The future is so much bigger than the past

- Sir Tim Berners-Lee

CONFIDENTIAL

GDPR and Solid

Overview

Solid is a web standards-based protocol that was designed to enable users to store their data in Pods and to control if, when, and with whom that data is shared. Solid provides the tools for organizations to comply with the GDPR effectively and elegantly. In a number of key respects, Solid also enables organizations to fundamentally exceed the GDPR's requirements and more meaningfully achieve the values at its core: true transparency, user control, and trust. This whitepaper outlines how organizations can use Solid to both meet and exceed GDPR standards.

Key Takeaways

1. Solid enables organizations to comply fully with the current standards and expectations of the GDPR. Moreover, by fundamentally restructuring the manner in which data can be controlled, Solid can enable organizations to significantly exceed these standards, moving beyond the GDPR's specific requirements to better fulfill the values at its core: transparency, trust and control.

2. Using Solid, organizations can enable users to exercise their data subject rights directly by seeing and acting on the data inside their Pod, enhancing the user rights model by providing real transparency and user control. Organizations using Solid can rely on any of the GDPR's lawful bases to process data, but in particular can obtain a user's clear, informed and specific consent in a way that integrates with the functionality of a product or service. Solid provides organizations with the tools to pinpoint the data they need and be clear about the purposes for which they are collected, and to significantly minimize the data that is collected and stored across multiple organizations because a single Solid Pod can be used to share data with a range of different companies or governments.

Introduction

The General Data Protection Regulation (GDPR) is a European law that sets out a range of legal requirements governing the way organizations can collect, process, store and transfer personal data. Although there are a large number of global privacy laws in place around the world, and multiple new privacy laws are being passed every year, the GDPR has become the de facto global privacy standard since it came into force in 2018.

The GDPR prescribes a number of significant compliance requirements for organizations that process personal data, ranging from broad principles like transparency and data minimization, to minimum standards for consent, to documentation requirements. The fear of being hit with a substantial penalty for non-compliance, and perhaps also of the negative publicity that an adverse finding can bring, has led most large companies to take data privacy and security more seriously. In practical terms, the GDPR has required most large companies to invest a significant amount of money and time into privacy lawyers and consultants, who in turn have expanded and standardized privacy policies, added privacy disclaimers and disclosures to in-product user flows, and implemented various back-end documentation processes and policies. While the GDPR has forced most organizations to take their privacy responsibilities more seriously, and to demonstrate increased diligence in relation to the data they collect, the tangible improvements to users have been surprisingly limited.

While most companies have sought to be legally compliant and to avoid the attention of privacy regulators, there is a significant gap between the GDPR's conceptual requirements and the ability of organizations to implement these requirements in a useful, practical way for users. Without the technology or tools to build for better privacy, most organizations have left users with extremely limited transparency and even less control over how their data gets used. At present, real world privacy compliance still relies chiefly on long general descriptions of how and why companies process data -- in the form of Privacy Policies and Terms of Service -- and requires users to accept these conditions in order to use the service. While compliance with the GDPR's back-end documentation requirements may help to improve organizations' diligence in relation to the data they use, they do little to improve transparency or understanding for users. In simple terms, users lack the tools to easily see or understand how their data is used within an organization or across multiple organizations.

If privacy is to be meaningfully improved for users, the large but necessary next step is for engineers to actually start building good privacy into the ways that users, and their data, interact with companies. What is needed, in other words, is a new and improved paradigm. But data privacy laws have established valuable and enforceable common standards and still have a very important role to play in the future of privacy. As

a result, the central challenge we now face is to build new privacy solutions that are consistent with laws like the GDPR but can also move meaningfully beyond them.

Enter: Solid

Solid is a web-standards-based protocol that enables users to store their data in Pods and then control if, when, and with whom that data is shared. Solid thus enables a fundamentally new framework within which users and companies can interact. Instead of every company individually collecting, storing and using its own copy of a user's data, Solid enables a user's data to be centralized in a Pod that they can see and control. With Solid, a user can be empowered to decide whether to share data with a particular application and which data should be shared.

As we will see below, Solid enables organizations to comply fully with the current standards and expectations of the GDPR. But by fundamentally restructuring the manner in which data can be controlled, Solid can also enable organizations to significantly exceed these standards, moving beyond the GDPR's specific requirements to better fulfill the values at its core: transparency, trust and control.

