
D: D-4.12 A4Cloud Tools, Liability and Compliance Investigations

Deliverable Number:	D44.1a
Work Package:	WP 44
Version:	Final
Deliverable Lead Organisation:	QMUL
Dissemination Level:	PU
Contractual Date of Delivery (release):	31/05/2015
Date of Delivery:	30/05/2015

Editor

Chris Reed (QMUL)

Contributors

Chris Reed (QMUL), Asma Vranaki (QMUL), Lorna Cropper (QMUL), Petra Zabudkova (QMUL), Lorenzo Dalla Corte (TiU)

Reviewers

Frederic Gittler (HP), Tobias Pulls (KAU)

Executive Summary

This Deliverable focuses on the A4Cloud accountability tools. It analyses the legal and regulatory consequences of supplying the tools to users, and also the consequences to those who use them. This analysis concentrates on liability and data protection compliance. Where there are differences in the liability approach of the two European legal traditions, common law and civil law, it illustrates the consequences of these differences by reference to English and Italian law as suitable representatives of the two traditions.

Liability is determined by the relationship between claimant and defendant. There are few court decisions specifically deciding how the law's liability principles are to be applied to software and information, and so much of the analysis is an attempt to predict how a court in the future might decide, using analogies from decisions in other fields.

Where the tool is purchased as software, its supplier will have potential contractual liability to the purchaser if the tool is (in broad summary) unsuitable for its intended purposes. Where use of the tool is supplied as a service, the supplier will be liable to the purchaser of the service if the supplier did not use proper care and skill in selecting and implementing the tool, and possibly also if the service is unsuitable for its intended purposes. The common law and civil law take different routes to this conclusion, but are broadly in agreement.

Producers of software and suppliers of services using the tools will also have potential non-contractual liability to those who rely on the correct working of the tools and on their outputs. This liability arises if there is insufficient care and skill in the design and production of the tools, and in their implementation and operation. Common law imposes this liability only if there is such a close relationship that the defendant appears to have undertaken responsibility to the person relying on the tool, whereas civil law may impose liability in respect of any person who was apparently intended to rely on the tool.

In deciding whether the tools are fit for purpose, and whether sufficient care and skill were taken, the courts are likely to take account of two factors over which tool producers and users have some control. Compliance with industry best practice is likely to be seen as strong evidence that reasonable care and skill has been exercised. More importantly, the information given to purchasers and users about the limitations of the tools will help determine the intended purposes of the tools and provide some assistance in assessing fitness. This information is also an important way of discharging any duty to use care and skill, in part at least.

Data protection law imposes absolute obligations on data controllers, with the possible exception of the security obligation, and so using accountability tools will not normally excuse a data controller for any breaches of data protection law. However, we suggest that if tools are implemented in good faith, using reasonable care and skill to ensure that they work to prevent breaches, this is likely to influence the enforcement action taken in respect of a breach, such as reducing the level of any fine.

Data Protection Authorities are increasingly conducting investigations into the data protection compliance of cloud service providers, and will be interested in the tools during this process. We suggest that investigations will focus on those elements of the A4Cloud tools which control processing (primarily the A-PPL Engine) and which log and take action in response to policy breaches (AAS, IMT/RRT, DTMT, Transparency Log). An investigation will examine whether the tools have been implemented properly in the data controller's system, and will also potentially include checks on the logic and the code implemented by the tools. This raises the question how such access to the internal working of the tools might be provided without compromising commercial confidentiality.

Data protection investigations are often ongoing, and so need to examine real-time compliance of data controllers' systems. The A4Cloud tools offer some of this functionality, and so investigators may want to link their analytical tools directly to the A4 tools. They may also seek remote access to the tools for

future assessment of compliance. Again, how to provide these facilities without compromising the integrity and confidentiality of the tools is an important question. The increasing reliance on independent internal and external auditors widens the group of those who will seek access.

The Deliverable concludes with some thoughts about taking the A4Cloud accountability tools to market. Because of potential liability concerns, those tools which might be seen as offering analysis and advice (COAT, DPIAT, possibly AccLab) could be unsuitable for normal software distribution and supply techniques, particularly to the intended market of SMEs. One way of managing liability risk is to supply the tools only through third party solutions or consultancy providers, or to supply only to experienced users such as cloud service providers. An alternative is to market the tools as suitable for use only by trained persons, in which case thought would need to be given to how that training might be provided.

Perhaps the most important finding is that for all the tools, liability can be to some extent managed through a high degree of transparency about the methods of using them and about their limitations. Achieving that transparency in the SME market is, however, challenging.

Table of Contents

1	Aims	6
2	Regulating people not technology.....	6
3	Methodology.....	8
4	Legal and regulatory relationships	8
5	Relationship between cloud customer/provider of tools or services	9
5.1	Obligations of the provider	9
5.1.1	Scenario 1 – running accountability tools on a cloud platform.....	10
5.1.2	Scenario 2 – use of accountability tool in a SaaS relationship	12
5.1.3	Scenario 3 – provision of an accountability service	18
5.1.4	Express contractual terms.....	19
5.2	Assessing reasonable care and skill.....	20
5.2.1	Tool producers	21
5.2.2	Tool users.....	23
5.3	Conclusions on liability.....	24
6	Relationship between data controller/DPA	25
6.1	Accountability tools as a means of compliance	25
6.1.1	Responding to a complaint.....	26
6.1.2	DPA-initiated investigations	27
6.1.3	Good faith and reasonableness	28
6.2	DPA’s use of outputs from tools during investigations.....	28
7	The A4Cloud accountability tools.....	31
7.1	Cloud Offering Advisory Tool (COAT).....	32
7.1.1	How COAT works.....	32
7.1.2	Legal and regulatory issues	34
7.2	Data Track.....	35
7.2.1	How Data Track works	35
7.2.2	Legal and regulatory issues	37
7.3	Data Protection Impact Assessment Tool (DPIAT).....	37
7.3.1	Legal and regulatory issues	38
7.4	AccLab	39
7.4.1	How AccLab works.....	39
7.4.2	Legal and regulatory issues	40
7.5	A-PPL Engine.....	41
7.5.1	Legal and regulatory issues	42

7.6	Transparency Log	43
7.6.1	How the Transparency Log works	43
7.6.2	Legal and regulatory issues	44
7.7	Experimental status of the A4Cloud tools	44
8	Business models to take the tools to market	45

Table of figures

Figure 1:	Tool functions	31
Figure 2:	COAT questions	33
Figure 3:	COAT pop-up explanation boxes	34

1 Aims

The original aim of this deliverable, as set out in the revised description of work, was to:

... analyse the implications of the findings in D: D-4.1.1 and MS D-4.4 for the A4 Cloud tools. Questions to be investigated will include: Are the outputs of these tools appropriate for use by regulators in their audit process, and if not how should they be modified to interface properly with regulatory compliance processes? What will audit mean for these tools – for example, will a regulatory compliance audit require an investigation into the coding of the tools, or the processes they implement, or will it treat them as ‘boxes’ whose output is simply a given for the regulatory audit process?

This work is undertaken in Section 6 below. It builds on previous deliverables D: D-4.1.1 and MS D-4.4, which researched how Data Protection Authorities use their investigatory powers when reacting to complaints in respect of a cloud service provider’s data processing activities, or deciding to undertake an investigation on its own volition.

Additionally, in the course of discussions with those research partners who are developing the tools, it became clear that it would be even more important, both to them and to other users of the tools, for there to be an analysis of the potential liabilities of tool creators and users. Thus a second aim was developed for this deliverable, to undertake that analysis – see Section 5 below.

Section 0 attempts to apply these analyses to the A4Cloud accountability tools.

2 Regulating people not technology

Before commencing the analysis, we think it important to emphasise that law and regulation never apply to technology directly. Instead, law and regulation are addressed to people (including legal persons such as corporations) and instruct them how they should behave. What is commonly referred to as technology regulation is in fact the regulation of how those people *use* technology. The first analytical step is therefore to decide which people we are interested in.

For the purposes of this deliverable we have identified four categories of persons whose legal and regulatory rights and responsibilities merit investigation:

- **Producers and/or suppliers of cloud accountability tools.** This group includes tool developers, cloud service providers, consultants who recommend or supply technology, etc. This group is interested in the legal and regulatory use of tools because it will wish to ensure that the tools are suitable for such uses. There is also potential liability for producing or supplying tools which are unsuitable.
- **Tool users.** The obvious example is a cloud customer processing personal data or confidential information, who will be using the tools as part of its system for complying with its legal and regulatory obligations. Cloud service providers might use the tools to provide services to cloud customers, and customers will use those services as part of their compliance systems. Data subjects and “owners” of information may also use those tools to check on compliance by other cloud actors.
- **External reviewers of tool outputs, such as data protection regulators and auditors.** This group is interested in the accuracy and completeness of tool outputs, particularly outputs recorded in logs. One of their roles is to check the legal and regulatory compliance, and they will use the tool outputs for these purposes.
- **Non-users who are owed legal duties by tool producers, suppliers and users.** This group, too, is interested in legal and regulatory compliance and the role which tools play in

achieving that. It is also potentially interested in the tool outputs, which will be important evidence if legal proceedings are undertaken. The most likely source of that evidence will be logs of the tool's activities.

In terms of accountability, the first of these groups is neither giving nor receiving an account. Rather, its members are providing technologies which will be used, or whose outputs will be used, by members of the other groups in their roles as accountant or accountee.

Tool users will be accountors if they are using the tool to provide an account of their data processing practices to an accountee. A typical example might be a data controller responding to a data subject access request using a tool to generate part of the response, or that data controller's cloud service provider using the tool to provide the controller with the information it needs to respond to the request. Alternatively, a tool user might be providing a service to an accountant by means of the tool, eg a cloud service provider which gives effect to its customer's data processing policies by using an accountability tool. Here the provider is offering both a service (processing of data) and an account (what was done with that data) by means of the tool.

The third and fourth groups, external reviewers and those non-users who are owed legal duties, will be accountees. Their accountant may use the tool to produce the account, or the accountee may use the tool to generate the account for itself.

The analysis in this deliverable does not focus on the accountability relationships, because that would tell us nothing about the legal aspects of tool supply and use. If there is a legal obligation to give an account, the law is concerned solely with whether that account was given, and if so whether its content was complete or accurate. The means used to generate the account (ie the tools) are relevant only to the legal penalties for failure (see section 6.1 for a fuller explanation).

For our liability analysis we need instead to focus on the other relationships which are created by supply and use of the tools. There are three relationships which are of interest here:

- The relationship between the supplier of a tool and its purchaser and/or user;
- The relationship between the supplier of a service (which incorporates tool use) and the recipient of the service.
- The relationship between the provider of information (by means of a tool) and the person who relies on the accuracy and completeness of that information.

In the supplier/purchaser relationship, the law's focus is about the promises the supplier has made about the tool's performance. But the technological functions which accountability tools perform are not directly relevant to the liability relationship between supplier and a non-purchaser user, or between suppliers and recipients of a service or between providers of information and relying parties. Here the law imposes obligations to behave in a particular way, and I therefore mainly interested in how humans *use* the tools. For this reason we need to identify at a human level the relevant functions of the tools.

The first of these functions is to provide information, and possibly even advice, which humans will use in conducting their activities. The law is interested in whether this information or advice was produced and used properly, and is also interested in whether it should have been used at all for a particular purpose.

