General Data Protection Regulation: Insight & Recommendations

The response to the fear of technology – why data protection law exists

The General Data Protection Regulation is clearly one of the biggest shake-ups in European Union data protection and privacy legislation in thirty years especially with the added confusion in the U.K of the Brexit vote. However, the earliest versions of European data protection law arise from the Council of Europe, as part of its human rights agenda. It's obvious that these laws were passed in reaction to a fear of the intrusive power of technology. The current legislation, the Data Protection Directive (95/46/EC), was agreed in 1995, long before the mainstream adoption of the Internet and the World Wide Web, and more than a decade before Facebook and Twitter were founded.

The world we work in has changed, but data protection legislation has not. Moreover, the existing rules are based on a directive, the significance of which is that each of the 28 EU member states interprets and implements the rules in its own way. The European Commission set out, in 2012, its Data Protection Principles and the draft text for a new regulation. The regulation will apply equally and simultaneously to all 28 member states on its enactment. After more than three years of discussion and debate (and not an inconsiderable degree of disagreement) the final version of GDPR was agreed in May 2016. The new law comes into force on May 25, 2018.

GDPR Requirements

Data Accuracy, Retention and Erasure

GDPR requires all personal data collected to be gathered lawfully, and for specific purposes only. In addition, it must be used for the purposes for which it was collected, and must be accurate and up-to-date (see Article 5).

This means that companies must have good records relating to personal data, and be able to review its currency and accuracy. This also relates to a company's ability to satisfy a Subject Access Request (see Article 15) — providing all of the data held on a data subject to that data subject, on request — and to affect the rights to rectification and erasure, more popularly known as the right to be forgotten (RTBF, see Articles 16 and 17).

RTBF presents particular challenges from a technology point of view. For example, data controllers must ensure that information relating to a data subject not only on their systems is erased but also that of third-party systems that have copied, replicated or linked to the original information. Additionally, it is conceivable that instances may arise where data is erased on request because it is no longer relevant to the data controller, but which then becomes relevant again in the future. This this may apply, for example, in areas of public interest or in legal proceedings. In these cases, organisations may find that they have to remember data that they were supposed to have forgotten. It is easy to see how organisations could quickly tie themselves in knots in trying to comply with both the letter and the spirit of RTBF.

It is also worthwhile remembering that RTBF and other retention requirements are not absolutes. Companies need to consider whether freedom of information rights could be affected, whether legal or public interest factors apply, and other concerns (Article 17, para. 3).

Consent

All organisations processing personal data must get prior consent from data subjects (unless consent is implied in a contract or processing is required by law or for official purposes). This consent must relate specifically to the purposes of the processing: companies cannot get consent for one purpose and then use the gathered personal data for another (see Articles 7 and 8).
A major change to existing legislation, the revocation of consent must be as easy to effect as the original consent. Many consumers today complain that it is easy to opt in to data gathering, but extremely hard to unsubscribe or opt out. This changes with GDPR. The technology implication is

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015

that the nature and attributes of consent must be recorded, and firms must enable consent to be easily withdrawn (and to confirm that consent has been withdrawn). When obtaining consent, GDPR also introduces mandatory additional collection of parental consent for any data subject below the age of 16 (member state countries have the option of lowering this to 13). Data controllers and processors therefore have the tough challenge of ensuring that all data subjects can provide proof of age.

Data controllers and processors must have an adequate system that requests consent when it is required, facilitates the revocation of that consent in an easy manner and has the ability to determine the age of the data subject. This will require an age verification capability, possibly through the use of credit checking or other identity verification services.

Data Classification

A primary premise of GDPR is that organisations know what personal data they hold. This is not as straightforward as it might seem: GDPR has a particularly broad view of what constitutes personal data, and is particularly sensitive about what it calls special categories — racial, ethnic, political, religious, health, biometric, sexual orientation, etc. (see Article 9).

Personal data itself includes obvious categories (name, identification number, etc.) but also includes location data, physical and physiological information, and mental, economic and cultural data. Profiling, which relates to a person's conduct and behaviour, is also within the scope of GDPR. For example, the fact that an individual liked a particular tweet or Facebook post would constitute personal data.

