: people whose engagement with your organization is not tied to a specific user account.

The type of data subject matters for a few reasons. Most importantly, it determines the possible entry points and verification options available to data subjects for exercising their data rights. Users, for example, may be able to request a copy of their data or schedule a deletion through a settings panel when they are logged into their account. Having a login credential also provides a helpful verification mechanism to ensure the person requesting access or deletion is the proper data owner.

Verification flows for non-users can be more challenging because they cannot authenticate through a traditional login process using a password. Both California and the EU law require organizations to verify someone's identity before complying with their request for a copy of their personal information or to delete their information. The key is to use methods of authentication and verification that are legal, effective, and easy for consumers to complete.

Additionally, if your verification process isn't automated yet, someone will have to do the verification manually. Plan your resources and priorities accordingly to ensure they have enough training and time to process verifications quickly and accurately.

QUESTION 2

Who needs to be involved in the design and implementation process?

Like many things in privacy, building infrastructure that enables someone to exercise their data rights **is a multi-disciplinary effort.** Before you make any concrete decisions about how to build your workflows, engage with all the stakeholders who will need to interact with your system in building and maintaining it.

Be mindful that taking on a project of this size can easily disrupt roadmaps for partner teams, so being flexible on delivery times and project milestones will go a long way to getting the resources you need from other organizations. You may have to compromise for a delayed start or finish date in order to accommodate other goals these teams need to achieve for the business.

Some of the most common stakeholders involved are **engineers and UX designers** who can align backend specifications with a seamless frontend experience for your users. You'll also need the help of whichever functions have visibility and authority over your data stores.

Obviously, **legal counsel** needs to be involved, not only for interpreting the regulatory requirements that apply to your organization, but they will be internal users of your privacy infrastructure, so it needs to be built in a way that's easy for them to use. Another important stakeholder is **your cybersecurity team** or whoever will manage the authentication protocols mentioned earlier. Failing to verify the identity of a requester could result in personal information being shared with the wrong person or being deleted by an unauthorized request.

Organizational structures may differ inside your business, so it's critical to proactively map out all the stakeholders who will be affected, either because they're needed for constructing the workflows or because they will be using them.

QUESTION 3

Where is your data stored?

One of the reasons cross-functional collaboration is so important when building privacy infrastructure is that no one has complete visibility or control over all the data inside a company. You may need to work with multiple data science or platform teams to locate where personal data exists in your environment. Datastores can consist of application databases like MySQL, data lakes such as Amazon S3 or Hadoop, and data warehouses like Snowflake.

The architecture of your datastores also matters. Smaller organizations with a monolithic infrastructure can follow a similar approach for privacy workflows, but this is less common. Today, it's more common for companies, especially tech companies, to use a distributed architecture with dozens, if not thousands, of microservices. With this approach, each service or system owns its own functionality, including privacy workflows, which need to be triggered by the backbone privacy infrastructure you're building.

Often, Kafka (or a custom queue system) spins up jobs for all microservices and confirms when the task is complete, or a more rudimentary architecture will use an additional notification system to confirm task completion. Exception handling is also critical if you're working with a microservices architecture.

Because data privacy laws dictate precise deadlines for responding to and fulfilling data subject requests, it's important that privacy jobs are automatically rescheduled for any microservice that's down. Map out the logic and rules for exception handling in advance to ensure all the relevant stakeholders are aligned on how they will be handled. For example, decide what the fault tolerance will be for error states and figure out how you will alert significant delays to legal stakeholders.

QUESTION 4

How is data used by your business?

Duplicate datastores, like analytic warehouses, are commonly used by analysts in combination with business intelligence tools to understand how your products are used by your customers. Deleting or otherwise manipulating that data could significantly impact the reliability of data for analysis, so coordinate with them and legal in advance to determine how you will treat data in duplicate datastores to comply with regulatory requirements. Where the law permits, metadata retention can help maintain the integrity of statistical analysis. However, those warehouses rarely have custom code connected to them. So, you have to stand up a new service to talk to the warehouse.

QUESTION 5

What vendor integrations do you have?