The second function, and the one which technologists would think of first, is to control data operations. Law and regulation in this field concerns itself mainly with the use and disclosure of information, and in the cloud all this is achieved by performing operations on data. The focus here will be on whether the tools perform well enough to allow the user to comply with legal obligations, and also on how far using the tools can substitute for human decision-making about use and disclosure.

The third function is logging and reporting what has happened to data. This information is clearly of interest to regulators who are making assessments of compliance, and to courts which are deciding questions of liability.

3 Methodology

The methodology used to undertake this research is a standard, four-stage legal analysis. Stage I identifies those people, hereafter referred to as actors, who are to be included in the analysis. Section 2 above makes that selection.

Stage II identifies relationships between actors. Law and regulation imposes obligations on actors, but those obligations do not exist in a vacuum. Each obligation is owed to some other actor, and the nature of the relationship between that pair of actors determines whether obligations exist and the content of those obligations. The number of potential relationships in our field of study is enormous because each of the four categories identified in Section 2 contains a wide variety of different types of actor. In addition, the nature of the relationship often derives from the interactions between actors, which adds further complication. It is therefore impossible in the space available to conduct an exhaustive analysis of all potential relationships. Instead, we have chosen a subset of relationships which we believe raise the most interesting legal questions and have the potential to illuminate other relationships.

Stage III identifies the legal obligations which arise from those relationships, including variations from their context (such as the communications between them, the agreements they have undertaken, etc).

Stage IV attempts to predict the legal and regulatory outcomes arising from use of the tools.

Stages I and II are conducted at a global level. Once we move to detailed analysis of legal obligations, though, it is more helpful to undertake stages III and IV together for each relationship.

The results of this analysis are applied in Section 0 to a selection of the A4Cloud accountability tools.

It is important that our analysis takes proper account of the legal diversity within Europe. EU law in the form of directives is implemented as national law, and generally applicable EU law such as regulations is interpreted by national courts. The result is similarity, rather than uniformity. Outside the harmonizing sphere of EU law national differences are more marked, particularly in the law relating to non-contractual liability.

To illustrate these differences, we have examined the liability issues from the perspective of two countries, each representing one of the major legal traditions within the EU. The common law tradition (consisting of the UK's jurisdictions¹ and Ireland) finds its liability law from the decisions in previous cases, whereas the civil law tradition (the other EU member states) finds its liability law in a written code and (in theory at least) cases are not sources of law but merely guidance as to how other judges have interpreted the code.

We decided to use English law as our primary example for common law, with occasional reference to US decisions where these provided helpful illustrations, and Italian law to represent the civil law approach. We had access to relevant subject expertise in each national law and our choice would accurately illustrate some of the divergences between the two traditions.

4 Legal and regulatory relationships

In choosing which relationships to analyse in depth, we had two aims. The first was to choose relationships in which use of the tools would play an important part. The second was that each of the

¹ Within the UK, England and Wales (usually shortened to England or English law) and Northern Ireland are both common law jurisdictions, while Scotland is a mixed common law and Roman law jurisdiction.

chosen relationships should involve all the most important legal issues which are likely to arise in other relationships, so that our analysis here might be useful in helping to understand those other relationships.

On this basis we chose two relationships for in-depth study:

- The relationship between a cloud customer and its supplier of accountability tools, or the supplier of a service which incorporates use of accountability tools; and
- The relationship between a data controller which uses accountability tools and its Data Protection Authority (DPA), the national regulator responsible for overseeing legal and regulatory compliance with data protection obligations.

In the cloud customer and tool/service provider relationship, the role of the tools/service is twofold. First, they provide information or advice about what is happening to data, and the customer makes decisions in part relying on the accuracy and completeness of that information or advice. Second, they control some aspects of data processing operations, and again there is reliance by the customer, this time on the correctness of the tools' functioning.

For accountability tools to be of any use to the customer they need to be reliable. But no technology is ever 100% reliable. The tool/service provider therefore needs to know what level of reliability is necessary to meet the obligations which the law imposes in its relationship with the customer. The customer wants to know whether, if the tools or service do not perform exactly as expected, it has some legal remedy against the provider. To answer these questions we need to investigate the law of contract and tortious liability for negligence.

There are two elements of the data controller/DPA relationship in which use of accountability tools or a tool-based service play an important legal role. The first is their use as a means of compliance – to what extent will the fact that a data controller used a tool or service, rather than relying on human observation and action, be accepted as complying with law and regulation? The second is the role of tool outputs, as recorded and logged, in investigations by DPAs of alleged breaches of law and regulation. Can DPAs demand to review those outputs, and if so, how will they be used in the investigatory process?

5 Relationship between cloud customer/provider of tools or services

5.1 Obligations of the provider

The obligations of the provider to the cloud customer are in theory different depending on whether or not the tools or service are provided under a contract. If not, those obligations have to be found from the law of tort or delict. In practice though, as we shall see, there is very little practical difference between the obligations which would be imposed in contract as opposed to tort. The main difference lies in the fact that contracts can be used to accept additional obligations or to exclude or limit those obligations which the law would otherwise insert into the contract. This is discussed in Section 5.1.4.

It is worth reiterating here the point already made in section 2, that these obligations derive from the supply/service nature of the relationship and not from the accountability functions of the parties. The most likely reason for the customer to enter into that relationship is so that the customer can improve compliance with its accountability obligations, but if an accountability tool performs other functions these obligations will not cease to apply to those other uses of the tool.

It will be helpful to examine three scenarios which demonstrate different types of provider/customer relationship:

- The cloud customer runs its own systems on a platform provided by a cloud service provider, and purchases a licence from tool producers to run one or more accountability tools on that platform. In this case the tool producer supplies the tool to the customer, and the relationship with the customer is contractual. We will assume that the cloud customer is a business; consumer activity in this area is unlikely.
- The cloud customer purchases the use of an accountability tool as a service from its cloud service provider. This is a SaaS relationship. The immediate provider of the tool is the cloud service provider, and its relationship with the customer is contractual. However, the ultimate provider of the tool is still its producer, and because there is no contract between the customer and the tool producer, any obligation owed by the producer must be a tortious one.
- The cloud customer purchases an accountability service, and the service provider uses accountability tools to provide the service. In this example there is a contractual relationship with the service provider for provision of the service. The customer never has use of the tools directly, and so we think there is no legal basis on which to impose any obligations towards the customer on the part of the tool producer.

5.1.1 Scenario 1 – running accountability tools on a cloud platform

In this scenario we have a contractual relationship between customer and provider, under which a digital product (software) is being supplied. What obligations does the law impose on the supplier?

Common law jurisdictions have had real difficulty deciding how to treat software for liability purposes. English law essentially divides the subject matter of commerce into goods and services, but software does not fit easily into either category.

So far as consumer contracts are concerned, this problem has been resolved in the UK by introducing a third liability regime based on the law relating to goods. If the consumer has paid for software then under the UK Consumer Rights Act 2015 the software is “digital content”² and accordingly must be of satisfactory quality, fit for any purpose the consumer buyer has made known to the seller, and correspond with its description.³ These rights apply if the consumer has paid directly for the software, or if it is supplied in conjunction with some other product (including digital content) or service which the consumer has paid for.⁴ However, most common law jurisdictions do not yet have specific legislation on this point.

The liability position is much more uncertain if the purchaser of software is a business. Goods are defined in s 61 Sale of Goods Act 1979 as personal chattels. This requires them to possess some tangible or corporeal element, so it is clear that pure information cannot be goods.⁵ In this scenario the tool software will be supplied as a download, so no physical property will be supplied. From this we might conclude that the supply must be that of a service, and if so the main obligation the supplier would have is to exercise reasonable care and skill⁶ in supplying that service (see Section 5.1.2 below).

However, we think it likely that the English courts will impose additional liability obligations, based on the commercial similarities between the supply of software and the supply of goods. In *St Albans City*

² Defined in Consumer Rights Act 2015 s 2(9) as “data which are produced and supplied in digital form”.

³ *Ibid*, ss 34-36.

⁴ *Ibid*, s 33.

⁵ *Oxford v Moss* (1978) 68 Cr App Rep 183.

⁶ This is very similar to the concept of “due diligence” as used in the technology industry (in legal circles, “due diligence” has a technical meaning in corporate takeovers, and so we use the standard legal terminology of reasonable care and skill). However, some technologists might consider due diligence to extend no further than complying with industry best practice, and as we shall see, sometimes a whole industry’s practices may be lacking in reasonable care and skill.

*and District Council v International Computers Ltd*⁷ the dispute was over the supply of an expensive software package, supplied purely in digital form. The software was admittedly defective, in that it did not perform the functions it was designed for. In the Court of Appeal Sir Ian Glidewell held that the contract was not a sale of goods, because property in the tangible medium of a disk had not been transferred, but that nevertheless:

In the absence of any express term to the contrary, such a contract is subject to an implied term that the program will be reasonably fit for, i.e., reasonably capable of achieving, the intended purpose.

The distinction here is between a matter which has been agreed and whose wording has been recorded, usually in writing (an express term), as compared to a situation where the conduct of the parties shows that they agreed on some matter, even though they did not record that agreement (an implied term). Common law courts are always open to implying terms⁸ when the parties act as if they have agreed on the matter in question.

Sir Ian Glidewell's statement was cited with approval in *Sam Business Systems Limited v Hedley and Company*,⁹ but the point has not subsequently been raised in other cases. The introduction of consumer rights of satisfactory quality and fitness for purpose are likely, in our view, to persuade the UK courts to imply similar rights into B2B contracts unless the express terms cover the matter.

The English courts are likely to be confirmed in this approach by EU law. The decision of the ECJ in *Usedsoft GmbH v Oracle International Corp*¹⁰ supports the suggestion that the supply of software in electronic form transfers ownership of that copy, and should therefore be treated at least as analogous to the supply of goods. There is further support in the Consumer Rights Directive,¹¹ which defines digital content as "data which are produced and supplied in digital form" and applies the same consumer protections to its supply as it does to goods.

Our conclusion is therefore that at common law, the tool provider will have an obligation to take reasonable care and skill in its supply, and normally also an obligation that the tool will be reasonably fit for its intended purpose and, if the tool is sold to a consumer user, that it is of satisfactory quality. Satisfactory quality means fitness for all the common purposes for which the software is bought.¹²

Care and skill is examined in Section 5.2 below, but how is fitness for purpose to be assessed?

Under English law the fitness for purpose obligation derives from s 14(3) Sale of Goods Act 1979, under which goods must be reasonably fit for the purposes the buyer has made known to the seller. The problem here is that the buyer will rarely express a clear purpose, and so the courts will have to examine the context of the transaction. In our view the most important factor will be the claims the tool producer has made in its advertising and promotional material. This may amount to a description of the goods, or may be some other relevant factor the courts may take into account in deciding if the tool software is of adequate quality or fit for purpose. It is worth noting that the obligation to consumers that digital content is of satisfactory quality does not cover anything which might otherwise make it unsatisfactory, such as failure to undertake a particular function, if that matter has been communicated to the buyer before the sale.¹³

⁷ [1995] FSR 686.

⁸ Though in B2B contracts it is possible to draft an express term which would prevent this happening.

⁹ [2002] EWHC 2733 (TCC) paras 50-51.

¹⁰ [2012] 3 CMLR 44.

¹¹ Directive 2011/83/EU of the European Parliament and of the Council of 25 October 2011 on consumer rights, amending Council Directive 93/13/EEC and Directive 1999/44/EC of the European Parliament and of the Council and repealing Council Directive 85/577/EEC and Directive 97/7/EC of the European Parliament and of the Council, OJ L304/64 22 November 2011.