Figure 1: Solid and GDPR Compliance

GDPR Requirements	Does Solid Enable GDPR Compliance?
Lawful Basis of Processing (and proper user consent)	Yes
Data Subject Rights	Yes
Keeping Records of Processing Activities	Yes
Requirements for Profiling and Automated Decision Making	Yes
Requirements for Third Party Processors	Yes
Data Minimization	Yes
Data Retention	Yes

In the sections below, we will assess the manner in which Solid enables organizations to comply with the GDPR and then examine the ways in which it enables them to exceed the current approach.

1. Consent and "Lawful Basis for Processing"

One of the fundamental tenets of the GDPR is the requirement that any time an organization processes personal data, it requires a lawful basis to do that processing. In simple terms, this means that for each use of a customer's data, the organization must be able to identify, and communicate, which of the six recognized legal bases it is relying on to do this. The six lawful bases that the GDPR recognizes as valid are the following:

1. Consent has been obtained from the data subject
2. The processing is necessary for the performance of a contract
3. The processing is necessary for compliance with a legal obligation
4. The processing is necessary to protect vital interests of a data subject or third party
5. The processing is necessary for performance of a task carried out in the public interest
6. The processing is necessary for the "legitimate interest" of the data controller, and is not overridden by the interests or fundamental rights and freedoms of the data subject.

Solid Meets GDPR Standards

Among the most commonly championed aspects of the GDPR is that it clearly sets out expanded requirements for what constitutes a valid consent: most notably that consent must be specific, informed, and conveyed by a clear affirmative action. Consent cannot, and should not, be inferred. Nor should the consent for various different acts of processing be bundled together.

Solid is able to meet these standards by enabling organizations to base potential processing actions on clear and specific affirmative consent as provided by users. The ability to categorize the data elements being processed, and to customize how consent is obtained, also enables organizations to process sensitive data.

Solid Can Meaningfully Exceed GDPR Standards

While the GDPR sets out high standards for consent, it also recognizes five alternative legal bases that companies can use instead to process personal data. Given that consent is often difficult for organizations to obtain in the context of a particular user flow -- and that asking for consent always comes with the possibility that a user may say no, or later revoke their consent -- many companies, and most large tech companies, purposely avoid relying on consent wherever possible.

To the detriment of user transparency, understanding and choice, the remaining legal bases under which data can be processed are not user-facing -- they are decided in the background, typically with little more than a general note in an organization's privacy policy that certain broad types of data are processed under these legal bases. In particular, the ability of companies to rely on a "legitimate interest" -- which they themselves determine is not overridden by the interests or fundamental rights and freedoms of the user -- means that users have no choice about, and little visibility into, the ways in which much of their data is processed.

Organizations utilizing Solid could in theory also choose to rely on alternative legal bases -- they might, for example, determine that they have a legitimate interest to process certain user data for a separate particular purpose. For example, an organization could process certain data for fraud prevention purposes as a legitimate interest. Using Solid would not prevent organizations from processing data pursuant to a legitimate interest, should they choose to do so.

However, Solid optimizes the ability of organizations to process data based on consent. By storing data in a Pod controlled by its user, Solid lets organizations operationalize consent: they can enable users to choose whether to share certain data for a particular processing purpose, make the particular elements of data being shared visible to the user, and facilitate a user's ability to withdraw consent at any time. Solid can extend a user's ability to consent by actualizing it as the ability to see, understand, and control how their data is used.

2. Data Subject Rights

Instead of considering personal data to be the exclusive property of the organization possessing it, the GDPR requires that users retain certain fundamental rights in relation to their personal information. This includes:

- The Right of Access -- users have the right to request copies of their personal information.
- The Right of Rectification -- users have the right to ask an organization to correct information they believe is incorrect, or complete information they believe is incomplete.
- The Right to Erasure -- users have the right to ask an organization to erase their personal data in certain circumstances (and subject to a number of exceptions).
- The Right to Restrict Processing, or to Object to Processing - users have the right to request that organizations restrict their processing, or the right to object to processing in certain circumstances.
- The Right to Data Portability - where data was processed based on consent, or pursuant to a contract, users have the right to receive a copy of their data in a

"structured, commonly used and machine-readable format" and to provide that data to another organization.

