

WHITE PAPER

Why a Good Additional Technical Safeguard is Hard to Find

A response to the consultation on the EDPB draft recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data

2020

 MMUTA™

The Court of Justice of the European Union (CJEU)'s judgment in Schrems II was not really a surprise.

Although the Advocate General tried to find a solution that would have been more deferential to the European Commission (EC), the CJEU held, for a second time, that the EC adequacy decision was invalid. This came while the CJEU upheld the EC decision on Standard Contractual Clauses, which will be soon replaced by a new implementing act. As a result, it is no longer possible to justify a restricted data transfer from the European Union (EU) to the United States (US) on the basis of the Privacy Shield.

The European Data Protection Board (EDPB) recently issued its draft recommendations in an attempt to shed some light on the implications of the Schrems II decision and explain how Standard Contractual Clauses and other 'appropriate safeguards' under Article 46 of the General Data Protection Regulation (GDPR) could ever help legitimize international data transfers. The main message stemming from the CJEU judgment is confirmed: Contractual and organizational measures will never be enough if the threat that must be prevented is access to data by intelligence services. This makes sense as by definition, legal obligations and organizational processes are useless in this context. The question that remains, though, is whether technical measures can ever usefully complement Standard Contractual Clauses and if so, how.

The EDPB draft recommendations are a key piece of guidance in that they attempt to detail and assess the potential of various additional technical safeguards. However, not many appear promising. In this response to the consultation on the EDPB draft recommendations, we aim to unpack the potential of three of these technical safeguards, i.e, encryption, pseudonymization, and split or multi-party processing, and call for clearer delineation between concepts.

1. Encryption

State-of-the-art encryption appears to be an interesting option for data in transit, as long as decryption is only possible outside the third country in question, for which there is no adequacy decision. As mentioned in our previous [blog post](#):

"encryption is a function that can be reversed with what is called a 'decryption key.' An encryption algorithm, also called a cipher, is what takes a readable chunk of text and turns it into seemingly random values that are not decipherable to others, at least, not without the decryption key."

Encryption however, can also be employed as a form of reversible masking for pseudonymization purposes, not simply to protect data in transit. This thus brings us to pseudonymization.

2. Pseudonymization

The EDPB recommendations are slightly confusing in that they seem to refer to the standard for anonymization rather than pseudonymization. We'll explain why.

To start, it is important to make sure key concepts are understood properly.

De-identifying personal data requires classifying the data into two groups: direct identifiers and indirect identifiers. As we explained in our previous [blog post](#):

identifiers are personal attributes that can be used to help identify an individual. Identifiers that are unique to a single individual, such as social security numbers, passport numbers, and taxpayer identification numbers are known as "direct identifiers." The remaining kinds of identifiers are known as "indirect identifiers" and generally consist of personal attributes that are not unique to a specific individual on their own. Examples of indirect identifiers include height, race, hair color, and more. Indirect identifiers can often be used in combination to single out an individual's records. Latanya Sweeney, for example, famously showed that 87 percent of the US population could be uniquely identified using only three indirect identifiers: gender, five digit zip code, and birth date.

A [more recent study](#) appearing in Nature Communications even found that 99.98% of Americans would be correctly re-identified in any data set using 15 demographic attributes.

Pseudonymizing personal data is often thought as a means to transform personal data in such a way that the individual is not directly identifiable anymore. This is confirmed by prior guidance issued by Art. 29 Working Party (which has been replaced by the EDPB), the French Data Protection Supervisory Authority (Commission Nationale Informatique et Libertés or CNIL), German Data Protection Supervisory Authorities, or even the UK Data Protection Supervisory Authority (the Information Commissioner's Office or ICO). The most recent guidance from the European National Security Agency (ENISA) adds that best practice mandates the adoption of a risk-based approach to be in a position to optimize the tradeoff between security and utility:

Hence, a trade-off between utility and data protection can be considered (...). When considering the application of pseudonymisation to real-world scenarios, this trade-off should be analysed carefully, so as to optimize utility for the intended purposes while keeping the protection of the pseudonym holders (data subjects) as strong as possible.

Importantly, even if the data is pseudonymized, it is possible that by combining the pseudonymized data with information that is publicly available or attainable, the individual to whom the data pertains remains (indirectly) identifiable.

The EDPB states, however, that in order to make pseudonymization an effective additional control:

the controller [must] establish by means of a thorough analysis of the data in question taking into account any information that the public authorities of the recipient country may possess that the pseudonymised personal data cannot be attributed to an identified or identifiable natural person even if cross-referenced with such information.