In addition to mapping out the location of personal data inside your own environment, identifying any vendor and third-party integrations before building new privacy infrastructure will save you a world of hurt. You need to know where personal data is exchanged with other organizations because, as a data controller, you're responsible for extending data access, deletion, and some opt-out requests to third-party systems where you store or share personal information.

Doing this early matters because vendors may not have an API that you can easily plug into your own workflow. They may require an email process, and knowing that upfront makes a big difference when you get to setting it all up later on.

QUESTION 6

How will you handle sensitive functions and product considerations for erasures?

One more pre-engineering task is to map out how you will run the necessary legal checks, product logic, and account states that may complicate your workflow. For example, if your product includes a group chat function, how will you handle requests from individual participants to delete that content? What are the retention requirements for data involved in safety or fraud investigations? How will you address deletion requests for accounts with an outstanding balance or credit? Will you automatically delete accounts after a specified period of inactivity?

03 The Setup

Once you have all the details ironed out above, it's time to start building. **The critical element to success here is a multi-disciplinary approach.** As we outlined above, this isn't just a legal or engineering challenge—it's also about the user experience and the ability to provably demonstrate compliance with regulatory requirements.

01 — Data workflows

The more comprehensive and complete your pre-setup work is, as outlined above, the better off you'll be when it comes to implementation. Doing so will reduce the volume of potential surprises for your engineering team, thus minimizing the volume of design changes and rework needed throughout the process.

With all the pre-work complete, you're ready to schedule integration building with internal and third-party systems. And you'll already know which ones don't have an API, so you can anticipate the specific integrations that will need an email process instead.

Most importantly, don't forget who will be using this tool on a day-to-day basis internally, usually legal. Give them admin rights to review and audit data subject requests as well as alerts of any delays in workflow jobs that could impact compliance with regulatory deadlines or completeness. Legal also needs discretion over the transactional email service and content that deliver the final report to data subjects. Building those capabilities into the design of your privacy infrastructure saves engineers a lot of time down the road as legal requirements and customer expectations change with regard to content and presentation.

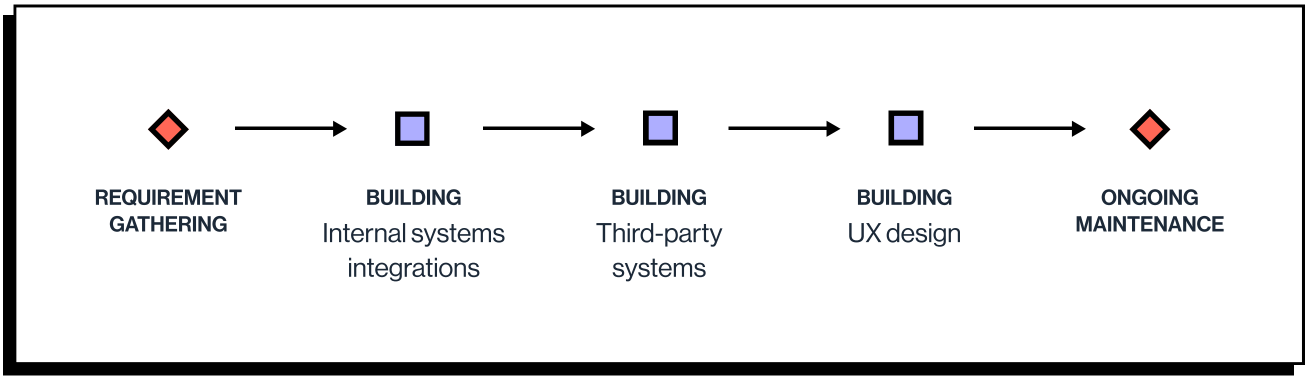

REQUIREMENT GATHERING → BUILDING Internal systems integrations → BUILDING Third-party systems → BUILDING UX design → ONGOING MAINTENANCE

02 — UX design

Designing the settings page or entry point for submitting a data rights request addresses two key priorities for any privacy initiative:

PRIORITY 1

Is it easy to use and thus easy to demonstrate compliance with legal obligations?