Solid enables organizations to comply with the GDPR with respect to each of these rights, as outlined below. But the Solid protocol also enables a significantly upgraded experience for users with respect to the exercise of their data subject rights. In most cases, GDPR-compliant organizations enable users to exercise their data subject rights in ways that are difficult to initiate, slow to complete, difficult to understand, and almost impossible for users to verify. For most companies, data rights can only be triggered when a user sends an email to a particular address, requesting a specific action. Some of the larger technology companies have built custom portals for users to submit a data rights request, though the ultimate results are largely the same. In either case, the user typically receives a formulaic response and needs to verify their identity, often by providing additional and potentially more sensitive data for the verification process. If a user requests the deletion of their data, it can be difficult if not impossible to confirm that all of their data was truly deleted, or to see or understand what data was retained for permissible legal reasons. If a user requests a copy of their data, that data is seldom provided in a format that renders it usable.

Solid enables organizations to provide users with their data subject rights in a far more practical, transparent, and useful manner, because users can see the data in their Pods and exercise these functions directly. As outlined in Figure 2 below, Solid Pods can be used to fundamentally shift the user rights model to one that favors real transparency and puts control in the hands of the user.

Figure 2: Solid and Data Subject Rights

GDPR Data Subject Right	How Users Can Potentially Exercise the Right
Right of Access	Users can access their data at any time by viewing it in their Pod.
Right of Rectification	With the ability to view their data, users are in a position to see errors and to make or request corrections.
Right to Erasure	Depending on the use case, users can revoke consent for a particular sharing, delete data, or request that particular data be deleted.
Right to Restrict Processing or to Object to Processing	Users can be enabled to restrict or revoke consent for a particular data

	sharing/processing from their Pod.
Right to Data Portability	The Solid Pod model is designed for interoperability: it allows a single data Pod to interact with multiple organizations or applications while remaining under the user's control. The data stored in a Pod is thus inherently portable. This functionality represents a significant improvement to current approaches. Depending on the use case, solutions could also be developed that would output data into other machine-readable formats.

Solid could also enable a number of additional advantages for users with respect to exercising their data rights. For example, users could be empowered to pre-program or schedule the exercise of their rights -- deciding to delete particular information after 90 days, for example, or setting other rules in advance. With broader adoption, use of the Solid protocol can also provide a significant practical improvement by centralizing the ability to exercise data rights across many different organizations into a single hub to which all of those organizations connect. At present, data rights are difficult for users to exercise because they may not know, remember, or track all the different organizations that are storing and using their data.

A Note On Data Deletion and Retention

It bears noting that using Solid does not change the ability of an organization to retain data that it needs to keep for a recognized exception, such as complying with a legal obligation. In simple terms, if an organization has a valid basis to retain particular information pursuant to the GDPR then using Solid does not change this. This could be the case with respect to certain financial records, tax documents, medical records, and other information that users may not be permitted to delete.

In simple terms, it is expected that there will be certain data that cannot be deleted or changed by users. Solid Pods should therefore be configured so that users can delete, modify or revoke access only to the data they are permitted to change. For data they cannot change, they could nonetheless be permitted to see this data and understand why it must be retained in its current form. For example, data could live in the *organization's* Pod, with a link to that data visible from the user's Pod.[1] Enabling this

[1] To provide one such example, users cannot be permitted to delete or change the balance in their bank accounts, but this balance data could be made visible to them. A link in the user's Pod could be provided to their balance information, which would be stored in the organization's Pod. Alternatively, the bank could issue the user a verifiable credential that confirms the current

level of transparency can represent a significant practical improvement on current data deletion practices.

3. Data Minimization

Data minimization is a fundamental privacy principle that dictates how data should be collected and used. Under the GDPR, personal data must be adequate, relevant, and limited to what is necessary for each specific purpose of the processing. In simple terms, organizations should only collect the data they need for the particular purposes they fulfil. Any excessive data collection is prohibited.

Solid can provide organizations with the tools to practice good data minimization and to improve on their current approach. Organizations using Solid are able to specifically pinpoint the data they need and be clear about the purposes for which they are collected. Because the collection and retention of data can be made transparent to Solid users -- they can view their data in their Pod -- users can see and understand that data is collected and used for specific purposes.

More broadly, the Solid Pod model encourages large-scale data minimization *across organizations*, because a single Solid Pod can be used to store one user's data and provide access to that data to a range of different organizations. By enabling users to reuse data that has already been collected by another organization, while still exercising control over whether that data is shared, Solid can enable significant data minimization improvements.