In other words, to make pseudonymization an effective complementary control, it is necessary to look beyond the additional information held exclusively by the data exporter and consider publicly available or attainable information to which an attacker has or could have access to. As a result, in order to assess whether pseudonymization is sufficient, it is necessary to determine whether the individual is indirectly identifiable as well. The end goal is thus now moving towards that of anonymization. As a reminder, data has successfully undergone anonymization if "the data subject is not or no longer identifiable," both directly and indirectly as per GDPR Recital 26, which should be read together with GDPR Article 4.

To determine whether the individual does remain (indirectly) identifiable, best practice suggests implementing a formal attack model. One model that could make sense for international data transfer to the US, where the concern is surveillance by US intelligence services, is the database cross match model (DBXM) (Elliot & Dale, 1999), which posits an attacker who aims to enrich their database by linking it to a target data set. This model is related to the common attack models used in privacy risk assessments, such as the prosecutor and journalist attack models (PAM and JAM, respectively) defined in (Marsh et al., 1991). As we explain in the following paragraphs, the relevance of DBXM or its variations is implicitly confirmed by the EDPB, which expressly states that "if you still wish to proceed with the transfer, you should look into other relevant and objective factors, and not rely on subjective factors such as the likelihood of public authorities' access to your data in a manner not in line with EU standards."

Under DBXM, it is assumed that an attacker wants to associate records within their database with individual records in a target data set. We may assume the attacker's database is quite rich and has already incorporated all publicly available or, if we would like to be conservative, publicly attainable information.

Under PAM, the attacker aims to locate the (de-identified) record corresponding to a single target individual, reminiscent of a prosecutor who hopes to convict their target. The journalist attack model, on the other hand, assumes that the attacker aims to re-identify at least one (arbitrary) individual as, perhaps, a means to discredit the release or the data exporter. These attack models are often used when anonymizing data for a public release or for clinical research purposes. The European Medicines Agency acknowledged this by distinguishing between the approach followed by Article 29 Working Party in its 2014 opinion and a risk-based approach that is seemingly not precluded by the 2014 opinion. While the mechanics of an attack under DBXM operate similarly to its PAM and JAM counterparts, it differs in that its goal is to maximize the number of linked records.

Note that the 'motivated intruder test' set forth by the ICO is less restrictive than PAM or JAM in that it assumes the attacker, called a motivated intruder, is not a specialist. Such a standard is an inappropriate model for attacks performed by intelligence services, which should be presumed to have specialist knowledge.

Under these models, the only way to claim that the data is de-identified is to implement stringent data controls, via privacy enhancing data transformation techniques. Marsh et al., 1991 propose viewing the re-identification probability as a conditional probability consisting of the product of two factors, the probability of identification given that an attempt is occurring, and the probability of an attempt. The former factor, which is amenable to theoretical analysis, can be viewed as the re-identification risk given the use of privacy enhancing data transformation techniques, (i.e., the *data risk*). The latter factor (i.e., the *context risk*) remains situationally dependent and is more subjective since it primarily depends upon situational, environmental, and/or incentive-based factors that often cannot be reliably quantified. This is especially true when the attacker, an intelligence service in this case, may not be well understood. Still, it remains possible to argue that the overall re-identification probability is low from analysis of the data risk alone, as the data risk provides a controllable mathematical upper bound on the overall re-identification risk.

A few data transformation techniques can potentially be used to achieve a negligible or remote data risk. The most obvious ones are k-anonymization and differential privacy. While k-anonymization perturbs indirect identifiers to make it possible for individuals to hide in groups of a k number of individuals, differential privacy randomizes query computation to produce query results in the form of safe aggregates. As defined in our previous [blog post](#):

k-Anonymization is a data generalization technique that ensures indirect identifiers match a specific number of other records, making it difficult to identify individuals within a data set (the total number of matching records is referred to as "k," and hence the name). For example, in data that's been k-anonymized, if k is set to 10 and where indirect identifiers include race and age, we would only see at least 10 records for each combination of race and age. The higher we set k, the harder it will be to use indirect identifiers to find the record of any specific individual.

Differential privacy is a family of mathematical techniques that formally limit the amount of private information that can be inferred about each data subject. There are two main flavors of differential privacy, offering slightly different privacy guarantees: “global,” which offers data subjects deniability of participation or inclusion within a data source, and “local,” which offers deniability of record content. Despite being slightly different, both operate by introducing randomization into computations on data to prevent an attacker from reasoning about its subjects with certainty. Ultimately, these techniques afford data subjects deniability while still allowing analysts to learn from the data.