It is therefore important for organisations to identify the personal data that they hold, and to treat it separately from other data held on its systems (financial, product information, and so on). They must also specifically identify the special categories of data and classify that at a higher level of sensitivity.

Many companies will struggle with such data classification, firstly by identifying all instances of personal data held across all systems they manage, and secondly by classifying correctly the variety of personal data types. Several automated data classification technologies are now available on the market, but the largest part of this exercise is currently done manually.

Data Portability

GDPR mandates that data subjects can request that their data is provided to them or a third party in a "structured, commonly used and machine-readable format" (see Article 20).

For data held in structured forms, for example in relational databases, this should be relatively straightforward for companies to achieve. Increasingly, however, data is held in unstructured formats. For example, data held in social media sites will be unstructured, as will data held in audio and video formats. Video in particular is notorious for being incompatible across formats, so companies will be challenged to provide information in a variety of ways that can be utilised by the data subject or a third-party controller.

For structured information we expect text, comma-separated and tab-delimited formats to be popular. For audio, mp3 is most likely. Video standards are nightmarish: Audio Video Interleave (.avi), Advanced Systems Format (.asf), QuickTime (.mov or .qt), Flash Video (.flv, .swf), MPEG-4 (.M4V and .MP4), plus various codecs and high-definition variants and emerging formats. Video converters abound but many are bug-ridden or handle only specific format-to-format conversions.

Compliance Auditing and Record Keeping

Companies processing personal data are obliged to keep detailed records of the data they hold, as well as the details of the processing conducted on that data. For example, information regarding a data transfer to a third country would be important to maintain (see Article 30).

Interestingly, the keeping of records is the only instance in GDPR where company size affects the obligations on a data controller. An organisation employing fewer than 250 people is exempt from keeping records (unless it processes personal data classified as high risk).

More importantly, there are indications in GDPR that good record keeping will be considered as a mitigating circumstance in the case of a data breach (see Article 83). This may reduce or eliminate an otherwise sizeable fine.

Auditing of activity relating to personal data is critical for GDPR, and while this is common in data life-cycle management products it becomes compliance-critical. It also gains importance in identity and access management (IAM) solutions: the ability to track who has access, to what, why and for how long. IAM auditing is relevant both as a preventative device and also in forensics post-breach.

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015

Data Residency and Data Transfers

Much of the focus on data protection in Europe has centred on the transfer of data from the EU to what GDPR calls "third countries" — that is countries that are not member states (see Article 45). The attention on data transfers emanates from two related situations: the high volume of data transferred outside the EU due to Internet-based services (particularly to the US) and the revelations by Edward Snowden that much of this data has been subject to surveillance by government agencies, particularly in the US.

Data transfers have increased rapidly in the past eight years due to the rise in social media and the adoption of cloud services. The current EU Data Protection Directive allows transfers only to third countries that demonstrate "equivalent" data protection laws: importantly the US is not one of those countries. (The countries currently assessed as having equivalent data protection laws are Andorra, Argentina, Canada [commercial organisations], Faroe Islands, Guernsey, Israel, the Isle of Man, Jersey, New Zealand, Switzerland and Uruguay. Data transfers may also be permitted where a data subject has explicitly provided informed consent to the transfer; see Article 49.)

Instead we have an agreement between the EU and the US, where US companies agree on an individual basis to adhere to EU data protection standards (see Article 46). This agreement "Safe Harbour" was struck down by the EU Court of Justice in October 2015 because of the Snowden revelations. Its replacement the EU-US Privacy Shield has not yet been formally adopted, and may in any case suffer the same fate as Safe Harbour.

GDPR essentially maintains the current position on data transfers. The concept of equivalence is retained: however, because GDPR raises the bar, a reassessment of the countries currently regarded as equivalent is likely. Importantly, GDPR evens up differences in data transfer law across EU countries by applying equally and evenly across all member states.