PRIORITY 2

What does it communicate about where privacy ranks in your company priorities?

An easy-to-understand, well-designed user interface helps you build trust with data subjects throughout every step of the process.

Additionally, our experts recommend applying the "user experience challenge" as you build out your infrastructure and controls. How easy is it for data subjects to find and use the tools you've built? In particular, take note of how much time it takes to complete your verification checks and fulfill a request. Ask yourself: "if a cross-functional team adds a new vendor or switches a product feature, is it easy or a nightmare to ensure consistency of user experience on data privacy?"

03 — Ongoing maintenance

It's crucial to build a budget and allowance for consistent, ongoing maintenance and adaptation. Pulling together the necessary components upfront and following best practices on the setup are critical for building a flexible, easy-to-manage system for legal and technical teams. This is another common bottleneck when it comes to roadmap disruptions and coordinating resources with other stakeholders. Long-term maintenance of in-house systems is resource-intensive, especially when it comes to patches or updates.

Keeping your newly minted data rights workflows up-to-date requires the flexibility to quickly and easily adjust to changes in your infrastructure or SaaS stack, as well as legal requirements. For example, if legal adds new retention requirements or the marketing department starts advertising on a new network, engineering functionality needs to be updated for all microservices.

The ideal state is for microservice functionality to be flexible enough that legal can update the workflow without needing new code deployments. Can they tweak reports and content without updating code? Building to this specification will make ongoing maintenance much easier and reduce the demand on engineers for continuous code updates.

Finally, be mindful of ongoing changes from third parties such as new data limitation or usage features or new API functionality. If you build your own in-house infrastructure, your engineers will need to keep up with those changes in order to manage integrations with each third party's dynamic frontend code.

04

Pro Tips

After years of building backend infrastructure with our customers, we've learned a lot about how even well-intentioned companies can get in their own way. Here are additional tips from our experts to help enhance your automated workflows and avoid costly mistakes.

01 — Build confirmation notifications and a grace period for erasures

Communication with data subjects is key throughout the entire process. Let them know what to expect and when the status of their request changes. Also, grace periods help protect against unauthorized requests in cases where a user account is compromised by a bad actor.

02— Don't use probabilistic or search-based deletion with automated erasure

When it comes to data deletion, making broad strokes creates a lot of problems. Instead, you want your privacy infrastructure to enact precise and accurate cuts.

03 — Don’t forget about advertising identifiers

These live on the user browser, so you have to extract the unique ID from your frontend to trigger action from your backend infrastructure.

04— Avoid regional workflows

When possible, universal workflows help reduce operational complexity by applying consistency. If not—since different laws and mandated data rights will inevitably conflict—be prepared to build separate backbones to run different, simultaneous workflows.

05 Summing Up

Increasingly, a modern, engineered data privacy request infrastructure isn't just nice-to-have for companies drowning in consumer privacy requests—it's becoming a crucial tool for companies needing to prioritize risk reduction, spend efficiency, compliance, and user experience.

Building scalable privacy infrastructure is an area **Transcend has deep expertise in**, and so, for those that choose to build internally vs. leveraging the support of a trusted privacy partner, our advice would be simple: the task is one that requires deep collaboration across teams and will require dedicated and ongoing resourcing.

You'll need to keep track of the evolution of your data environment, the various connections across systems, and the governance layer atop your systems that will need to respond to the evolving legislative environment in the long term.

Either way—our team of experts **is always available** to offer advice on the best path forward and advice along the way.

This Privacy Playbook is brought to you by Transcend.

At Transcend, we're big believers in investing in engineered solutions to privacy's greatest challenges.

We provide an enterprise-grade, full-stack solution to receive, manage, and automatically fulfill privacy requests from your users—trusted by leading brands including Patreon, Indiegogo, Robinhood, and more.

Free up engineering resources, meet—exceed—regulatory obligations with auditable automation, and build brand trust through respectful data and consent transparency and control.

Speak to our solutions team »

Learn more at transcend.io »

Transcend

Transcend