Broadcom comments on EDPB guidelines 01/2021 on Data Breach Notification Examples

March 2nd, 2021

Introduction

EDPB guidelines 01/2021 (henceforth “the Guidelines”) are a very welcome, practical complement to the prior data breach notification guidelines of Article 29 Working Party Opinion WP250rev1[1]. By and large the types of breaches and the specific examples listed in the new Guidelines are well aligned to our own field experience as well as to the feedback we have received from our customers in preventing, and as the case may be, responding to data incidents, or building technology that enables our customers and partners to do the same[2].

Accordingly, we share the views expressed in the guidance that certain organisations should be held to a higher standard when it comes to cybersecurity due to the nature of the data they process. Maintaining a high level of cybersecurity is critical to achieve accountability

A reasonable level of transparency, where feasible, is another key element of cybersecurity in our view that we endeavour to achieve by publishing some of our cybersecurity practices as well as by disclosing certain product features through transparency notices.

The realistic assessment of risk is the third critical component that is required in delivering privacy-friendly results. The resources of organisations are not limitless, this is the case even for the big ones. Treating everything as high risk ultimately results in being able to defend nothing effectively as historical experience shows[3].

We have taken note of certain elements in the guidance and we have already started looking into how we can adjust the assessment of our own breach readiness so as to better map to the new Guidelines. At the same time as daily practitioners of cybersecurity we need to base our assessments on risk scenarios and potential impact that would go beyond the theoretical possibility of a “residual risk”, into risk that can have an impact and a probability to materialise. Depending on its gravity, such risk should be sufficient to trigger legal obligations. If any residual risk can be sufficient grounds for notification it would eliminate any kind of processing activity or relationship because there will always be some residual risk as there cannot be 100% cybersecurity. The precautionary principle of the GDPR provisions require a tangible and

[1] https://ec.europa.eu/newsroom/article29/document.cfm?action=display&doc_id=49827

[2] The one example in the Guidance which some might find questionable is case number 12: Some might be tempted to argue that a loose paper notebook with no access controls or organisational measures around it does not involve any automated processing and does not constitute a “filing system”, so it is not subject to the GDPR per the criteria of Article 2(1). In reality the processing should certainly be subject to the GDPR. However if a case like that were to occur in real life, chances are that the controller -- who is very obviously unaware of any of the GDPR’s requirements -- would not be concerned with notifying the breach, even if it came to notice it. The fundamental issue with case 12 does not start with the breach of the data, but more generally and before that with the violation of practically every requirement of the law.

[3] “*He who defends everything, defends nothing”* Frederick the Great

reasonably realisable risk. We are pleased to see that in most scenarios described in the guidance this is indeed the case. Our experience shows that in practice national supervisory authorities take a similar approach pursuing action for incidents with impact in terms of number of affected individuals and reasonable risk factors materialising. The alternative would be counterproductive as it would drown the limited resources of national supervisory authorities in an avalanche of “prudential” notifications.

As a leading global provider of enterprise solutions for identity, network and information security technologies, we also have a unique perspective on how the regulatory compliance requirements around data breach notification translate into operational business practices. We will not be discussing in greater detail than the previous comments the legal and conceptual aspects of assessing risk and determining notifiability, which the text of the GDPR describes in great detail and which the Guidelines cover in sufficient depth.

Instead, we would like to use the opportunity of this consultation to share our pragmatic experience of how the appropriate technical and organisational measures that can best help organisations to prevent, detect and respond to security incidents may be deployed in enterprise environments, what difficulties tend to arise, and what trade-offs organisations have to consider. Our goal is to encourage further discussion in the privacy community around these issues, and hopefully, in due course, to help forge a consensus between market operators, data subjects and data protection authorities on what the appropriate balance should be between the various regulatory and operational constraints at play.