There are important technology implications from the data transfer rules in GDPR. The main one is the requirement to know where personal data is at all times. Crucially, this includes not only the data itself but also metadata and indexing information, which is often held in separate data stores. It is also quite common in cloud services to move data from its normal location at rest to a different location for processing. For example, data held at rest in the EU may be transferred to the US for processing. Ironically, this processing may include the encryption and decryption of the personal data, encryption being an important data protection technology called out in GDPR.

Securing Data

As stated earlier, GDPR is particularly light in its prescription of security requirements. It does however call out one set of technologies as being an example of appropriate technology, namely **encryption** and **pseudonymisation** (see Article 32, section 1 [A]). The implication is that data held in an encrypted or pseudonymised form is not deemed to be personal data. There is a clear direction (though not a mandate) here: encrypt or pseudonymous personal data.

Encryption and pseudonymisation, though, come with several health warnings.

- Firstly, it is conceivable that data encrypted and considered secured using today's technology may become readable in the future with the availability of more advanced decryption capability.
- Secondly, managing encryption keys is a serious exercise, with the consequences of mismanaged or (worse) lost keys potentially catastrophic to a business.
- Thirdly, encrypting personal data may mean that useful processing of that data may no longer be possible. For example, the ability to search, index and correlate personal data may be impossible with encrypted data.
- Fourthly, pseudonymisation, which removes identifiable attributes of personal data, can be reversed through correlation with other third-party data. The EU is particularly concerned at this re-identification possibility (see recital 26).

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015

GDPR: Technology

GDPR is far more explicit than its predecessor, the Data Protection Directive when it comes to its focus on the role of technology. If it is to be properly effective, however, the GDPR must assist in the delivery of business transformation and legal compliance. It does this in a number of ways. It requires the use of Privacy by Design techniques and the performance of risk assessments. It also identifies data management techniques, such as data mapping, and techniques for how to handle operational failure, such as breach disclosure.

Key Technology goals

Driving data protection principles into technology, through appropriate technical and organisational measures

When developing technical and organisational measures, organisations must have full regard to the 'nature, scope, context and purposes of processing' and 'the risks of varying likelihood and severity for the rights and freedoms of natural persons'

Adopting a proper approach to technology design and deployment

GDPR's helpfully includes a set of requirements that provide organisations with practical assistance in how to flow data protection into technology.

These are:

- Accountability;
- Records of processing activities;
- Data protection by design and default;
- Data protection impact assessments;
- Breach notification.

Collectively, these new requirements provide a 'user manual' for delivering operational success.

Ensuring the technology environment can protect individuals' rights

The core individual rights are the 'right of access', 'right to rectification', 'right to erasure' (or the 'right to be forgotten'), 'right to restriction of processing', 'right to data portability' and 'right to object'.

In a functional sense, these rights require the technology to

- Connect individuals to their personal data;
- Categorise personal data by type and processing purpose;
- Map or trace the full information lifecycle;
- Perform search and retrieval;
- Enable rectification, redaction, erasure and anonymisation;
- Enable freeze and suppression;
- Enable the transmission of personal data from one technology stack to another.

All of this must be protected by appropriate security

GDPR and Security

Much of GDPR is about process and the operational aspects of data protection. However, some elements can only be enabled by technology, and others are made more manageable or cost effective through technology.

Whilst GDPR is highly generic when it comes to technologies to be deployed, and in the consideration of security GDPR has very little to say (see Article 32). It only suggests: "The pseudonymisation and encryption of data; ability to ensure confidentiality, integrity, availability and resilience of processing; the ability to restore data after an incident; and a process for testing, assessing and evaluating effectiveness of security".

Most of this is already standard security practice. Confusingly, GDPR introduces the concept of "state of the art". This is a term which is designed to render GDPR future-proof (the criticism of previous legislation was that it became quickly out of date due to technology innovation). While this position is understandable it does make interpretation difficult: what is the state of the art, who decides, and is one company's view of this arguable in a court of law?