¹² UK Consumer Rights Act 2015 s 34(3).

¹³ *Ibid*, s 34(4)(a).

Civil law countries have faced similar difficulties in deciding how to treat software for liability purposes.

Italian doctrine classifies software licenses as atypical (or unnamed) contracts – that is, contracts whose content is not pre-determined by the law but is left to the parties' contractual freedom.¹⁴ When applicable, and if the party did not agree otherwise in the contract, the obligations of the contract are modelled after typical (or named) contracts, such as contracts of *vendita* (sale) where the licence is perpetual, ie non-expiring) or *locazione* (rental as an approximate translation) where the licence is for a fixed term. Similarity with a contract for sale has to a certain extent been assumed, particularly in cases in which the licenses were granted for an infinite period of time. The classification of the contract has relevant effects¹⁵ on the producer's liability and on what is implicitly guaranteed by the contract,¹⁶ when the express contractual terms do not override, and that overriding is allowed by the statutory provisions applicable by default to that particular contract type.

While Section 5.1.4 below will touch upon the general rules of contractual liability exclusion and limitation clauses with respect to B2B contracts, it seems worth clarifying how the different ways of classifying a license contract can affect the provider's obligations. Classifying – as the dominant doctrine seems to do¹⁷ – the license contract as a *locazione* (rent/lease) would allow the applicability of article 1579 of the Civil Code, which states that clauses aiming at excluding or limiting liability for a good's defects are void if those defects have been kept secret deliberately, or if they make it impossible to use that good for its intended purpose. Classifying the basis of the license contract as a sales one, however, would trigger the applicability of article 1490 of that same Code, whose second sub-clause provides that clauses aiming at excluding or limiting liability for the good's defects are void if entered into *mala fide* (in bad faith), without specific reference to the 'fitness for purpose' criterion mentioned by article 1579.

5.1.2 Scenario 2 – use of accountability tool in a SaaS relationship

In this scenario there is a contract between provider and customer for the use of the accountability tool, and this contract is clearly constructed as the provision of a service. The common law position is exemplified by s 13 Supply of Goods and Services Act 1982 which will imply into service contracts a term that the service provider will take reasonable skill and care in the service provision. Because this service is also the use of a digital product (digital content) there is also scope for common law courts to imply terms about the fitness for purpose of those products.¹⁸

So far as the liability of the tool producer is concerned, because there is no direct relationship with the customer the most likely source of liability is in tort. The drafting of the Directive on Product Liability¹⁹ seems to preclude any strict liability on the part of software producers,²⁰ and so the most likely source of obligations lies in the tort of negligence.²¹

¹⁴ See Art. 1322 of the Italian Civil Code (Codice Civile, Regio Decreto 16 marzo 1942, n. 262).

¹⁵ See Pietro Gobio Casali, *I contratti di software: qualificazione, responsabilità e garanzie*, I Contratti 4/2014, pp. 389 ss.

¹⁶ See Marco Montalbano, *Il regime delle garanzie e della responsabilità nella licenza d'uso del software applicativo*, Giustizia Civile 4, 2003, 121.

¹⁷ *Ibid.*, note 5.

¹⁸ See Section 5.1.1 above. See further Clarice Castro, Chris Reed, Ruy de Queiroz, *On the Applicability of the Common European Sales Law to some Models of Cloud Computing Services* 4(3) EJLT (2013) <http://ejlt.org/article/view/186>.

¹⁹ Council Directive 85/374/EEC of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products, OJ L 210/29, 7 August 1985.

²⁰ Both because software cannot be a product, and also because the Directive's regime imposes liability only for personal injury and physical property damage. See Chris Reed & Alison Welterveden, *Liability* in Chris Reed & John Angel (eds), *Computer Law* (4th, Oxford, OUP 2000) Ch 3.

²¹ The strict liability of the tool producer is likely to be excluded by the Italian *codice del consumo* (consumer's code, Legislative Decree 6 September 2005, n. 206), in which the norms transposing Directive 85/374/EEC are contained, as well. The objective liability contemplated in that piece of legislation does not substitute the general

At common law, liability in negligence depends on whether the defendant owes a duty of care to the claimant, which is based on the foreseeability of loss. If a duty exists, the defendant has an obligation to take reasonable care.²² The test for the existence of a duty differs depending on the nature of the relationship.

In our view, the nature of the relationship between tool producer and tool user is that the user is (and is known to be) relying on the tool to produce accurate outputs. This suggests that common law courts will apply a test similar to that used in the cases on negligent misstatements. The results produced by the software tool are similar to (and in many cases treated by the user as) statements of fact. If reliance on those statements is foreseeable it would seem reasonable to assume that the considerations influencing the courts in negligent misstatement cases will also be relevant in deciding the extent of the duty owed by a tool producer. The concern with software outputs is the same as for eg books and maps which contain incorrect information – potentially these might be relied upon by anyone in the world, and that is too large a class of persons in favour of which to impose liability.²³ Thus the common law's concern is to narrow down the class of users to whom a duty of care is owed.

In 1998 the courts took the approach of analogising software output to incorrect statements of information in deciding whether a software package which assisted users to decide which stocks and shares to buy constituted the giving of investment advice under the Financial Services Act 1976.²⁴ In *Re Market Wizard Systems (UK) Ltd*,²⁵ the Secretary of State petitioned for the compulsory winding up of the company on the ground that it was engaged in an unlawful activity, the carrying on of an investment business without authorisation contrary to s 3 of the Act. The company's business consisted of supplying computer software to end users. The software required users to enter share prices and other information about a selected set of securities on a daily basis, and then calculated the financial futures positions which the user should hold in respect of those stocks. The software was advertised on the basis that users could expect to make substantial profits if they followed its recommendations. The user manual described the Market Wizard program as a 'computerised trading tool' which 'generates detailed trading advice for a specific range of exchange traded securities'.

Carnwath J, holding that the company should be wound up, made two findings which are relevant here:

1. That the output of the Market Wizard program constituted advice:

I have no doubt that the signals generated by the use of the system constitute the kind of advice with which [section 15 of the Financial Services Act 1976] is concerned. The signals provide guidance as to the course of action which the user should take in relation to the buying or selling of the investments. Such guidance, in the ordinary use of English, is 'advice on the merits' of purchasing those investments. It matters not that the user is free to follow or disregard the advice; nor that he may receive further advice from his broker before making a final decision.

2. That the company was responsible for that output, and was therefore giving investment advice to users through the medium of the program's output:

remedies foreseen by the civil code, but provides a sectorial and alternative discipline whose application is left to the choice of the claimant. See further n 51.

²² *Donoghue v Stevenson* [1932] AC 562.

²³ The policy issues are explained in the detail in the judgment of Denning LJ in *Candler v Crane Christmas & Co* [1951] 2 KB 164, discussed below.

²⁴ Activities which required authorisation under s. 3 of the Act included:

'Giving, or offering or agreeing to give, to persons in their capacity as investors or potential investors advice on the merits of their purchasing, selling, subscribing for or underwriting an investment, or exercising any right conferred by an investment to acquire, dispose of, underwrite or convert an investment.' Financial Services Act 1976 Sch. 1 Pt. II para. 15.

²⁵ (1998) *The Times* July 31.

The question is whether the company is carrying on the business of giving advice. It is not necessary to identify a particular point in time at which the advice is given. It is enough, in my view, that it is providing the customer with a medium by which its purported expertise in the analysis of historical trading patterns is communicated in the form of advice related to a particular investment. If the programme were being operated by the company itself to produce the signals, in response to specific requests from customers, there would be no doubt that it was the company which was providing the advice. The fact that it is the customer who is operating the programme does not change the nature of the advice or its source.²⁶

This judgment is the first acceptance by the English courts that the output of software might constitute advice, although it must be recognised that there were a number of special factors in this case which led to that finding, in particular the way in which the company advertised the software. Nonetheless, the decision in *Re Market Wizard* is strongly supportive of the argument that non-contractual liability for defects in the output of a program should be dealt with on the basis that they are, or are closely analogous to, negligent misstatements.

The difficulty which the common law faces in this area is that statements may, potentially, be relied upon by anyone in the world. Similarly, the output of software may be relied upon by anyone. This level of liability is considered far too great, and so in a series of decisions the courts have attempted to explain how close a relationship is required in order to impose a duty of care, and thus potential liability.

The leading English case is *Hedley Byrne & Co Ltd v Heller & Partners*,²⁷ where the claimant suffered loss when he gave credit to a firm called Easipower in reliance on a financial reference given by the defendant bank. The court held that because the plaintiffs and the defendants were in a direct and close relationship, such that the defendant knew the claimant would rely on the statement, the defendant owed a duty to take care in giving the reference. This case established the possibility of claiming for negligent misstatements. The problem that troubled the court most was the danger of 'opening the floodgates' to litigation. The difficulty with careless words as opposed to careless actions is that the range of those affected is potentially very large indeed. The example given by Denning L.J in *Candler v Crane Christmas & Co*²⁸ of the marine hydrographer is instructive: we are asked to envisage that the hydrographer, in drawing up a chart of a particular part of the oceans, negligently fails to mark in a reef that is a danger to shipping. The chart is published, and is used by the masters of ships sailing in those waters. One or more ships run on the reef, entirely due to the fact that it is not marked on the chart. Should the hydrographer be liable to compensate the master of the ship, the ship-owners, and any passengers or cargo owners, all of whom will suffer loss because of his carelessness? Clearly the hydrographer satisfies the test of foreseeability laid down in *Donoghue v Stevenson*.²⁹ Clearly, also, his liability is potentially so wide, and extends so far in time (the charts might well be used for many years) that it seems wrong to say that he ought to be held liable.

The solution adopted in *Hedley Byrne & Co Ltd v Heller & Partners Ltd* was to limit the duty of care to those who were in a 'special relationship' with the maker of the statement. This special relationship was variously defined as being 'equivalent to contract' (per Lord Devlin), a voluntary undertaking given to the plaintiff to undertake skill and care (per Lords Morris and Hodson), or knowledge by the defendant that the plaintiff *would* rely on the statement (per Lord Reid). In any event, it was clear that the mere fact that it was foreseeable that some person in the defendant's position *might* rely on the statement would not be enough to establish a duty of care.

²⁶ The link between the company and the advice was, in the judge's opinion, reinforced by a requirement to update the software on a daily basis via the Internet from the company's website.

²⁷ [1964] AC 465.

²⁸ [1951] 2 KB 164.

²⁹ N 22.

The position has been somewhat clarified in *JEB Fasteners v Marks, Bloom & Co*,³⁰ where the defendant accountants negligently over-valued a company's assets in a report prepared for the company. As the defendants knew, the report was intended to be shown to prospective investors in the company. The plaintiffs, who were the eventual purchasers, brought an action against the defendants based on the negligent misstatement in the report. The court held that although the defendants did not specifically know that the report was to be shown to the plaintiffs, they did know that the report would be shown to, and relied on by, the class of intending purchasers, a class of which the plaintiffs were a member. There was, therefore, sufficient proximity of relationship for the defendants to owe the plaintiffs a duty of care, though in the event the action failed as the plaintiffs had not relied on that statement in deciding to purchase the company and so the negligence did not cause the loss.³¹

An important element in deciding the proximity question appears to be the purpose for which the advice was produced. In *Caparo Industries v Dickman*³² the House of Lords held that a company's auditors owed no duty of care to the shareholders in respect of the accounts because the accounts were not produced for the purpose of being relied on when making investments (even though it was foreseeable that they would be relied on).³³ We think the courts would usually find that software was produced for the purpose of being relied on – after all if the output of the software is not intended to be relied on, it is difficult to see why the user would want to buy it.