Monitoring For Data Breaches Is Itself A Privacy-Impacting Activity

Even though it concerns the aftermath of security incidents, GDPR Recital 87 is still perfectly unambiguous in terms of what controllers -- as well as, in their own remit, processors -- are expected to do before the breach even occurs:

“*It should be ascertained whether all appropriate technological protection and organisational measures have been implemented to establish immediately whether a personal data breach has taken place and to inform promptly the supervisory authority and the data subject.*”

Obviously, investigating and establishing what happened would be impossible without the appropriate technical and organisational measures having been in place beforehand. Accordingly, many of the measures recommended by the Guidelines include *ex ante*, protective and preventative measures such as:

- “*appropriate, up-to-date, effective and integrated anti-malware software*”;

- “*firewall and intrusion detection and prevention [and other perimeter defence] systems (...) even in the case of home office or mobile work*”;

- "*[logging of information security events and] forwarding or replication of all logs to a central log server*";
- "*[timestamped] record[s] of all [software and firmware] updates performed*";
- "*proper access control policies and forcing users to follow the rules*";
- "*[forcing] [strengthened] user authentication when accessing sensitive personal data*";
- "*checking unusual data flow between the file server and employee workstations*";
- "*disabling open cloud services*";
- "*forbidding and preventing access to known open mail services*";
- "*functionalities of highly mobile devices that allow them to be located in case of loss or misplacement*".

We wholeheartedly welcome this explicit enumeration of some of the most adequate technical and organisational solutions which the EDPB considers as relevant and advisable to defend against, and react to data breaches. As a matter of fact, we also have ample empirical evidence of such tools and techniques not only being mature and readily available, but also very highly effective when implemented correctly and maintained on an ongoing basis.

At the same time we must emphasise that, as with any other business activity, and by the very letter of GDPR Article 32 ("*taking into account the state of the art, the costs of implementation, and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons*"), the use of any such solutions will always be the result of a calculation between what is theoretically feasible, what is legally permissible, what is proportionate to the risk, what is affordable for a specific organisation and what makes business sense. In other words, not all of these tools will be used -- or even usable or relevant -- in every situation, and nor will they always be employed to the fullest of their capabilities. We will illustrate that with observations commonly experienced by organisations when deploying these various types of measures, grouped here into conceptually similar or related batches:

Security Measure	**Practical Observation**
● *appropriate, up-to-date, effective and integrated anti-malware software*	As indicated earlier, security measures themselves need to be prioritised based on risk, on the value of the information to be protected and on the ability as well as on the maturity of an organisation to manage them. Deploying and enforcing the use of malware and intrusion detection technologies on corporate-

● *firewall and intrusion detection and prevention and other perimeter defence systems even in the case of home office or mobile work*	managed devices -- including even properly managed BYOD -- is not a problem today. However organisations realise that in the face of ever changing, fast evolving, highly customised, nimble malware and targeted intrusions, using only static, locally managed detection based solely on pre-defined malware signatures will not be successful. In order to meaningfully detect and to effectively block malware and intrusions, countermeasures today need to rely on a complex combination of multiple techniques. Those include signatures, but also heuristics, behavioural analytics, isolation, sandboxing, detonation, and other highly advanced, sophisticated procedures based on in-depth defence and zero-trust paradigms. Because of the complexity, fluidity and scale of the threat facing organisations, no cyber-defence strategy will be conclusive unless it is able to leverage a combination of big data threat intelligence, automated machine learning, near-real-time artificial intelligence and last resort human intervention by highly qualified experts. In practice, no endpoint security technology in the form of antivirus or firewall on an endpoint device and no single end-user will have the computing power, storage capacity, bandwidth, software capability and technical skills to self-defend against all the threats encountered. Not to mention that almost no organisation in the world possesses in-house the all-round competency and capability to handle and respond to the malware and intrusion detections that occur across its environment. Instead, endpoints need continuously to share telemetry about the activities and events they encounter, for that data to be transferred to, analysed by and reacted to through cloud-powered security operations performed on -- and delivered from -- remote servers, often by specialised third-party service providers. The reason this can have privacy implications for organisations is that a lot of that indispensable threat telemetry contains elements that may under certain circumstances be directly or indirectly revealing of the actions of a specific end-user (e.g. the employee), and/or of any other person with whom they may communicate (e.g. co-workers, external interlocutors, and even a cybercriminal posturing as a trusted email correspondent). Some level of pseudonymisation can be implemented to limit the amount of personal data collected, but in the end there are limitations to it because the purpose of cybersecurity is to be able to pinpoint a particular activity, device, or network location and to take corrective action in the form of blocking, patching, restricting or remediating. In other words, the telemetry that is necessary to prevent and detect malware incidents and cyber intrusions that could lead to