The Oncore IT attitude to GDPR looks beyond compliance, seeing it not just as an obligation but seeing GDPR as a catalyst for improved maturity in Data Privacy and Security

Article 24 (1) – Responsibility of the controller Taking into account the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for the rights and freedoms of natural persons, the controller shall implement appropriate technical and organisational measures to ensure and to be able to demonstrate that processing is performed in accordance with this Regulation. Those measures shall be reviewed and updated where necessary

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015

Recommendations

Boundary Security

This is defined by hardware or software based solutions that protect your network from connections or to external sources, such as the Internet, Email or cloud services

Enable filtering of connections to and from the internal network. Running the IPS Software pack provides additional benefits of incident reporting and intrusion detection.

Firewalls rely heavily on professional configuration and adhering to security policies to make sure that external access in to the internal network is restricted to a minimal level. The same is also applied to traffic going from the internal network to the internet

Our Recommendation

- Juniper SRX Firewall
- IPS Software pack
- Managed configuration

Email Security

Users should be protected from SPAM, Malware and Viruses by ensuring that suspect messages are stopped at the boundary and never reach the end users inbox.

Additional security measure can also be taken against URL`s that are linked to Internet services know to contain Malware. Oncore Email Defence blocks over 60% of all emails at the edge, meaning that 60% emails are not ever transmitted. It then rejects a further 20% of messages following security scans.

Our Recommendation

- Oncore Email Defence featuring
- Anti-SPAM
- Outbreak Filters
- Grey Mail Filters
- URL re write
- Macro Detect

Encryption using Two Factor Authentication

Users connecting remotely to internal systems have to authenticate using a username and password.

Providing passwords meet complexity requirements this is quite secure. However, if a user's credentials are compromised, there is no form of defence and an attacker would have unfettered access to the internal network.

Adding a second authentication level requires the user to be in possession of an encrypted token that is unique to the user. This is provided as an application on the user's phone. When a user enters their credentials they then have to authorise the request by responding to a message or typing a code which is provided by the app on the phone.

Our Recommendation

- Duo 2FA Access
- RD Gateway 2 FA
- Workstation Logon 2FA
- Device Insight policies
- Open VPN with 2FA

Server & Workstation Security

Even with all of the security at the boundary, user error, files from USB disks and browsing of legitimate websites can still pose a threat.
Anti-Virus: Anti-Virus (AV) should be deployed on workstations, servers. AV will protect from all know virus and Malware attacks. However, AV solutions rely heavily on being up to date with that latest virus definitions to be effective. Oncore AV Defender not only provides a remote deployment and managed service, it monitors for new definitions and deploys them instantly. In addition, a new Ransomware Vaccine will soon be available to provide added protection from known Ransomware.

Our Recommendation

- AV Defender (Managed)
- Ransomware Vaccine

Web Filtering

Traditionally this is deployed at the boundary, however this technology is dated and does not protect against zero hour vulnerabilities (Where threats are emerging and filters are not aware of the threat).

A Web filter allows an enterprise or individual user to block out pages from Web sites that are likely to include objectionable advertising, pornographic content, spyware, viruses and other objectionable content. No longer is it possible to maintain databases of potentially or known suspect websites. Therefore, Web Filtering is now better deployed at the network and workstation level and has become a hybrid component.

Our Recommendation

- Cisco Cloud Umbrella
- Workstation Agent
- Network proxy

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015

Data Residency, Storage and Archive

Unstructured data growth, application proliferation and increasing variation in data types continue to accelerate. This leads to increased server and storage sprawl with numerous silos of infrastructure supporting different workloads.

Our storage platform provides a massively scalable, multitenant object storage environment. It provides the security and broad support needed to enable organizations to gain the benefits of archiving, backup reduction, cloud, consolidation and more without compromising security or completely rewriting critical applications.

Our Recommendation

HCP (Hitachi Content Platform) Private, Public or Hybrid Deployment

Data Protection

- Leverage advanced replication, erasure coding and edge device integration to ensure collection and protection of content.
- Select the right level of protection with dynamic data protection levels.
- Locate the right information with ease thanks to metadata query tools, versioning and full content search.
- Move data protection or deep archive copies to spin-down disk or make tape copies if desired.