Subsequent cases have attempted to clarify the test for imposing a duty of care. There appear to be two main elements to that test: (a) that the maker of the statement has assumed responsibility for the accuracy of the statement, as evidenced by the relationship between maker and person relying; and (b) that the court considers it to be fair and reasonable to impose a duty.³⁴ In applying this test the courts will pay particular attention to how far the recipient relied on the statement, and any elements of the relationship which suggest that responsibility had not been assumed or that the particular reliance was unreasonable. This last point arose in *JP Morgan Chase Bank and others v Springwell Navigation Corp*,³⁵ where Springwell had purchased various poorly-performing investments through one part of the Chase Manhattan group, and argued that another part of the group owed it a duty of care to advise on the quality of those investments. In deciding that no duty was owed, the judge took account of the relationship between the relevant members of the Chase group (as known to Springwell) and the contractual documents signed between Springwell and various Chase entities.³⁶ The relationship was assessed as a whole, rather than breaking it down into a set of individual relationships and analysing each separately.

Applying these principles to the question of whether a software tool producer owes a duty of care to the ultimate user, the following position seems a likely one:

- (a) If the tool was commissioned by, or modified for, the user specifically, the producer is likely to owe him a duty of care unless there are other relevant factors, such as marketing

³⁰ [1983] 1 All ER 583.

³¹ See also *Haig v Bamford* (1977) 72 DLR (3d) 68 (Canada).

³² [1990] 2 AC 605.

³³ See the judgment of the Court of Appeal, [1989] QB 653.

³⁴ The cases are reviewed at length in *BSkyB Limited v EDS* [2010] EWHC 86, paras 328-356. See also *Henderson v Merrett Syndicates Limited* [1995] 2 AC 145; *HM Customs & Excise v Barclays Bank* [2006] UKHL 28, paras 5 & 88.

³⁵ [2008] EWHC 1186.

³⁶ ...on every occasion on which the parties came to document their contractual relations, and for whatever purposes, they agreed that Chase was not required to give any advice, was not to assume investment advisory obligations or responsibilities, and that Springwell acknowledged that it was relying on its own judgment in entering into the transaction. The documentation extends over a lengthy period of time, all of which was to the same effect, and is, in my judgment, inconsistent with the alleged, or any, advisory duties of care, of whatever scope. *Ibid*, para 480.

or contractual statements, which negate the existence of a duty. This appears to follow from *Hedley Byrne*.³⁷

- (b) If the software was produced for use by a limited class of users, eg for a particular cloud service provider which would make the tool available to its customers, that provides a basis on which the courts could decide that a duty of care exists. *JEB Fasteners v Marks, Bloom & Co.*³⁸ indicates that an important factor would be that the producer apparently intended the tool to be used and relied on by that group.
- (c) If the tool was produced for use by cloud users generally, rather than with any specific user or group in mind, then it appears unlikely that a common law court would hold that the producer owes a duty of care to the user, as the class of users is too indeterminate to satisfy the tests laid down in *Hedley Byrne* and *JEB Fasteners*.

Civil law courts may be less restrictive than common law courts on this third point. Dutch law imposes a duty of care based on the relationship between the parties, using the concept of 'relativity', but is likely to find a sufficiently close relationship if the producer intended a product to be used by the general public.³⁹ Italian law seems to take the same approach. The highest court, the Corte di Cassazione, decided in 1980⁴⁰ that the producer/seller of a product is responsible both in contract towards the buyer and extra-contractually towards third parties (including professionals) who are damaged by the defective product.

However, even if there is no tortious duty of care the producer might still have indirect liability, in that the user might have a contractual claim against its supplier and the supplier a contractual claim against the producer.

Even if a common law duty of care in negligence can be established, it is still necessary for the claimant to prove that the defendant was in breach of that duty, and that there is a sufficient causal connection between the breach and the loss that the claimant has suffered. The defendant will be in breach if he has failed to take as much care in producing his software as a reasonable person in the same position, professing the same expertise, would have done. This issue is discussed in Section 5.2 below.

The requirement that there be a sufficient causal link between breach and loss may also raise problems. The test for sufficiency is a simple one - is the loss a foreseeable result of the breach? This test was laid down in *The Wagon Mound (No 1)*,⁴¹ where oil that had carelessly been discharged from the defendants' ship was ignited by sparks from the plaintiffs' welding operations and burnt down the plaintiffs' wharf. On the evidence before it, the court held that fire damage was not a foreseeable consequence of the discharge, and thus the plaintiffs' case failed.

More recent cases such as *Anns v Merton London Borough Council*⁴² and *Junior Books Ltd v Veitchi Co. Ltd*⁴³ have emphasised the close connection between duty and causation. The test for a duty of care includes, in part at least, whether the defendant ought to have foreseen that damage of that type might occur.

It is possible to envisage situations where a duty of care is owed but the loss is an unforeseeable consequence of the breach of duty. For example, suppose that an accountability tool fails to notify the

³⁷ N 27. And is further supported by *Junior Books Ltd v Veitchi Co Ltd* [1983] 1 AC 520.

³⁸ N 30.

³⁹ See AS Hartkamp & CH Sieburgh, *Mr C Assers Handleiding tot de beoefening van het Nederlands Burgerlijk Recht. 6. Verbintenissenrecht. Deel IV. De verbintenis uit de wet*, Deventer: Kluwer 2011.

⁴⁰ Cassazione civile, sez III, 13 March 1980, no. 1696.

⁴¹ [1961] AC 388.

⁴² [1978] AC 728.

⁴³ [1983] 1 AC 520.

user that data has been wrongly deleted, and that for some reason the user's backup system overwrites the backup copies of the data before the loss is identified. Assuming the tool producer owes the user a duty of care (eg the tool was specially modified for that user) then it is clear that temporary loss of data is foreseeable. Nevertheless, it is arguably not foreseeable that all data would be lost, as it is standard practice to maintain backups for long enough to allow lost data to be restored. It follows that further consequential losses - perhaps loss of business, penalties imposed by a DPA, etc - will also be unforeseeable consequences of the breach.

However, such an approach is unappealing to a court because it absolves the careless defendant of all liability. The alternative is for the court to find that the loss was foreseeable but that the claimant's own actions were also a contributing cause, and thus reduce damages. This occurred in a real case of defective backup systems, *Logical Computer Supplies Ltd v Euro Car Parks Ltd*,⁴⁴ where the court (in effect) found that it was foreseeable that a computer user might fail to implement adequate backup systems, and therefore imposed liability on the contractor who negligently caused the data loss, but reduced the claimant's damages by 50%.

According to some authors, under Italian law the SaaS contract between the cloud service provider and the customer would be classified as an atypical – or unnamed – contract, whose terms would be seen as a combination between a software license contract and a service-provisioning contract (*appalto di servizi*).⁴⁵ Those agreements, in practice, lean more towards one of the aforementioned named contract types or towards the other depending on what their terms state expressly.⁴⁶ Alternatively, the courts might treat such contracts as not falling within any of the existing classifications.⁴⁷ DigitPA, the former Italian Agency for the Digital Agenda, has issued guidelines on cloud computing and public procurement which suggest that until the uncertain contractual classification of the contract is clarified, pragmatically it is possible to apply the *appalto di servizi* regime.⁴⁸

In the absence of an explicit agreement between the cloud customer and the tool producer, we believe that, in a civil law system such as the Italian one, the most likely source of liability for the producer would be the extra-contractual regime of general liability,⁴⁹ especially given the fact that the discipline of the consumer's code⁵⁰ is likely inapplicable (on several grounds) to the A4Cloud toolset.⁵¹

The general rule is contained in article 2043 of the Civil Code, which states that anyone committing an action that causes harm due to negligence or intentionality is legally bound to redress the unjust damage caused to others. The subject asking for redress would need to prove the damage suffered, the software failure and the causal link between them; the producer would need to prove he acted with the diligence

⁴⁴ Unreported, Queen's Bench Division 19 June 2001.

⁴⁵ See E.Belisario, *Cloud computing* (Pistoia, 2011); S Bendandi, "Software as a Service: aspetti giuridici e negoziali" www.altalex.it, 2008.

⁴⁶ Mantelero, Alessandro, *Il Contratto per L'Erogazione Alle Imprese di Servizi di Cloud Computing*, Contratto e Impresa, 2012, 4-5, 1216 ss.

⁴⁷ See Guido Noto La Diega. "Il cloud computing. Alla ricerca del diritto perduto nel web 3.0 (Cloud computing. In Search of Lost Law in the Web 3.0)" *Europa e diritto privato XVII*.II (2014) 577.

⁴⁸ Raccomandazioni e proposte sull'utilizzo del *cloud computing* nella pubblica amministrazione, 28 June 2012, http://www.agid.gov.it/sites/default/files/documenti_indirizzo/raccomandazioni_cloud_e_pa_-_2.0_0.pdf.

⁴⁹ In Italy extra-contractual liability is informed by the principle of atypicality, according to which every damage caused (at least) negligently is wrongful; common law traditions, differently, follow the principle of typicality, according to which there is a number of illicit behaviours (torts) that have a specific discipline and are based upon a precedent that allows the corresponding action. The common law tort of negligence, discussed in this Section, is in practice not dissimilar to the Italian approach, though because it requires a duty relationship between the parties it may apply in fewer cases.

⁵⁰ Legislative Decree 6 September 2005, n. 206.

⁵¹ The product liability discipline anchors the right to obtain redress to physical damages to people or things, and would therefore be applicable, even if commercial software programs are to be interpreted as 'products', only to that kind of software whose output has a physical effect and, when defective, is able to cause harm to people or property: see Montalbano, n 16.

required by the case.⁵² Generally, the threshold is determined by referring to the concept of diligence of the '*bonus pater familiae*' – an average degree of diligence that results in a standard level of common sense. Some of the A4Cloud tools, such as DPIAT, can however be framed for legal purposes as expert systems, which in effect simulate the output of a professional's expertise. As a consequence, it is likely that the degree of diligence required to meet the minimum threshold would be a higher one, based on professional standards⁵³ – both for the tools' output and for the developers' diligence and good faith requirements.

5.1.3 Scenario 3 – provision of an accountability service

This scenario is by far the simplest to analyse. The relationship between the provider and the customer is a contract for the provision of accountability services, and therefore the provider will have an obligation to provide that service using reasonable care and skill. Use of the tools is the mechanism selected by the provider, but this is irrelevant to the customer. The provider might choose to use different tools, or to use no tools at all.

Thus the legal question is, assuming the provider does use accountability tools, whether doing so is sufficient to amount to reasonable care and skill. The provider is relying on the tools to perform functions which would otherwise have to be undertaken by humans, ie deciding how data is to be processed and disclosed, and we can therefore narrow the issue down to a simpler one – is reliance on the tools alone an exercise of reasonable care and skill? This is discussed in Section 5.2.