	data breaches may contain some transient elements of personal data, which need to be transferred in real time on an ongoing basis, often internationally in a follow-the-sun model, typically to one or several third-party providers. Moreover, because this kind of telemetry constitutes the benchmark against which future occurrences of the same or similar threat can be detected, the processing of the telemetry beyond the sole purpose of blocking one specific attack against one specific organisation is indispensable. Typically organisations that are security conscious understand the need for cybersecurity and why such telemetry is being transmitted, shared, or further processed in a purpose-compatible manner. They are uncertain however of the legal basis justifying such processing under privacy law. Security professionals in those organisations realise that by failing to permit such processing to be pursued, they undermine the very efficiency, and the viability even on the short term of the technology they utilise. The most common concern raised is "how could this be done so as not to violate our privacy obligations". We strongly believe that this is possible (see further below), and we do not encourage positioning the problem in terms of a trade-off between information security and privacy compliance. However we also believe that the uncertainty, fear and doubt in the marketplace around this issue can only be dispelled through explicit and unambiguous guidance from data protection authorities recognising the necessity of such processing. We are also of the view that clear recognition of the challenges in this area would ultimately benefit privacy because it would reduce the number of incidents and thus privacy risk. In addition it would also ensure that the use of security technology is done in a privacy-friendly and proportionate manner.
● *logging of information security events and forwarding or replication of all logs to a central log server* ● *timestamped records of all software and firmware updates performed*	The inherent difficulties of logging are twofold: (1) It is impossible to foretell what specific information will subsequently turn out to be relevant to an incident and should have been logged; and (2) it is equally impossible to determine in advance how long after the actual event a particular log entry will become useful to an investigation and any legal obligations that may result out of that (e.g. legal preservation requirements). Similarly to the threat telemetry discussed earlier, end-user activity logs are replete with information that not only directly or indirectly identifies end-users, but also often reveals details about the circumstances and contents of their actions, activities and behaviours. Consequently security log collection itself is by its very nature a highly privacy relevant and personal data intensive

processing activity.

As such, it is naturally subject to the principles and safeguards of the GDPR, among others as they relate to the collection, retention, protection and destruction of logs, as well as to their transfer across borders for example for the purpose of centralising or replicating them on a secure and tamperproof server. Among organisations' primary concerns with logging are the questions of:

(i) *what legal basis to rely upon to collect the logs in the first place* (see below for further remarks on that point);

(ii) *what exact data to collect without running afoul of data minimisation*;

(iii) *how much log data to collect without infringing the necessity and proportionality principles*; and very importantly

(iv) *how long to keep log data without breaching the retention limitation principle*, all the while not actually knowing when the data may come to serve, possibly as evidence in defending a breach notification case before the supervisory authority or even in court.

Our field experience suggests that organisations at large remain highly confused and uncertain as to what is actually required for compliance, versus what would be viewed as excessive and therefore unlawful. The difficulty is further compounded by the complexity of the multilayer regulatory environment under which an organisation may have to perform logging not just for the purpose of being able to investigate potential breaches of personal data, but also for the purpose of managing and demonstrating compliance with other, cumulative regulatory mandates and contractual requirements of information security, such as (and without limitation) under the Electronic Communications Code, the ePrivacy Directive, the NIS Directive, the eIDAS Regulation, or the European Banking Authority guidelines on IT risk management and outsourcing. Moreover local requirements either for data preservation, or for labour law compliance create further complications that can have a direct impact on the compliance posture of an organisation. This may result in organisations not collecting enough logfiles that would enable them to prove that a notifiable breach occurred or that would enable them to conclusively demonstrate that data was not compromised.

Therefore we would very much welcome further dialogue among the privacy and security community, and further clarification from regulatory authorities on how to calibrate the right balance between *necessary* and *sufficient*. Or in other words how to define what is both reasonably adequate for good privacy and security, and legally acceptable for compliance.