4. Data Processors

The reality of modern data processing is that most companies make use of a range of other companies to efficiently handle the performance of various functions. These service providers typically require that an organization share certain user data with them so they can provide their services, which can include payment processing, cloud storage, fraud detection, customer service, or a broad range of other services. In these arrangements, the GDPR calls the organization that makes the decisions about the data (or, specifically, that "determines the purposes and means of the processing of personal data") the data controller. The service provider with which the organization shares data, to help it perform one or more tasks, is called a data processor.

Where an organization intends to share data with a processor, it can only do so by ensuring that its obligations to protect data are flowed to the third party via a contract.

---

balance and store that credential in the user's Pod. This way, if a party needed to rely on that data, it could verify the credential to ensure authenticity. Approaches can vary based on the use case.

The contract between a controller and a processor ensures that the processor only processes data in accordance with the controller's instructions, maintains the controller's GDPR compliance, and protects the privacy and security of the data. The contract must also ensure that the processor cannot further contract out work to a subprocessor without the authorization to do so by the controller.

These obligations in relation to processors exist whether or not an organization uses Solid. In narrow terms, the use of Solid does not change the ability of an organization to use processors or its GDPR obligations when it does so.

However, Solid provides organizations with the opportunity to provide improved transparency and functionality for users in relation to data processors. At present, most organizations provide little transparency with respect to the third party processors they share data with and the purposes for that sharing beyond broad categorical descriptions in their privacy policies. Using Solid, however, organizations could choose to make third party sharing visible to users, and could potentially provide consumers with choices with respect to if, and where, their data is shared with a processor. Even where the use of a third party is essential in order for an organization to provide its services, the visibility into third party sharing that Solid can enable would represent a significant leap forward.

5. Records of Processing Activities

The GDPR also features certain requirements that require backend (non user-facing) diligence on the behalf of organizations in relation to the data they use, requiring them to create and maintain records of their data processing activities. These records are created and maintained internally and are seldom, if ever, made visible to users, although they may be reviewed by regulators. The GDPR requires that organizations document the purposes of the processing they do, describe the categories of personal data they process, identify the categories of recipients to whom they disclose data, and document the instances in which they transfer data.

Where more "high risk" processing takes place, as in the case where profiling or systematic monitoring activities are conducted, sensitive data is collected on a large scale, or automated decision making is used to make significant decisions, the GDPR requires that organizations conduct a more detailed Data Protection Impact Assessment (DPIA).

Organizations using Solid can of course maintain Article 30 and DPIA records in the same manner that any organization does and should approach these assessments with the same diligence. However, because Solid enables organizations to provide users with significantly improved transparency and control over their data, this type of information need not be solely relegated to internal processes; users can instead be

shown, and helped to understand, how and why their data is being collected, processed and shared. This can be accomplished by bringing messaging forward to explain to users why their data is being collected and processed, designing control mechanisms that let them decide how their data can be used, and by enabling them to view the data that is being used in their Pods.

6. Profiling and Automated Decision Making

The GDPR defines profiling as any form of automated processing used to analyse or predict things about a person, such as a person's "performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements."

Where an organization engages in profiling, the GDPR imposes certain requirements that govern how these acts can be conducted. There are transparency requirements, which require organizations that carry out automated processing, including profiling, to provide meaningful information about the logic involved, as well as the significance and the envisaged consequences of this processing for the user. Users also have the right not to be subject to a decision based solely on automated processing, including profiling, that produces legal or similarly significant effects.

Automated processing with significant effects can be used in certain circumstances, including where the organization has the user's explicit consent. Organizations using this type of automated processing are required to implement "suitable measures" to safeguard the user, including the right for the user to "express his/her point of view", contest the decision, and obtain human intervention.

Organizations using Solid can and should comply with the GDPR's requirements for profiling and automated processing in the usual course. However, Solid also enables organizations to exceed these requirements in order to animate the values at the heart of the GDPR. In particular, Solid can enable organizations to provide visibility into the data used for profiling and AI processing, and to provide messaging explaining how such processing works and identifying the safeguards and rights that it has established in relation to these processes. By enabling real transparency to processes that, at present, operate largely in the dark, Solid can help build the trust and understanding of users.

7. Cross-Border Data Transfers

While the modern business landscape is increasingly international, and most users take for granted that information moves seamlessly around the world, the laws governing the security and privacy of user data remain bound by national borders. The GDPR applies to organizations that are established in the EU or who process the

personal data of people in the EU. If any of these organizations need to transfer the personal data of EU users *outside* of the EU, the GDPR requires the organization to establish that both the recipient country and the recipient company provide an acceptable level of data protection. The ways in which this threshold of acceptability can be met has been the subject of dispute and has fostered a significant amount of uncertainty since the GDPR came into force.