Italian law, as an example of a civil law country, in this case and with respect to the toolset developed by the A4Cloud consortium, would probably frame the contract between the provider and the customer as a hybrid contract, with some elements from an *appalto di servizi*⁵⁴ (service-provisioning contract) and others from a software licensing contract. Several elements of a SaaS contract – such as the nature of the obligation undertaken by the provider in exchange for a fee to be paid by the customer and the former's likely organisational setup – would indeed call for its classification as an *appalto di servizi*,⁵⁵ particularly those concerning periodical or continuative provisions of services. Some other elements, on the other hand, are to be traced back to software licensing contracts. In this respect, it seems interesting to point out that, even if the default discipline for an *appalto* contract would lead towards the classification of the provider's obligation as a result obligation, which would impose strict liability for failing to achieve that result, in practice those contracts are drafted in a way that tries to frame the obligation as a means one,⁵⁶ which might in theory only require the provider to prove that it used reasonable care and skill in good faith. In practice the distinction has a merely descriptive function, without any impact on the regime of liability (strict, negligence, etc.),⁵⁷ and even those judges who maintain the distinction, do not infer from the classification of the contract in terms of "appalto di servizi" that the relevant obligation is necessarily a "result" one.⁵⁸

⁵² This is an important difference from the common law system, where it is normally the claimant who has to prove negligence rather than the defendant disprove it.

⁵³ See Art. 1176 of the Italian civil code.

⁵⁴ See Arts 1655 et seq of the Italian civil code.

⁵⁵ According to article 1655 of the civil code, an *appalto di servizi* is the contract in which a person undertakes to provide a service to or to build a particular piece of work for another subject in exchange for a certain amount of money, using its own organisation and at its own risk. Art. 1677 states that the discipline of arts 1655 et seq applies to service provisioning as long as compatible, along with the norms regulating the *somministrazione* typical contract (arts 1559 et seq), which is the contract with which a subject undertakes to provide an amount of goods to another subject, periodically or continuatively, and in exchange for a fee.

⁵⁶ Stefano Bendandi, *Software as a Service (SaaS): aspetti giuridici e negoziali*, Altalex, 2008, available on <http://www.altalex.com/index.php?idnot=44076>, last accessed 11 May 2015.

⁵⁷ See Cass. sez. un. 28-7-2005, n. 15781, Resp. Civ., 2006, 229 and Cass. 13-4-2007, n. 8826, Danno e resp., 2007, VII, 811. See also Noto la Diega, n 47.

⁵⁸ Trib. Arezzo, 30 January 2013, unpublished.

5.1.4 Express contractual terms

We noted in the previous Section that contractual terms disclaiming responsibility for the performance of an accountability tool might be relevant in deciding whether a duty of care is owed. It is important to recognise that national laws may contain provisions which invalidate such clauses. If the clause is not legally binding, we suggest that it should not be taken into account when assessing the duty question.

At the EU level the law has been harmonised by the Directive on unfair terms in consumer contracts.⁵⁹ Under the Directive a term in a B2C contract which is unfair will not be enforceable against the consumer, although the contract will still subsist so far as is possible and its remaining terms will be enforceable.⁶⁰ Although the intended market for the A4Cloud accountability tools is SMEs, some are likely to be used by consumers⁶¹ and, of course, many consumers are cloud customers and might receive an accountability service which makes use of the tools under terms of service which attempt to limit liability.

Art 3(1) of the Directive provides that a term is unfair if (a) it has not been individually negotiated and (b) “contrary to the requirement of good faith, it causes a significant imbalance in the parties’ rights and obligations arising under the contract, to the detriment of the consumer.”

To assist interpretation of this test, the annex to the Directive contains an indicative and non-exhaustive list of terms which “may be regarded as unfair”.⁶² Of these, the most important to software and computing services contracts are provisions:

- (b) inappropriately excluding or limiting the legal rights of the consumer ... in the event of total or partial non-performance...;
- (i) irrevocably binding the consumer to terms with which he had no real opportunity of becoming acquainted before the conclusion of the contract;
- (k) enabling the seller or supplier to alter unilaterally without a valid reason any characteristics of the product or service to be provided;
- (q) excluding or hindering the consumer’s right to take legal action or exercise any other legal remedy, particularly by...unduly restricting the evidence available to him or imposing on him a burden of proof which, according to the applicable law, should lie with another party to the contract.

In practice, this means that any attempt to exclude or limit liability to a consumer is likely to fail. However, as we have seen, the scope of any duty of care to a consumer is determined by how reasonable it is for the consumer to rely on the tool’s outputs – thus explaining clearly the limitations on using a tool safely can have the effect of preventing liability arising at all, assuming the explanation is accurate and understandable.

Controls on clauses in B2B contracts are not harmonised, but many EU Member States have national laws which impose controls. As an example, the UK Unfair Contract Terms Act 1977 provides in s 2(2) that a clause which excludes or limits liability for negligence, including breach of a contractual duty of care and skill, is not enforceable unless it satisfies the Act’s test of reasonableness. Section 3 subjects

⁵⁹ Directive 93/13/EEC, OJ L 95 April 21 1993.

⁶⁰ Art. 6(1)

⁶¹ Data Track is aimed primarily at consumers, and it is conceivable that a consumer might use COAT.

⁶² Art. 3(3). Note that the DTI in *Implementation of the EC Directive on Unfair Terms in Consumer Contracts* (DTI: London October 1993), its consultative document on the implementation of the Directive, takes this wording to mean that the terms in the list may be, but are not necessarily, unfair. Other Member States may take a stronger position on this point, and in any event including any of the terms in the Annex, para. is likely to give rise to a presumption of unfairness which will have to be disproved.

exclusions and limitations in standard terms to the same test. The UK courts have taken the position that as between businesses such clauses are likely to be reasonable, precisely because those businesses are the best judges of how to allocate risks between them.⁶³ However, if a tool supplier fails to explain clearly the risks of using an accountability tool, but nevertheless tries to exclude liability if the tool fails to perform properly, this might well be a factor which would persuade the courts that the term was unreasonable.

According to Italian legislation, outside the dispositions specifically aiming at safeguarding consumers (primarily the Italian implementation of the Unfair Terms directive),⁶⁴ parties are free to limit or nuance (or increase, for that matter) the liability deriving from a contract between them.

Article 1229 of the civil code, however, in its first paragraph sanctions the nullity of every agreement that excludes or limits the debtor's liability in cases of intentional misconduct or gross negligence, and in its second one the nullity of every pact aiming at excluding the liability of the party and of its auxiliaries in case of violation of legally sanctioned obligations. In a B2B contract, therefore, if the liability does not arise from *dolus*, gross negligence or the violation of the law, it can be contractually limited, both capping the maximum amount for which the debtor shall be liable and limiting the remedial options of the other party. The contractual terms limiting a party's liability, though, according to art. 1341, would need specific approval by the other side in written form or by equivalent means in order to be valid. The provider of the A4Cloud toolset would therefore have to make sure that they are agreed upon in an explicit and separate manner even when the contract is entered into through digital means. Moreover, reasonableness, fairness and the general good faith obligation (which binds the parties from their pre-contractual negotiations to the execution of the contract to its interpretation⁶⁵) can mitigate the effect of contractual clauses and hence also of exclusions and limitations of liability. This means that, even if a liability clause is valid, it will not be effective *per se* if under the specific circumstances invoking such a clause might be contrary to the aforementioned principles of good faith, reasonableness and fairness, and therefore void.

The limitations briefly sketched above only apply to liability exclusion and limitation clauses, and not to other contractual terms that have, in practice, similar effects, such as indemnification clauses, or clauses limiting the scope of the contract, unless those terms directly circumvent the statutory provisions applying to liability limitation or exclusion clauses. As mentioned in Section 5.1.1 above, moreover, article 1229's general rule aside, other special norms limiting what can be excluded by the provider can be eventually applied, depending on how the contract is concretely interpreted by the courts. Moreover, the user's behaviour, when concurring in determining the damaging event, has been deemed sufficient to exclude the producer's liability, and the aforementioned behaviour has to be assessed in light of the intended use of the product – in this instance, of the A4Cloud toolset.⁶⁶ Hence, the specification of the tools' purposes and instructions by the producer and the adherence to that specification by the user (or lack thereof) is a critical issue to evaluate, just as would be the case under common law.

5.2 Assessing reasonable care and skill

As we saw in Section 5.1, there are circumstances in which those who produce and supply accountability tools will owe a duty of care to the customers who use those tools. The next question is how to decide whether a breach of duty has occurred. The question here will be whether the tool was produced, selected or recommended using reasonable care and skill.

⁶³ *Watford Electronics v Sanderson* [2001] 1 All ER 696.

⁶⁴ Eg those contained in artt. 33 et seq. of the Italian Consumer Code (Legislative Decree 6 September 2005, n. 206).

⁶⁵ See artt. 1337, 1375 and 1366 of the Italian civil code.

⁶⁶ See Consiglio Superiore della Magistratura, "*La responsabilità da prodotto difettoso*", Incontro di studio sul tema: "Tutela dei consumatori", coord. dott. G. Grasso, Roma 14-16 Nov. 2005, p. 8.

Breach of duty is also an important question for those tool users who owe duties of care in respect of their use of the tool (eg the service provider in scenario 3 who uses the tools to offer an accountability service to customers). The issues here are whether reasonable care and skill were used in deciding to use the tool in the first place, and in how the tool is actually used.

So far as the common law is concerned, assessing breach is exactly the same process, irrespective of whether the duty of care is contractual or tortious. The test, established over 150 years ago, is as follows:

Negligence is the omission to do something which a reasonable man, guided upon those considerations which ordinarily regulate the conduct of human affairs, would do, or doing something which a prudent and reasonable man would not do.⁶⁷

To prove breach of duty the claimant first has to identify the act or omission which is alleged to be negligent, and then it is for the court to say whether in all the circumstances that act or omission meets the test. In theory there is no more that can be said, because the court's decision is entirely context-dependent, but in practice the courts are guided by principles that have been developed in previous cases. The most relevant of these are:

- Those who profess special skills, such as doctors (and presumably software developers⁶⁸) are expected to perform at a higher standard than unskilled persons.⁶⁹
- Acts or omissions which present no foreseeable risk of loss to those to whom a duty is owed cannot amount to a breach of duty, even if they in fact cause loss.⁷⁰
- Known risks must be balanced against the cost of avoidance, including the loss of benefits from the risky activity.⁷¹
- The greater the potential loss, the more precautions a reasonable person is expected to take.⁷²

5.2.1 Tool producers

If tool producers owe any duty of care to tool users (see Section 5.1.2), then there are likely to be two main elements to that duty. The first is to design, code and test the tool using reasonable care and skill. The second is to explain to users how to implement the tool so that it operates properly, and also to disclose information which is relevant to the user's reliance on that tool, such as limitations on its workings or accuracy. Failure to do either is likely to amount to negligence, and thus give rise to liability to any user who is owed a duty of care.

⁶⁷ *Blyth v Birmingham Waterworks Co* (1856) 11 Ex 781, 784 per Alderson B.

⁶⁸ See Conor Ward, Liability of computer consultants (1995) CTLR 68, 69:

In determining the standard of skill and care expected of the computer consultant, the court will take into account what computer consultants do in fact achieve ordinarily, as well as what members of the profession ought to achieve. Expert evidence will be relevant to what is the norm in the profession, as will professional literature published by relevant professional bodies.

More recently, a consultant advising on the procurement of a computer system was held negligent through failing to exhibit the expertise expected from other members of the profession - *Stephenson Blake (Holdings) Ltd v Streets Heaver Ltd* [2001] Lloyd's Rep. PN 44 (QBD (OR)).

⁶⁹ *Bolam v Friern Hospital Management Committee* [1957] 1 WLR 582.