● *proper access control policies and forcing users to follow the rules* ● *forcing strengthened user authentication when accessing sensitive personal data*	Strong and differentiated access controls based on the systematic enforcement of robust policies through auditable technical and organisational measures are a very common baseline good practice, with a plethora of tested and proven solutions of many kinds used successfully in all types of environments and industry fields. Having said that, governing (granting, challenging, denying, revoking) a user's access to specific data on the basis of the existence, probability, absence or lapse of legitimate need-to-know is a construct that is not easy to put into practice. It is simple to understand conceptually but it is difficult to implement operationally. On the face of it, it seems possible to define a rigorous data access regime where every user has the access permissions necessary to their role. In fact however, the real life use cases are extremely rare where the permissions and privileges of an individual user can be defined in a way that will always exactly and flawlessly match the legitimate business needs of that user to access *all* the data, and *only* the data they need in order to accomplish their task. It is also unlikely that during the employment relationship those access rights will not have to be dynamically adjusted to business requirements or changes in professional roles. There is a whole series of factors at play which require introducing a reasonable and always variable share of flexibility and tolerance. Among many others: (1) Tremendous amounts of data are unstructured by nature (e.g. contents in emails and other interpersonal electronic communications, in word processing documents, in presentations, on collaboration platforms, etc.). In such settings, the data tends to be inextricably co-mingled with other confidential information, not all of which may be necessary or relevant to a given user's task. Therefore user permissions can hardly be set at the level of the data elements: practically they can only be defined and enforced at the level of files, or more realistically systems (e.g. specific applications, specific repositories, specific subdomains within a platform, etc.), or on the basis of requirements/limitations directly applicable to content, e.g. inability to copy/paste sections of a document or to forward content that contains sensitive data (for instance credit card numbers). (2) As the Guidelines themselves recognise, no technical or organisational measures can fully protect from, or compensate for certain human factors that may range from ignorance and negligence to recklessness and malicious intent. Even perfectly legitimate access permissions can be misused for illegitimate

purposes. A considerable proportion of data breaches can be partially or fully traced back to ill-advised, careless or overly bold user actions, all the way to malicious insiders abusing their privileges.

(3) Even the nature and therefore the accessibility of the very same piece of data may vary depending on the context in which it is being used, and/or on the purpose for which it is being accessed. To give a very simple example, the browsing logs of a user using their corporate device to browse an adult website will be very different depending on whether those logs are processed:

(i) by corporate security to check that the queried website does not harbour malicious content, on account of the fact that such websites are notorious cyber-threat vectors;

(ii) by the HR department to conduct an ethics and compliance investigation for disciplinary purposes, on account of the suspected violation of the company's acceptable use policy;

(iii) by a law enforcement agency exerting its lawful authority to combat illegal or harmful content online;

(iii) or by the user's manager for validation and approval purposes for instance if the user's job is to track and help flag and take down child abuse imagery from online platforms.

For all these reasons -- and many others --, effective access controls in today's highly complex digital data ecosystem need to be not only user-specific and object-oriented, but also highly context-relevant. The same user may very legitimately have a business need to access a piece of information in one situation, and no reason at all to access that same data in another situation. In some cases, a risk-based judgment call may need to be made, balancing the risk of data exposure on the one hand, and the risk of business disruption on the other. This recognition that there is no "one-size-fits-all" and "once-and-for-all" solution to govern user access to data is what has led to the emergence of paradigms such as information centric analytics, zero-trust security and context-aware access control. These in turn underpin concepts like data loss prevention technology, secure cloud access brokering and risk-based multi-factor authentication.

Now as the above example illustrates, the use cases for differentiated access controls are often eminently linked to the user's own privacy. Therefore even the very act of granting or denying an individual access to a certain data in a certain context may have privacy impacting consequences, for that individual and/or possibly for others.