By way of example, k-anonymization has been used by the US Research Data Assistance Center to protect data derived from two large insurance programs administered by the Centers for Medicare and Medicaid Services in the US. It is also referred to by the European Medicines Agency as a key anonymization technique. The US Census Bureau also announced that “2020 Census results will be protected using “differential privacy,” the new gold standard in data privacy protection.” Reading Art. 29 Working Party’s opinion on Anonymization Techniques of 2014, it appears that differential privacy is one of the most promising anonymization techniques, as it has the potential to mitigate singling out, inference, and linkability.

Since these techniques can have serious utility implications, however, it is essential to explore and further develop their dynamic implementation so that they maximize utility for a given set of queries.

3. Split or multi-party processing

The last technical control worth mentioning is what the EDPB calls ‘split or multi-party processing.’ However, it is difficult to understand from the EDPB’s description when such a split or multi-party processing would make sense.

For context, secure multi-party computation (SMC) is a branch of cryptography concerned with designing protocols that enable somewhat adversarial parties to jointly perform a computation, while keeping their respective inputs secret from each other. For example, two hospitals may wish to determine which patients they treat in common without revealing to each other the names of other patients. As another example, a government may wish to enact a voting system whereby all members vote to elect a leader while keeping their votes secret from each other.

Roughly speaking, the parties participating in these protocols are only able to learn whatever is jointly inferable from the knowledge of their own input together with the output, if permitted to see the results.¹ This is of particular relevance here, as we are directly concerned with what is inferable by parties in other jurisdictions.

The attack modes considered in SMC include one or more corrupted party members, who possibly collude with other corrupted party members. Attacks may be *passive*, in that corrupted party members only share knowledge of their input and/or intermediate state with other corrupted party members, or could be *active* (“malicious”), in that corrupted party members attempt to lie to (and possibly collude with) other party members about their data and intermediate results in order to achieve a desired outcome.

¹ Typically, all parties learn the output of the joint computation. However, protocols can be designed so that only some subset of the parties learn the results.

The EDPB ‘split or multi-party processing’ use case seems to envision a scenario in which the data exporter is in possession of personal data to be outsourced for processing. Before discussing this in detail, let’s outline two scenarios where cryptographic computation may be useful in light of the Schrems II decision:

1. A data exporter wishes to utilize the computational resources operating in another jurisdiction.
2. A data exporter wishes to engage in joint processing with contributing² collaborators residing in other jurisdictions, and party members would like their input to remain secret from the data exporter and/or the jurisdiction thereof.

Item 1 can be addressed by fully homomorphic encryption (FHE), which enables the data exporter to provide encrypted inputs to processors in untrusted jurisdictions who can then, without decryption, compute the encryptions of results. In turn, those results can only be decrypted, and therefore read, by the data exporter.

Item 2 can be addressed by SMC. The parties utilize SMC to perform the desired processing over their private inputs. Upon completion of the protocol, the parties learn nothing more than what is jointly inferable from the processing results and their respective inputs. If desired, the protocol may be designed such that only the data exporter learns the results.

As written, the ‘split or multi-party processing’ scenario seems to have elements of both Items 1 and 2. However, in adopting mutually exclusive requirements from both, it fails to satisfy either.

The scenario description appears to suggest that the data exporter is in possession of the entire input, which they split and distribute to processors in other jurisdictions. This is curious for a few reasons:

1. If the data exporter is already in possession of the entire input (specifically in the sense that they do not require the private contributions of extra-jurisdictional party members), then there is nothing preventing the data exporter from simply performing the processing themselves. If they simply do not wish to carry out the processing “in house,” they may reasonably delegate processing activities to a processor residing in their jurisdiction. Instead, they split and distribute the data to extra-jurisdictional parties.
2. If the split portion of the data received by a party is encrypted beyond their means to manipulate it, and they possess no data of their own to contribute, then this situation is equivalent to FHE with the additional and artificial constraint that the evaluation be distributed across jurisdictions. It is not clear to the authors whether or not such a scheme is known, but the point is moot. FHE is sufficient regardless of the jurisdiction where the computation takes place. In other words, the data need not be split in the first place, even if the entire FHE computation is to take place in a hostile jurisdiction.