Content Preservation

- Store, retrieve and manage the lifecycle of data for long-term storage for corporate governance
- Protect and secure content for long-term preservation; content is continually checked throughout its retention period for integrity.
- Aggregate content from a variety of sources, including file servers, email and collaboration tools such as Microsoft SharePoint.
- Grow business horizontally to support multiple applications and content types and scale vertically to support continued data growth.

Content Distribution

- Connect multiple, distributed sites to a centralized content repository.
- Share content from one edge location to another via Hitachi Data Ingestor.
- Control data placement and distribute content to appropriate audiences.
- Provide bottomless, backup-free storage to branch and remote offices via Hitachi Data Ingestor.

Cloud Deployment

- Employ a single, multipurpose, unstructured data platform for archive, cloud and backup capabilities.
- Monitor and report on storage and bandwidth use for chargeback.
- Employ user tools as well as application and management interfaces for cloud and distributed environments.
- Service providers can offer bottomless, backup-free storage to customers via Hitachi Data Ingestor.

HDI (Hitachi Data Ingestor) Physical or Virtual Deployment

Hitachi Data Ingestor provides a standard connection, or on-ramp, into the core data centre without requiring application recoding and without changing the way users interact with storage today. Because HDI acts as a caching device, it provides users and applications with seemingly endless storage and a host of newly available capabilities.

Hitachi Data Ingestor presents a standards based file system interface that is tightly integrated with Hitachi Content Platform to provide seamless access and a wide range of advanced storage features.

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015

HDI uses Reduce Cost and Complexity at the Edge and Simplify Cloud Deployments. Hitachi Data Ingestor use HTTP/HTTPS to securely move data over a local or wide area network and into HCP.

- Provides local and remote access to a HCP for clients over CIFS and NFS
- Delivers seemingly bottomless storage capacity, back-ended by HCP
- Migrates content to a central HCP and maintains a local link to the migrated content
- Provides file restore
- Allows content sharing between HDI systems
- Provides a management API that enables integration with the HCP management user interface (UI) and 3rdparty or home-grown management user interfaces
- Supports Active Directory and LDAP authentication for HCP clients
- Scales to 400 million files per HDI
- Employs intelligent local cache to accelerate access to HCP content over CIFS and NFS

Unstructured data growth, application proliferation and increasing variation in data types continue to accelerate. This leads to increased server and storage sprawl with numerous silos of infrastructure supporting different workloads. Our storage platform provides a massively scalable, multitenant object storage environment. It provides the security and broad support needed to enable organizations to gain the benefits of archiving, backup reduction, cloud, consolidation and more without compromising security or completely rewriting critical applications.

Most significantly, our platform eliminates the need for a siloed approach to storing unstructured content. Thanks to its massive scale, multiple storage tiers, Oncore reliability, cloud capabilities, multitenancy and configurable attributes for each tenant, the platform can support a wide range of applications on a single physical cluster. By dividing the physical cluster into multiple, uniquely configured tenants, administrators create “virtual content platforms” that can be further subdivided into tens of thousands of namespaces for additional organization of content, policies and access control.

Conclusions

Companies should use GDPR as a cornerstone for a risk mitigation process. For example, GDPR specifically states that the conditions for imposing fines depend upon (among other things):

- Intentional or negligent non-compliance
- Action taken to mitigate the damage suffered by the data subject
- Previous infringements
- Degree of cooperation of the supervisory authority
- The manner and extent to which a data breach was notified
- Adherence to codes of conduct
- Any aggravating or mitigating factor (see Article 83)

In other words, actions and evidence count toward mitigation: trying hard to comply with GDPR is important. Understanding the risks, and mitigating factors, of data protection should be a part of everyday business.

Paul Cook Sales Team
Oncore IT
5th Floor
1 Portsoken Street
London
E1 8BT

T: 020 3818 3411
E: pcook@onoreit.com

1 Source: IDC, FutureScape, 2016 2 Source: Uptime Institute/451 Group, June 2016 3 Source: IDG Enterprise Computing Study, 2015