⁷⁰ *Roe v Minister of Health* [1954] 2 QB 66. (Storage of sealed glass containers of anaesthetic in disinfectant was not negligent because the possibility of contamination via hairline cracks in the glass was unknown at the time. Note that now the risk is known, such storage *would* be negligent.)

⁷¹ *Bolton v Stone* [1951] AC 850 (Cricket club not liable to passer-by injured by cricket ball because the costs of avoidance – building a fence or abandoning playing the game – outweighed the small risk).

⁷² *Paris v Stepney Borough Council* [1951] AC 367 (One-eyed workman needed greater eye protection because the consequences of eye injury would be more severe than for other workers).

Design, coding and testing are, nowadays, well-understood activities. In deciding whether a tool producer failed to use reasonable care, the courts would receive expert evidence from software designers and producers. Adopting acknowledged industry best practice is likely to be a relevant factor as well. The question for the court will be whether the defect in the tool which caused loss to the user is one which should have been identified and corrected by a competent software producer.

But how is innovation likely to be treated? By definition, innovation goes beyond industry best practice. So the question for the courts is whether a reasonable software producer would have adopted that innovation, and whether it was properly tested to see if it operated as expected. Again, these are questions of fact which depend on the context of the particular case. Adopting an innovation is not negligent per se, even if it does not work perfectly, so long as it is at least no less effective than the industry standard alternative.⁷³

The second element of the duty of care is to provide sufficient information to users so that they can use the tool effectively. There is no doubt that it is likely to be negligent to fail to warn users of known defects or limitations which are likely to result in loss.⁷⁴ We think that the courts are likely to go further, and expect producers of innovative software to warn users about *potential* limitations which have a non-trivial risk of causing loss.

An obvious example, in the context of the A4Cloud accountability tools, is the limitations of those tools in achieving legal and regulatory compliance by the user. One clear finding from the project is that it is not possible to capture in the tools the full complexity of the legal rules and the factual context in which they are to be applied. This is particularly obvious in the case of COAT and DPIAT, but is potentially relevant to all the tools – see Section 0 below. The consequence is that those tools which contain legal and regulatory elements can only contain a partial approximation of the relevant law and regulation, and may not identify relevant facts. We think that failure to alert users to this limitation may well be negligent. Technical and operational limitations, eg that the A-PPL Engine relies on inputs from external sources which are not tested or validated by the tool, also need to be explained.

A helpful illustration, though in the context of a service rather than a software product, is the US case of *Brown v United States*.⁷⁵ There, two fishing boats were caught in extreme weather due to an inaccurate weather forecast, which caused both vessels to lose crew members and one of the boats to ultimately sink. The National Oceanographic and Atmospheric Administration, which produced the forecast using its computer systems, knew that the input data for the forecast was defective because a relevant weather buoy was out of action. It was held that issuing the forecast without adding a warning that it was potentially inaccurate amounted to negligence, and thus the plaintiff's case succeeded.

It is worth noting in this context that the practice in the IT industry, derived from its US-centric structure,⁷⁶ is to deal with the information problem by means of contractual terms limiting or excluding liability. A typical term is as follows:

Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS BASIS", WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the

⁷³ See eg the US case of *Gosney v State of California* (1970) 89 Cal Rptr 390, where the court refused to grant an injunction to force the state to operate its new computerised system of social security payments properly and to introduce new checks and 'fail-safe' procedures. Overall the new system reduced the number of errors; the question of negligence in respect of individual cases could be left to trial of those actions.

⁷⁴ *Walton v British Leyland* (1980) Product Liability International 156 (known manufacturing defect in car).

⁷⁵ (1984) 599 F Supp 877.

⁷⁶ US laws tend to give strong weight to the parties' freedom of contract, and therefore enforce contractual limitations and exclusions of liability far more frequently than do the courts of EU states.

appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.⁷⁷

This alerts the user to the possibility that the software might not achieve what the user expects, but does not give any specific information and how and why it might fail. We think that this is insufficient to discharge any duty of care; a more detailed explanation of known and anticipated limitations is needed. A further defect in this approach is that in Europe such clauses may well be unenforceable – see Section 5.1.4. An explanation which discharges any duty of care, and thus prevents liability arising, is far preferable to an attempt to exclude liability after it has arisen, particularly because the exclusion may well not work.

5.2.2 Tool users

Tool users, such as the service provider in Scenario 3 which uses accountability tools to provide an accountability service to users, owe a different duty from tool producers. Here there are three questions for the court: would a reasonable person have decided to use that tool; was the tool implemented and operated properly; and was reliance on the tool's output justified. The latter is particularly important in the case of the A4Cloud tools – there is no point in using them if their outputs cannot, at least to some extent, be relied on.

An important factor in deciding these issues will be the reputation of the tools in the industry. If their use were supported by a reputable industry body, such as the CSA, this would suggest that electing to use the tools would not in itself be negligent.

We think though that, in all cases, the information provided by the tool producer about the tool's known and expected limitations will play the most critical role. It will explain what the tool is supposed to do, how to implement and operate it and, if our recommendation in Section 5.2.1 is adopted, how far its output can safely be relied on.

Where software is apparently operating correctly it will not be negligent to rely on its output in the absence of other indications of error. In the US case of *Scott v District of Columbia*⁷⁸ a police officer arrested the plaintiff on the basis of an erroneous entry in a database which showed an outstanding arrest warrant. The court held that the officer had not been negligent because “[the plaintiff's] protests gave [the officer] no factual basis for questioning the accuracy of a computer system regularly relied on by officers throughout the metropolitan area.”

However, if the person relying on the tool's output has knowledge which indicates a problem, then reliance on the tool's output will potentially be a breach of duty. The UK decision in *Prendergast v Sam & Dee Ltd*⁷⁹ related to information on paper, but could equally apply to software or database output. In that case a pharmacist misread a doctor's writing on a prescription form and as a result supplied drugs to the plaintiff which caused him to suffer brain damage. The pharmacist argued that a reasonable person would have made the same mistaken reading. However, on that reading the prescription was still clearly defective because the drug was not made in the strength stipulated on the prescription. The court held that this knowledge gave the pharmacist a duty to check the prescription with the doctor before he dispensed it. His failure to do so was therefore negligent.

Use of, or reliance on, a software tool to achieve a particular function may also be negligent if the user should have known that the tool might fail to achieve that function. A graphic example of this can be seen in the use of the (then) innovative technology of radar for collision avoidance at sea. In *The Lady Gwendolen*⁸⁰ a ship was fitted with the new technology of radar to assist in avoiding collisions. However,

⁷⁷ Apache 2.0 license, <http://www.apache.org/licenses/LICENSE-2.0.html>.

⁷⁸ (1985) 493 A 2d 319.

⁷⁹ (1988) *The Times*, 24 March 1988.

⁸⁰ [1965] P 294.

the master of the ship misunderstood the technology and used it to travel at high speeds, even in fog. At the time he collided with another ship, in dense fog and in the restricted channel of the Mersey, he was not only travelling at top speed, but was operating the radar incorrectly. The court was clear that even if the master had operated the radar properly, he would still have been negligent because the radar did not give sufficient warning of other shipping to allow him to proceed at such a speed. Similarly in *Central Maine Power Co v Foster Wheeler Corporation*,⁸¹ a US power company brought an action in negligence against the designer of a condenser which leaked and damaged other parts of the plant. It was held that the power company was guilty of contributory negligence⁸² as its employees had relied solely on the plant's computer control system to bring any alarms to their notice, though at the time the system was not set up to do this and thus failed to alert them in time to prevent the damage.

On a more positive note, though, once the tools are recognised as performing well enough to provide appreciable benefits to users, it might actually be negligent for a user not to adopt them. The courts have no hesitation in holding that serious risks should be guarded against by obtaining or even developing new technology, provided this can be done at an appropriate cost in relation to the risk.⁸³ Another US marine safety case, *The TJ Hooper*,⁸⁴ illustrates this clearly. In that case the plaintiff's barges were lost in a storm at sea whilst being towed by the defendant's tugs. If the tugs had been fitted with radios, they could have received warning of the storm and taken shelter, thus avoiding the loss of the barges. In spite of the fact that it was not common industry practice to fit radios to tugs, the court held that the defendant shipowner was negligent - the technology was easily available, comparatively cheap, and its utility was clear. More recently in *United States Fire Insurance Co v United States*⁸⁵ it was held that the actions of the Coast Guard in calculating the site of a navigation beacon by manual rather than computerised means, when the computer system was both available for use and known to be many times more accurate, was potentially negligent and thus an issue to be decided by the jury at the trial.

5.3 Conclusions on liability

As we have seen, both producers and users of accountability tools are at risk of being liable to those who suffer loss if the tools do not work correctly. There will always be potential liability to contractual partners, though the liability risk can to an extent be controlled through exclusion and limitation clauses. In addition, there is also a real likelihood that producers and users owe a duty of care to those who use the tools, or are affected by tool use or reliance on tool output. Whether a duty of care exists, and if so what constitutes breach, can only be decided in the context of the particular relationship; we have attempted to explain the principles which will guide the courts in making those decisions.

Two factors over which tool producers and users have some control will be particularly relevant, and it is worth emphasising them here:

- Compliance with industry best practice is strong evidence that reasonable care and skill has been exercised. In particular, if the A4Cloud tools are accredited by some reputable body, that will be a strong factor in persuading the courts that their design and production was suitably careful, and that a reasonable cloud service provider or cloud user would adopt the tools.
- Most critically, it is essential for tool producers to provide users with adequate information about the limitations of the tools, including limitations on their role in data protection compliance. It is also important to bring this information to users attention in a way which

⁸¹ 684 F Supp 724 (D Me 1988).

⁸² If the claimant's own negligence was a contributing cause of the loss, then common law courts will apportion the loss according to each party's share of fault (the English law position), or in some jurisdictions deny the claim altogether, as happened in this case.

⁸³ *General Cleaning Contractors v Christmas* [1953] AC 180.

⁸⁴ (1932) 60 F 2d 737.

⁸⁵ (1986) 806 F 2d 1529.

makes it likely that they will read and act on the information.⁸⁶ This will go a long way to discharging any duty of care which tool producers owe. Users who rely on the tools outside their limitations, or do not implement them properly, are at real risk of being found negligent.

6 Relationship between data controller/DPA

There are two aspects of the relationship between a data controller and the DPA to which it is responsible which are likely to be affected if the controller uses accountability tools. First, one of the motives for adopting these tools is to reduce the risk that the controller is not complying with its data protection obligations. Thus we need to examine how DPAs are likely to react if, despite using the tools, a breach occurs. Second, we have identified in Deliverable D-4.11 and its update⁸⁷ that DPAs are increasingly using reactive and proactive investigations into data controllers' processing activities and policies as a regulatory technique whose aim is to detect past, current or potential breaches and improve the controller's future compliance. From that work we are able to make some extrapolations about how a DPA might want to make use of the accountability tools as part of its investigation.

6.1 Accountability tools as a means of compliance

The most important thing for controllers to note is that their data protection obligations are, for the most part, absolute obligations. Thus it is no defence that the technologies adopted for compliance failed to act as expected. This is so even if that failure is the fault of some other person (such as the tool producer) and the controller acted entirely reasonably in relying on the technology.