Accordingly, an effective access control system is one that grants legitimate access where necessary while preventing unauthorised or even simply undesirable access where that could be risky or unnecessary. Such a system can only be effective if it is able, in real time, to make automatic and well-informed determinations based on the user's existing permissions against a predefined policy and their prior risk history, as well as to take additional measures (e.g. force multi-factor authentication) or to enforce exceptions (e.g. override an existing permission and deny access) in light of:

- The credibility of the access request that is made;
- The circumstances (time, place, environment, surrounding events and activities) in which the access is invoked;
- The characteristics of the data that is to be accessed (even if that data is unstructured, and has no metadata attached to it that would give indications about its nature, substance or intended use);
- The stated, apparent or probable purpose for which the access is requested;
- The assessed risk of subsequent misuse that could result from the access being granted in the given situation.

In providing this functionality the relevant technology may have to rely on a number of capabilities such as inspecting communication content, identifying and classifying particular data sets as privacy relevant, or even taking remedial actions to compensate for the risk or the possibility that a user acts in error, such as by displaying warnings that personal data are about to leave an organisation's perimeter or by requiring the use of encryption before sending a particular data set to a third party.

Operationalising an effective access control system where users' permissions need to be fluid and flexible, where data usage can occur in many different scenarios, and where data itself may be totally unstructured requires that the system be able to understand and benchmark certain user attributes and actions, to identify and evaluate relevant circumstantial factors, and to assess and classify the data even if it is unstructured. That means not only collecting information about the user actions and processing the very data that is to be accessed before access is actually granted, but also doing all of that in a demonstrable and auditable fashion, so that the automatic decisions that are taken can be explained, and any subsequent disputes can be fairly adjudicated. These again, as mentioned before, are also subject to further local requirements either for data preservation, or for labour law compliance, creating

	further complexity for organisations to manage. So in short, the very activity consisting of defining and managing users' access to data can be a privacy impactful process, entailing the observance of all the requirements and safeguards foreseen in the GDPR. That is no trivial task. Just like with topics discussed previously, our experience is that for fear of regulatory complications, many organisations will stop short of undertaking the full effort of properly developing and calibrating the right levels of context-aware access controls that the state of the art could afford. This tends to result in organisations defaulting either to overly restrictive access controls that undermine business efficiency, or on the contrary to overly permissive policies that fail to provide the necessary -- and differentiated -- levels of access control and protection to various classes of confidential information such as different categories of personal data. Once again, further debate among practitioners and guidance and reassurance from regulators would be highly desirable to improve general hygiene and corresponding compliance efforts around the technical and organisational measures to control and document access to personal data.
● *checking unusual data flow between the file server and employee workstations* ● *functionalities of highly mobile devices that allow them to be located in case of loss or misplacement*	Practically all of the arguments raised earlier could be reiterated and applied to the detection of unusual activities and behaviours, and to the tracking down of stolen or lost devices. However, beyond all the points that have already been made about the privacy-impacting nature of any such security measures, and about the regulatory and compliance consequences that attach accordingly under the GDPR, there is a critical additional element to highlight here: Both the analytics involved in managing identity or checking unusual data flows, and the geotracing of corporate-managed devices are predicated on a certain level of monitoring and benchmarking of what "usual" or "normal" device activity looks like. To the extent that such monitoring and benchmarking unavoidably relate to the habitual actions of the end-user concerned, and to the extent that they are used to detect any meaningful or suspicious departure from the established baseline, these measures may potentially be privacy-impacting. Moreover they are at a very sensitive crossroads between: ● Legitimate-interest based processing in the meaning of GDPR Article 6(1)(f)[4];

[4] Especially given that in such circumstances, valid consent may not be obtained from data subjects who are employees of the data controller or processor (see: Article 29 Working Party Opinion WP249, Section 6.2)