While the EU-U.S. Privacy Shield framework was created to allow certified U.S. companies to lawfully receive personal data from the EU, it was subsequently ruled to be invalid. The GDPR permits the use of "standard contractual clauses" to transfer data outside of the EU -- essentially a comprehensive standard contract that establishes safeguards for data, to which both parties must agree. But the Court of Justice of the European Union has continued to require modifications to the language of these clauses and to impose new conditions on their use. The court has also ruled that even where standard contractual clauses are in place, organizations need to perform a comprehensive case-by-case assessment before transferring EU data across borders.

The shifting legal landscape has made cross-border transfers the subject of significant uncertainty, if not outright confusion, for most companies. More broadly, however, users have been left largely in the dark with respect to understanding when their data is going to be transferred across a border and why. The GDPR permits data to be transferred outside the EU where explicit consent is obtained from the user for that transfer, once the user has been informed of the possible risks. This approach is rarely taken, however, despite the fact that it would provide the most transparency and could also accompany the use of other safeguards. This is likely due, at least in part, to the fact that most organizations lack an operational ability to show users where their data would be transferred or enable them to make informed choices in respect to such transfers.

The Solid protocol can provide organizations with the tools to deliver a much improved user experience in relation to cross-border data transfers. To begin with, any Solid Pod can be easily stored in a particular location, which in itself can provide the flexibility to process data without moving it across borders. More broadly, geo-tagging solutions are currently being explored that could be used to show users where their data is being stored and where it may be transferred. The ability to understand where data is stored and where it would be moved -- and the ability to make choices in respect of such transfers -- would represent a significant improvement to the current approach to cross-border data transfers.

8. Transparency

Transparency is a foundational privacy principle, a core component of the GDPR, and a fundamental element of ensuring that users understand and trust the use of their

data. The GDPR also recognizes transparency as the key to fairness, reflecting the need for data controllers to be open and clear about how they use personal information. The GDPR's requirements for transparency are broad but purposive:

> The principle of transparency requires that any information addressed to the public or to the data subject be concise, easily accessible and easy to understand, and that clear and plain language and, additionally, where appropriate, visualisation be used...This is of particular relevance in situations where the proliferation of actors and the technological complexity of practice make it difficult for the data subject to know and understand whether, by whom and for what purpose personal data relating to him or her are being collected... [GDPR Recital 58].

However, providing users with real and meaningful transparency, in a way that lets them see and understand what's happening to their data, has continued to prove elusive in the post-GDPR world. In most cases, organizations meet GDPR transparency requirements by providing general descriptions of the data they use and their processing purposes in lengthy privacy policies. In some cases, particularly where a particular act of processing may be unexpected, organizations may provide more specific "just in time notices" that give more specific explanations. But the use of such notices tends to be inconsistent, and where there is a significant amount of data processing to explain they can be cumbersome and lead to "consent fatigue," discouraging users from understanding or even reading the notices.

Solid enables organizations to meet GDPR transparency requirements and to significantly exceed them. Using Solid Pods, organizations can enable users to see and control their personal data that is being used -- moving away from a reliance on privacy policies and towards an integrated solution that is both intuitive and functional.

Conclusion

Solid provides organizations with the tools to comply with the GDPR and to exceed its requirements in a number of fundamental ways, letting them better achieve the values at its core.

Using Solid, organizations can enable users to exercise their data subject rights directly by seeing and acting on the data inside their Pod, enhancing the user rights model by providing real transparency and user control. Organizations using Solid can rely on any of the GDPR's lawful bases to process data, but in particular can obtain a user's clear, informed and specific consent in a way that integrates with the functionality of a product or service. Solid provides organizations with the tools to pinpoint the data they need and be clear about the purposes for which they are collected, and to significantly minimize the data that is collected and stored across multiple organizations because a single Solid Pod can be used to share data with a range of different companies or governments.

Organizations using Solid could choose to design solutions that make third party sharing visible to users, or could provide users with choices as to whether and where their data is shared. They could use Pods to localize data into a single geographic location and to show users where their data is being stored and where it would be transferred. They could build products that show users what data may be used for profiling and AI processing and provide messaging to explain how such processing works.

By using Solid, organizations can empower users to see how their data is used and to exercise real and meaningful control over it, building for transparency and control and engendering real trust.