The absolute nature of obligations is clear from the wording of the relevant law. Art 6(1) of the Data Protection Directive (DPD)⁸⁸ states: "Member States shall provide that personal data *must be* ..." (emphasis added), and art 6(2) provides that "It shall be for the controller to ensure that paragraph 1 is complied with." Similarly absolute language is used in the DPD's other provisions imposing obligations on data controllers,⁸⁹ and in the Proposed General Data Protection Regulation (GDPR).⁹⁰

This means that acting reasonably by implementing compliance technology is only a defence if the law specifically so provides. Examples from the DPD include the obligation in art 6(1)(d) that personal data should be accurate and up-to-date ("every reasonable step must be taken to ensure that data which are inaccurate or incomplete ... are erased or rectified") and the data security obligation in art 17(1)⁹¹:

Member States shall provide that the controller must implement appropriate technical and organizational measures to protect personal data against accidental or unlawful destruction or accidental loss, alteration, unauthorized disclosure or access, in particular where the processing involves the transmission of data over a network, and against all other unlawful forms of processing

Having regard to the state of the art and the cost of their implementation, such measures shall ensure a level of security appropriate to the risks represented by the processing and the nature of the data to be protected.

⁸⁶ See further Section 8 for a discussion of the difficulties of making users aware of this information.

⁸⁷ See n 94.

⁸⁸ Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data [1995] OJ L281/31.

⁸⁹ See eg arts 7 ("personal data may be processed *only if* ..."), 8(1), 10, 11.

⁹⁰ See n 122

⁹¹ It is important to note that this article only outlines the principles to be adopted by Member States. There is a substantial body of specific regulation and guidance material in most states, and from the EU Article 29 Working Party, which will be taken into account when deciding whether a data controller has taken appropriate measures.

It is worth noting that appropriateness is not always congruent with reasonable care and skill – for example, if even the best available security is inadequate for the data processing in question then it will be an infringement to undertake the processing because of this inadequacy – but in most cases, security will be adequate if a reasonably careful and skilled controller would have adopted it.

Unless such a defence is available, the relevant DPA may want, or even be obliged to, sanction the breach. The question is thus how far the use of accountability tools might influence the discretion to impose sanctions, if such discretion exists, or the penalty imposed.

As already explored in D-4.11, EU DPAs⁹² often regulate the processing of personal data by deploying one specific regulatory tool, namely, investigations,⁹³ in the context of cloud computing.⁹⁴ Personal data means “any information relating to an identified or identifiable natural person”.⁹⁵ DPAs are statutory independent⁹⁶ public regulatory bodies which have various functions including applying and enforcing EU data protection laws in member states.⁹⁷ Investigations refer to one of the enforcement powers of DPAs, namely, their power to investigate data controllers,⁹⁸ such as companies which offer cloud computing services or technologies (ie cloud service providers), in specific circumstances (eg when an individual complains).⁹⁹ DPAs can trigger Investigations in various circumstances, such as after an individual has filed a complaint against a cloud service provider or even out of their own volition. As we will explain below, whether an investigation is a reactive or an *ex officio* investigation may at times impact on whether a DPA imposes a sanction (and if so, which type of sanction) after its investigation. Other factors, such as the particular DPA’s specific powers of intervention, the nature of the breach (eg is it a past or ongoing breach?), and the level of public interest in the investigation can also impact on how a DPA deals with a breach which it has detected during an investigation.

6.1.1 Responding to a complaint

Where DPAs have triggered a cloud investigation in response to a complaint filed by an individual, most DPAs are bound to impose the relevant sanctions set out by their national data protection laws (eg a fine or an order to halt the processing). Where imposing a sanction is compulsory DPAs have little room to exercise their discretion to take into account possible extenuating circumstances, such as defective performance by accountability tools, when deciding the regulatory responses to such compliance breaches.

⁹² The study in D-4.11 was limited to EU DPAs, but we think it likely that non-EU DPAs will take similar approaches. Thus in the remainder of this discussion we simply refer to DPAs.

⁹³ For more on investigations as regulatory tools in the data protection arena, see Peter Carey, *Data Protection: A Practical Guide to UK and EU Law* (OUP 2011) 69 and 127; Philip Schütz, ‘The Set Up of Data Protection Authorities as a New Regulatory Approach,’ in S Gutwirth et al *European Data Protection: In Good Health?* (Springer Netherlands 2012) 125.

⁹⁴ Asma Vranaki and Chris Reed, ‘Cloud Investigations by European Data Protection Authorities: An Empirical View,’ (A4 Cloud, WP 44, D-4.11, 28 February 2015), Asma Vranaki and Chris Reed, ‘The Rise of Investigations by European Data Protection Authorities in the Context of Cloud Computing,’ (A4 Cloud, WP 44, D-4.11, 30 September 2014).

⁹⁵ Article 2(a) DPD.

⁹⁶ For more on DPAs’ independence from the influence of government, legislature and other stakeholders, see Lee A Bygrave, *Data Protection Law: Approaching its Rationale, Logic and Limits* (Kluwer Law International 2002). See also Article 28, DPD.

⁹⁷ In different jurisdictions, various labels are used to denote the statutory independent public regulatory body which has the function of applying and enforcing data protection laws. For example, in the UK the DPA is referred to as the ‘Information Commissioner’ where as in Italy the DPA is referred to as ‘Il Garante per la protezione dei dati personali’.

⁹⁸ Article 2(d) of the DPD (n3) defines the ‘data controller’ as a ‘natural or legal person, public authority, agency or any other body which alone or jointly with others determines the purposes and means of the processing of personal data.’

⁹⁹ Article 28(3), DPD (n3). It should be noted that Article 28 has been inconsistently transposed by various European member states. For more on this, see Bygrave (n 4), 71ff.

If the DPA has both the power to fine and the power to issue orders then it has to consider which power would more effectively sanction the relevant data breach. For many DPAs fines are often considered as the last resort, but a fine is often the only available response to sanction data breaches which have already occurred (eg a specific instance of security breach).¹⁰⁰ However, national implementations of the DPD take different approaches to the question of sanctions; for example, under the UK Data Protection Act 1988 (as amended) the Information Commissioner only has the power to impose monetary penalties for compliance breaches if (a) the controller should have anticipated the risk of the breach occurring and that it could have serious consequences, and (b) the controller failed to take reasonable steps to prevent the breach.¹⁰¹ Thus if the compliance breach is of UK law and it was caused by failure of an accountability tool, the controller will not be fined if it can show that use of the tool amounted to reasonable steps.

The level of any fine is usually a matter for the discretion of the DPA or court, and here the reasonableness of adopting the accountability tool as a means of compliance is likely to influence the level of penalty imposed. The reason that the law provides a range within which a fine can be set is to reflect the seriousness of the misconduct in question, and a breach where the controller has made good faith efforts to comply must be less blameworthy than an intentional or reckless breach.

One might question the effectiveness of fines where a DPA has discretion to issue them, simply on utilitarian grounds. In many cases where DPAs have imposed fines after cloud investigations, such fines have not brought about a systematic change in terms of the provider's data protection operations and policies.¹⁰²

Other regulatory responses, such as issuing administrative orders (such as an order requesting the cloud service provider to amend its current data processing operations), tend to be deployed in cases where the DPAs seek to alter the ongoing or future processing operations of the cloud service provider.¹⁰³ If the DPA has discretion to choose between a fine to sanction the past breach and an order that the problems with the tool be fixed to prevent future breaches, again the reasonableness of adopting the tool will be an influential factor in making this decision.

6.1.2 DPA-initiated investigations

Where a DPA conducts a cloud investigation out of its own volition, it may often have more leeway in terms of how it addresses the data breaches detected during the Cloud Investigation. When dealing with large multinational cloud service providers, many DPAs will attempt to persuade the provider to alter its processing operations, amend its existing contractual documents (eg its privacy policy) or include more detailed information (eg inclusion of a pop-up box to explain why the cloud service provider requires the individual to disclose specific types of personal data) rather than bring about

¹⁰⁰ Interview of a representative of a DPA conducted by Dr Asma Vranaki on 30 May 2014 ('Interview 1').

¹⁰¹ UK Data Protection Act 1988 s 55A, as inserted by UK Criminal Justice and Immigration Act 2008 ss 144(1), 153:

This subsection applies if the data controller—

(a) knew or ought to have known —

(i) that there was a risk that the contravention would occur, and

(ii) that such a contravention would be of a kind likely to cause substantial damage or substantial distress, but

(b) failed to take reasonable steps to prevent the contravention.

S 55A also applies to telephone and electronic direct marketing communications by virtue of Regulation 31 of the Privacy and Electronic Communications (EC Directive) Regulations 2003/2426, reg 31. In this application there is as of 2015 no need to prove the likelihood of causing substantial damage or substantial distress, but it is still necessary to show that the person engaged in the communication failed to take reasonable steps to prevent the contravention (Privacy and Electronic Communications (EC Directive) (Amendment) Regulations 2015/35).

¹⁰² Interview 1, n 100.

¹⁰³ Interview of a representative of a DPA conducted by Dr Asma Vranaki on 31 July 2014 ('Interview 2').

such changes by applying a sanction. The existence of cloud accountability tools will be of relevance here, as these are exactly the kinds of technologies a DPA might try to persuade a provider to adopt. This will, of course, only happen if the DPA is convinced that the tool in question is likely to improve the provider's compliance performance. Thus it would be worthwhile for tool producers to enter into dialogues with DPAs during the final stages of tool development, in order to discover what features and functionalities DPAs have found to be desirable during their investigations. Once cloud service providers have developed a relationship with their DPA as the result of an investigation, they often discuss potential new ideas for technology which has data protection implications with the DPA at the very outset,¹⁰⁴ and this suggests that DPAs might be open to a dialogue with accountability tool producers.

As further explored in D-4.11,¹⁰⁵ DPAs regularly use the threat of sanctions in the course of cloud investigations as a negotiation tactic to persuade cloud service providers to bring their operations in line with the relevant data protection laws.¹⁰⁶ If the cloud service provider keeps resisting the changes recommended by the DPA, the latter may often have no other choice than to impose a specific sanction to bring about the relevant changes,¹⁰⁷ though for most DPAs this level of escalation is seen as the 'last resort'¹⁰⁸ when they are dealing with large multinational cloud service providers that are unwilling to bring their activities in line with the relevant laws.

6.1.3 Good faith and reasonableness

As we explained at the beginning of this Section 6.1, using accountability tools will not normally excuse a data controller for any breaches of data protection law. However, we have seen that if tools are implemented in good faith, using reasonable care and skill to ensure that they work to prevent breaches, this may well influence whether a fine is imposed (if the DPA has the discretion not to do so) and we think this will almost certainly reduce the level of any fine.

There is, so far as we can discover, no published guidance from DPAs about what they will consider to be good faith and reasonableness. But the language used in national implementations of the DPD¹⁰⁹ is very similar to that used by the courts in deciding claims for negligence. We therefore think that the discussion in Section 5.2 above is likely to be helpful to cloud service providers and other data controllers who adopt accountability tools.

6.2 DPAs' use of outputs from tools during investigations

During cloud investigations, DPAs deploy various techniques to investigate whether cloud service providers comply with relevant data protection laws. The investigative techniques which are deployed during a specific cloud investigation depend on various factors including the resources of the DPA,¹¹⁰ its investigative powers,¹¹¹ the business models of the cloud service provider (eg single service or multiple services),¹¹² and the data protection concerns raised by the organisation in question.¹¹³

¹⁰⁴ Interview 1, 100.

¹⁰⁵ See n 94.

¹⁰⁶ Interview 1, n 100.

¹⁰⁷ Interviews 1 and 2, nn 102 and 103.

¹⁰⁸ Ibid.

¹⁰⁹ See eg UK Data Protection Act 1988 s 55A, n 101; Irish Data Protection Act 1988 (2014 version) s 10(9).