	● The associated right to object under GDPR Article 21 and, where necessary, the ability of the controller to demonstrate a countervailing legitimate interest that is compelling enough to be legally overriding; ● Automated processing and potentially profiling in the meaning of GDPR Article 22, even though the purpose of the processing is to achieve network and information security, and not to produce legal or other significant effects on the data subject, nor to evaluate any personal aspects of theirs, nor to predict their performance at work, economic situation, health or preferences. For all these reasons, even though robust technological solutions do exist, many organisations of lesser privacy maturity and/or with lesser risk appetite consider that the privacy compliance risk involved even in just scratching the surface of this topic is simply not worth the while. The Guidelines expressly recommending the use of such measures to prevent, detect and remedy data breaches is extremely helpful, however most organisations would certainly welcome additional assurances from their supervisory authorities on how, in practice, these measures can be implemented in a proportionate, effective and compliant manner.
● *disabling open cloud services* ● *forbidding and preventing access to known open mail services*	*Prima facie* the definition, the transparent communication and the effective enforcement of such policies seems rather straightforward, and there is no shortage of technological tools to automate, manage and document their implementation. However putting in place any such measures will very often create tensions with other regulatory requirements typically in the area of labour law, whereby employees must be afforded a certain degree of freedom and discretion in the use of corporate resources for personal purposes, including at the workplace and during business hours. Unfortunately, this almost inevitably means creating and maintaining gaps in the organisation’s perimeter defences, which is an additional risk that cannot be fully mitigated and therefore has to be measured, accepted, documented and managed. In our experience, some data leaks happen through disgruntled or ill-meaning employees gaining duly authorised access to business confidential information, and then leveraging their mandatorily permitted “private use allowance” to exfiltrate that data via services which the employer is prohibited from blocking or even monitoring. The issue is further compounded in larger organisations which operate across borders by the fact that national or local requirements even just within the GDPR territory are so different and fragmented that devising and enforcing the same permissible

	use policy, and implementing the same compensating risk mitigations across all jurisdictions is almost impossible. Therefore we encourage the EDPB to conduct further analysis of this issue, possibly involving also employee protection and labour inspection authorities from the member states, and to develop harmonised guidelines that organisations can more confidently rely upon when negotiating the implementation of these technologies with their employees' representatives, as well as when defending the use of such measures in front of their supervisory authorities.

Recommendations

As a large organisation with reasonable privacy maturity and resources, and as a global leading provider of information security technologies, we firmly believe that implementing the measures listed above in full consideration of their privacy implications, and in full compliance with applicable data protection legislation is perfectly feasible, but it does involve some rather heavy lifting at the outset, and significant ongoing maintenance and compliance efforts.

The Guidance is very welcome in that it lends further credence to the importance of implementing technical and organisational measures that can help organisations successfully prevent, promptly detect and adequately respond to data breaches. At the same time, we feel that there remains a significant number of friction areas, typically at the intersection of various guidance materials put forward by different sectoral supervisory authorities at the European and national levels. This Guidance will cover distinct but closely interrelated topics and will result in privacy compliance challenges that are difficult to reconcile especially as regulatory focus on cybersecurity increases.

At the heart of most of these tensions is the question of how to ensure that the processing of personal data necessary to implement effective security measures in the work environment can be done in an accountable manner. In our view, the topics of highest priority about which organisations are the most interested in receiving further guidance and reassurance from supervisory authorities are the following:

1. <u>Legal basis of processing</u>

The recommendations as put forward, and the emphasis on multiple occasions that the examples and proposed measures are by no means comprehensive, give a clear signal that implementing breach prevention, detection and response forms a necessary part of organisations' efforts to comply with the GDPR's requirements, chiefly under Articles 32, 33 and 34.

This is a very welcome development that we would like to strongly endorse. We are comforted by the supervisory authorities' clear statement that in order to protect personal data from breach, such measures may be taken to the extent necessary and proportionate. The logical conclusion that flows from the acceptance that the measures in question stem from requirements of the aforementioned GDPR articles is also that the legal basis for processing personal data in this case should be the organisations' obligations to comply with applicable Union law they are subject to (GDPR Article 6(1)(c)).

Nevertheless we have not seen such a clear statement in the Guidance. In fact we have some concerns that both in WP249 about data processing at work and in WP250rev1 on data breach notification, the Article 29 Working Party stopped short of stating this explicitly, offering as the only practicable option the legal basis of legitimate interest.