² Here, contributing means that the collaborator will contribute input data to the joint processing activities, as opposed to simply carrying out a computation on behalf of other parties. In other words, some of the input data comes from data which resides in this jurisdiction.

Instead, the split appears intended to force alignment to a setting reminiscent of SMC. This is made clear in the list of conditions under which the EDPB would consider split processing an effective supplementary measure. The first condition states:

a data exporter processes personal data in such a manner that it is split into two or more parts each of which can no longer be interpreted or attributed to a specific data subject without the use of additional information.

while a later condition continues,

the processors optionally process the data jointly, e.g. using secure multi-party computation, in a way that no information is revealed to any of them that they do not possess prior to the computation,

It should be noted that neither of the above conditions explicitly require encryption, only that the data not be attributable to a specific data subject without the use of additional information. This rules out the possibility for participating extra-jurisdictional processors to augment their received portion of the data split with additional personal information, simply because joining to the received data implies being able to match on de-identified records, which in turn implies the existence of an extra-jurisdictional means of re-identification.

In the absence of requiring party members to keep their inputs private from each other, the problem is trivially solvable outside of this use case by having these collaborators send their inputs to the data exporter's jurisdiction for processing. Thus, to the extent that an extra-jurisdictional processor possesses additional information which they wish to contribute, we may assume those willing to share directly have done so, and are thus not parties in a joint computation via SMC. We may therefore assume that all parties in any SMC protocol under this use case are contributing parties with private input.

So that's it, right? This use-case provides adequate protections when all parties are contributors of private data, provided that all extra-jurisdictional partners are unable to learn anything new about the exporter's data from their participation? Not so fast — there appears to be a gap between satisfaction of the use case criteria and requirements for safe implementation of SMC.

Typically speaking, protocols can be designed to detect active adversaries, but may require that a (possibly large) fraction of the party members participate honestly. This is potentially problematic, as it seems to necessitate involving other jurisdictions for the sole purpose of ensuring an honest majority. Further, it requires a basis for believing that the external jurisdictions will not collude in an attempt to subvert the protocol. Along these lines, the EDPB conditions,

...there is no evidence of collaboration between the public authorities located in the respective jurisdictions where each of the processors are located, which would allow them access to all sets of personal data held by the processors and enable them to reconstitute and exploit the content of the personal data in a clear form in circumstances where such exploitation would not respect the essence of the fundamental rights and freedoms of the data subjects. Similarly, public authorities of either country should not have the authority to access personal data held by processors in all jurisdictions concerned.

From a technical standpoint, the requirement is too weak. In particular, it requires that no public authority be able to access all sets of personal data. However, in order to be effective, it should instead be required that the fraction of possibly colluding authorities not be allowed to exceed the threshold at which the security of the protocol can no longer be guaranteed. In general, though protocol-dependent, the number of corrupt parties required to subvert the protocol is typically much less than “all” and often a simple majority will suffice. Protection against adversaries with less access to other parties’ data provides a necessary higher bar for security.

Ultimately, assuming split or multi-party processing makes sense, it cannot be used to enrich data sources located in third countries or generate new insights in these third countries.

To recap, effective additional technical controls thus appear very limited. The EDPB draft recommendations seem to suggest that unless the data recipient located in the third country is a protected recipient, which will have to be specifically ascertained, restricted transfers will require a standard leaning towards anonymization. This holds true, unless it is possible to rely upon derogations under GDPR Article 49, which should mean that the transfer is occasional and non-repetitive, or the EDPB agrees that if a data controller is subject to GDPR as per Article 3, the transfer is not a restricted transfer within the meaning of Chapter V, which we don’t know yet as the question has been left open in prior guidelines. Until a new adequacy decision targeting the US is issued by the EC, this means that data controllers should rely upon localization-based access control and make sure only data analysts located in the EU have access to EU data each time dynamic anonymization does not fit their use case’s requirements.

Sophie Stalla-Bourdillon

Sophie Stalla-Bourdillon is Senior Privacy Counsel and Legal Engineer at Immuta, where she works on analysing and developing solutions to support smart data governance and risk assessment within data analytics environments. Sophie is also a Professor in Information Technology Law & Data Governance at the University of Southampton, where she co-directs the Web Science Institute.

Alfred Rossi

Alfred Rossi is a theoretical computer scientist and Research Scientist at Immuta, where his efforts are currently focused on mathematical privacy, including analysis and development of privacy enhancing technologies and techniques, and risk analysis. Alfred is also a Senior Lecturer in the Ohio State University Department of Computer Science and Engineering.