This is alarming for multiple reasons. Silence to explicitly recognise that European privacy law creates a legal basis for certain forms of processing permits an "*argumentum a contrario*" that this may indeed not be the case. Suggesting that compliance with European privacy law itself may not be compelling enough to justify relying on Article 6(1)(c) creates a lot of uncertainty for organisations, who find themselves between a rock and a hard place. Organisations must do certain things specifically to comply with applicable law (e.g. Articles 32 and 33 GDPR), and yet they are not permitted to use the relevant legal basis but instead have to rely on a balance of interests test when the real test should be the proportionality (and to a lesser degree the appropriateness) of the implemented measures, and not the necessity or the balance of interests of the data subject, who in some cases may even be the attacker.

Moreover legitimate interest-based processing is subject to certain limitations ("reasonable expectations of the data subjects", "right to object") which can very seriously hinder or even damage the effectiveness of the security measures at stake, making it a particularly tricky area for organisations to navigate, especially as the processing of employee data is concerned. A good example for that is the scenario of data leakage prevention in the case of a hospital. A data leakage prevention tool in a medical environment would likely be required to scan medical data (a special data category). The purpose of the scanning would be to protect the data from leakage and, admittedly, to protect data subjects from serious risk. Yet the legitimate interest clause would be unsuitable in this case because processing of special categories of personal data (in this case conducting protective processing/scanning of medical data to prevent a leak) cannot happen on a legitimate interest legal basis (GDPR Article 9 does not offer that option).

We fully appreciate that Recital 49 establishes that network and information security constitutes a legitimate interest, however we do not believe that the existence of that recital should mean that no other legal basis may be relied upon for security processing. It would be useful to understand whether the supervisory authorities interpret Recital 49 as permitting a reversal of the burden of proof where cybersecurity of the controller is clearly and always recognised as a legitimate interest, provided that the processing happens for what is strictly necessary and proportionate. This is a possible way to read Recital 49 in comparison to the marketing use case, about which GDPR Recital 47 clearly states that it "may be" a legitimate interest unlike the affirmative statement in Recital 49. We would very much appreciate further exchanges among regulators, IT security specialists, as well as labour law and privacy practitioners, with a view to obtaining further clarification and harmonisation from supervisory authorities on this point.

2. Data Protection Impact Assessment

Regardless of the legal basis used, the deployment of pretty much every one of the recommended security measures in isolation, and even more so the implementation of any combination (or the full suite) thereof involves the large scale processing of the personal data of vulnerable data subjects in the meaning of GDPR Article 35 as interpreted by Opinion WP248.

In our view there is no doubt that, even if the regulatory compliance legal basis is relied upon, a DPIA is required for the deployment of any such information security solutions. Therefore a very detailed balance of interests test would always be conducted, in the form of a full DPIA. Accordingly, we do not perceive any risk of organisations skipping the “balance of interests” test by not relying on the legitimate interest legal basis.

Meanwhile it is also important to take into account that the regulatory environment around cybersecurity becomes more crowded by legislative initiatives impacting the horizontal cybersecurity of critical infrastructure, or of particular sectors. Therefore, there will be new legal obligations created by EU law. Eventually, organisations will be required to point to those obligations as a legal basis for relying on GDPR Article 6(1)(c). Such laws may even be outside of the scope of privacy law, or of the protective framework it creates.

We would strongly recommend that supervisory authorities engage in further, specific dialogue with the business community, and develop additional guidance to complement the recommendations in Opinion WP248. In particular, more in-depth dialogue would be helpful on what the regulators’ expectations are in terms of permissible use of security technology, acceptable levels of monitoring, sufficient privacy safeguards, and adequate risk assessment and management documentation. This would be critical not only for organisations to conduct better impact assessments in a particularly sensitive area, but also to be more comfortable when running consultations on this delicate subject with their workforce, and even as the case may be with their supervisory authorities, whether on a voluntary basis or under the requirement of GDPR Article 36.

3. Privacy requirements in the context of employment

For well known and well understood reasons, GDPR Article 88 empowers member states further to specify under national law the rules applicable to the processing of personal data in the context of employment. Additionally, the employment domain itself has a rich history of preexisting labour protection regulations mandating various measures of employee representation, consultation, involvement and co-determination, many of those requirements being highly country specific, and often significantly misaligned between jurisdictions.

Obviously it is not within the remit of the EDBP or of any privacy supervisory authority to resolve such tensions, to arbitrate differences or to harmonise the interpretation and implementation of national, local and sectoral labour regulations. However in our experience an increasingly frequent bone of contention between employers and employee representative bodies is the deployment of information security technologies in the workspace, and the suspected or feared impact that those may have on employee privacy. In some extreme cases, the level of tension can reach such heights where the effective implementation of even the most basic, obviously necessary and reasonable security measures is undermined by the controller's fear of ruining employee relationships or triggering privacy complaints and litigation.

In our view, the new Guidelines are the first clear indication from privacy regulators that a certain level of security monitoring at the workplace is not only permissible, but even actually required under the GDPR. However we believe that further education backed by sound advice from the supervisory authorities would be highly desirable in order to:

- help organisations build more robust cases for the measures they implement, and better document and explain them;

- help allay the concerns of employees through explanations around how accountability, the balance of interests, and privacy risk assessment and mitigation work under the GDPR;

- provide credible and authoritative assurances to both sides about what is necessary, proportionate and permissible, and how the appropriateness of the measures is checked and maintained over time, including through regulatory supervision.

Currently, guidance materials on what data subjects are entitled to under the GDPR are abundant, whereas there is rather scant guidance on what controllers and processors are lawfully permitted to do. As a result, many employer-to-employee discussions in the area of IT security start on the wrong and biased premise that any effort by the employer is necessarily a dangerous threat to employees' rights. This is both unhealthy for the labour relationships, and highly counterproductive for privacy compliance and effective protection of personal data by organisations. We believe that by providing evenly balanced guidance articulating both sides' legitimate rights, interests and obligations, privacy regulators can play a very positive and constructive role in de-fusing such undue tensions. This could promote an employer-employee dialogue that is more conducive to the development, acceptance and implementation of the necessary security measures which are best suited to protecting the organisation's personal data assets, including the data of the employees themselves, in a manner that is privacy friendly, and can enable employees to abide by the data handling policies and practices which they are responsible for observing.

Conclusion

Ultimately, when it comes to breach notification, a key question that is likely going to be litigated in the coming years as a result of the growing number of cybersecurity incidents and data leakages is:

When -- if ever -- would it be acceptable defence from an organisation that a personal data breach occurred because taking the measures that could have prevented it would have disproportionately interfered with employees' or other data subjects' privacy? What would be the risk criteria and the necessity and proportionality considerations that would make such a position defensible if at all?

We believe this is a difficult question that in the end it is bound to be put before a judge and will need to be answered. It is important that every data subject, including employees, have clarity on this as they are the privacy right-holders, but also have an individual (and collective in the case of employees) responsibility when it comes to information security. We fully understand that it is each organisation's responsibility to make the right and accountable decisions appropriate to their specific circumstances. Yet it is important to ensure that organisations have adequate indications of what the acceptable balance looks like in the light of the latest technological developments and risks. As EDPB further considers how to develop its guidance and the areas it further needs to clarify, we think it is critical to avoid the risk that:

- The new Guidelines being interpreted as a mandatory baseline of measures that organisations must implement to comply with the GDPR's data breach provisions;

- Organisations decide not to go beyond what is required in the new Guidelines despite meeting the conditions laid out in Article 32 GDPR and their risk profile justifying such action;

- The substance of the aforementioned Opinions WP248, WP249 and WP250rev1 as currently interpreted resulting in some challenges in the practical implementation of the measures this new Guidance calls for.

We know this is not the intention of the regulators, but we must stress that in our field experience this is a key challenge in which many organisations are finding themselves today when trying to make sense of the requirements facing them. This situation can be further improved and we believe that significant progress can be made through dialogue efforts. We look forward to further discussing this important matter with the EDPB, in the hope of fostering the emergence of a consensus in the shared interests of all stakeholders concerned, and to the ultimate benefit of better data protection and more meaningful privacy safeguards.

++++++++++++++++++

We remain at your disposal to provide additional information. Please feel free to contact:

Ilias Chantzos, LLM, MBA,

Global Privacy Officer and Head of EMEA Government Affairs

Ilias.chantzos@broadcom.com