¹¹⁰ See Vranaki and Reed, n 94.

¹¹¹ Ibid.

¹¹² Interview 1, n 100.

¹¹³ See Vranaki and Reed, n 94.

Most DPAs use a questionnaire at the start of the cloud investigation to evaluate if and to what extent the cloud service provider complies with the relevant data protection laws.¹¹⁴ Such questionnaires seek to evaluate how well the cloud service provider complies with its obligations under national data protection laws, such as the obligation to specify all the purposes for which it processes the personal data of its users.¹¹⁵ Additionally, many DPAs may also require the cloud service provider to provide them with a copy of its relevant internal and external contracts and documents, such as privacy policy, security policy, data processing agreements, and data protection and security guidance for staff.

Other documents may also be provided by the cloud service provider, at its own discretion or at the request of the DPA, such as examples of data breach logs, and screen shots of the various technical operations implemented by a specific technology in case of a data breach. Depending on how the cloud service provider answers the questionnaire, DPAs may also ask for further information such as the relevant source code which implements specific data protection and security tasks, such as cookie deletion.¹¹⁶

Some DPAs will review or test the relevant algorithmic sequences provided by the cloud service providers. Reviewing portions of algorithms means that the DPA only reviews the logical consistency of the code sequence. Testing portions of algorithms refers to an in-depth technical testing of the code sequences to evaluate whether their outputs match the intended logic, for example by checking that the retention period of a specific cookie is correctly implemented.¹¹⁷ In both cases, DPAs aim to review or test whether the organisation has technically implemented the obligations contained in its relevant internal and external contracts and documents. For example, in one cloud investigation, a DPA tested whether the computer instructions related to cookie deletion technically implemented the cookie deletion policy of the organisation as set out in its cookie policy.¹¹⁸

DPAs can often investigate, when they are on the premises of a cloud service provider, how the organisation deals with specific data breaches in reality. For example, one DPA told us that during one of its Cloud Investigation, specialist units of the cloud service provider showed to its investigative team in real-time how they dealt with data breaches as they occurred.¹¹⁹

We therefore think it likely that DPAs will be particularly interested in those elements of the A4Cloud tools which control processing (primarily the A-PPL Engine) and which log and take action in response to policy breaches (AAS, IMT/RRT, DTMT,¹²⁰ Transaction Log). Not only will DPAs want to check that the tools have been implemented properly in the data controller's system, but they may also want to dig into the logic and the code implemented by the tools themselves. Customers for the tools will need to provide these facilities to their DPAs, and so are likely to demand them from tool providers. It is therefore important for those producing these tools to consider how access to the internal working of the tools might be provided to DPAs without compromising commercial confidentiality in the code and logic.

DPAs conducting investigations are particularly interested in real-time compliance, rather than proceeding by after-the-event auditing. The A4Cloud tools have real potential to provide this

¹¹⁴ Ibid.

¹¹⁵ Ibid. A template of the general questionnaire used by the Irish DPA during its investigation can be found at section 8.2 of its audit resource. Audit Resource (Office of the Data Protection Commissioner, January 2009, Version 1.0) <<http://www.dataprotection.ie/documents/enforcement/AuditResource.pdf>> accessed 8 July 2014.

¹¹⁶ Interview 1, n 100.

¹¹⁷ See Vranaki and Reed, n 94.

¹¹⁸ Interview 1, n 100.

¹¹⁹ Ibid.

¹²⁰ The AAS, IMT/RRT and DTMT tools, which are not analysed here, all plan to include functionality to detect potential policy breaches, and to initiate some remedial action as a result. The workings of this functionality will clearly be of particular interest to DPAs, as data controllers will be using that functionality to substitute for human judgment about legal and regulatory breaches.

functionality, and again consideration should be given to introducing facilities which allow DPAs to link their analytical tools directly to the A4 tools.

Many DPAs aim to achieve an ongoing relationship with the investigated cloud service provider after the cloud investigation is completed. This is more prevalent in cases where DPAs are dealing with large multinational cloud service providers that have a sizeable share of the European market.¹²¹ This means that those DPAs might find it useful to have direct, remote access to the outputs of technological tools such as the Transparency Log so that they can detect when data breaches are taking place.

Others might benefit too from direct, remote access. The various draft versions of the GDPR¹²² place strong emphasis on the roles of independent internal and external auditors to assess the compliance of cloud service providers with the relevant provisions of the GDPR.¹²³ Additionally, the GDPR also enables DPAs to accredit specialised third-party auditors that can evaluate the compliance of data controllers, such as cloud service providers, with the relevant data protection laws.¹²⁴ Consequently, direct access to tools such as the Transparency Log would be most helpful to such auditors.

Independent accreditation by trusted third-parties could be extremely useful in establishing a market for the A4Cloud accountability tools. Many DPAs are constrained by a lack of resources, and while they would welcome the tools as mechanisms to adequately monitor and detect the non-compliance of cloud service providers with the relevant laws, they may not have the capability or funds needed to make their own assessment of how far the tools can be relied on for these purposes. Although third-party verification can present various opportunities (eg low cost, generation of customer trust), it can also raise a number of challenges, such as ensuring that the values and rights that are protected within the legislative frameworks are actually still comprehensively protected when private actors perform regulatory functions. Here we can look to other regulated fields, such as climate change, to learn lessons about how private actors can perform regulatory tasks successfully.¹²⁵ In particular, strong DPA involvement in selecting, accrediting, training, and overseeing third-party verifiers will be key to harnessing the potential of this partial privatisation of regulation. This suggests that the A4Cloud tool producers should be developing contacts with potential verifiers and with DPAs, in order to smooth the path to independent accreditation of the tools.

¹²¹ Ibid.

¹²² European Commission, 'Proposal for a Regulation of the European Parliament and of The Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)' COM 2012 (011) final <<http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:52012PC0011:en:NOT>> ; and draft European Parliament Legislative Resolution on the Proposal for a Regulation of the European Parliament and the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data COM 2012 (011) ('GDPR').

¹²³ Eg Draft European Parliament Legislative Resolution on the Proposal for a Regulation of the European Parliament and the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data COM (2012/0011), Amendment 36, Recital 60; Amendment 117, Article 22.

¹²⁴ Eg *ibid*, Amendment 136, Article 39.

¹²⁵ For third-party verification and the regulation of climate change see Lesley K McAllister, 'Regulation by Third-Party Verification,' 1 (2012) *BCL Rev.* 53.

7 The A4Cloud accountability tools

In order to conduct our legal analysis it was necessary to investigate the functioning of the A4Cloud tools and to conceptualise them in terms of their functions at a human rather than a technical level (see Section 2 above). This produced the following diagram (Figure 1):

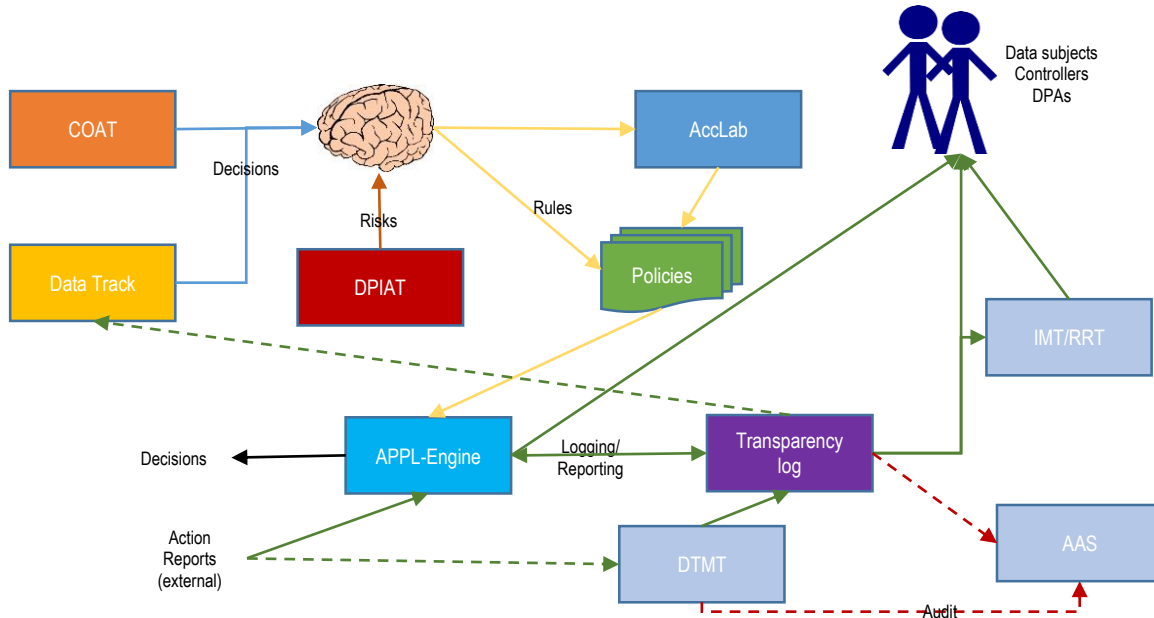


Figure 1: Tool functions

Three of the tools perform the information/advice function, in that their outputs are addressed to the mind of the user and are relied on for decision-making. COAT (Cloud Offering Advisory Tool) assists Small and Medium-sized Enterprise (SME) users to choose a cloud service on the basis of its accountability. Data Track informs individual users about how particular cloud offerings share their personal data. DPIAT (Data Protection Impact Advisory Tool) assists users to assess how their planned data processing raises data protection concerns.

Once a user has decided what cloud data processing operations it will undertake, and the restrictions which it wants or needs to place on them, the second function comes into play. Data operations need to be controlled or constrained, and this is undertaken by the A-PPL Engine. Controls or constraints are expressed in policies, either written by a human coder on the instructions of a user or by using the AccLab tool.

Logging and reporting is a function shared by a number of tools. At the time of conducting this research, only the Transparency Log was far enough developed to be included in the analysis. However, the other tools (DTMT, Data Transfer Management Tool; AAS, Audit Accountability System; and IMT/RRT, Incident Management Tool/Regulatory Response Tool) will raise very similar issues to those discussed in respect of the Transparency Log.

Our understanding of the tools, as explained below, was developed through first reviewing the various A4 Cloud documents which relate to the tools, and then seeking further clarification and explanation through dialogue with the tool producers. This dialogue was undertaken during January and February 2015, and the development of the tools has continued since then and will not be completed until well after the date of this deliverable. What appears below is a snapshot of the state of the tools as at that

time and may well not be accurate as to their final shape, particularly if some of the recommendations in this deliverable are adopted by the tool producers.

7.1 Cloud Offering Advisory Tool (COAT)

7.1.1 How COAT works

The Cloud Offering Advisory Tool (COAT) is aimed at both consumers and SMEs. This tool aims to perform a similar service to that of an insurance or utilities comparison website, but for cloud services and concentrating on factors relevant to accountability.

When users initially access the tool their country of origin is presumed and a default is selected, although this can be changed if it is not correct. Users then need to select whether they are a consumer or a business.

The user is then asked a series of questions which will be used in order to determine which cloud service provider is the most appropriate for them. For example, they will be asked questions such as their Price range, or where is an “Acceptable Storage Location including Backup”. The last currently offers the following series of answers: Europe (EU); US; Europe (Non-EU); China; Local; Any.

Other questions asked include:

- Data transfer in case of an emergency;
- Do you want encryption?
- Is it important that any disputes are resolved in your own country?

Figure 2 below is a screen shot which shows some of the other questions.

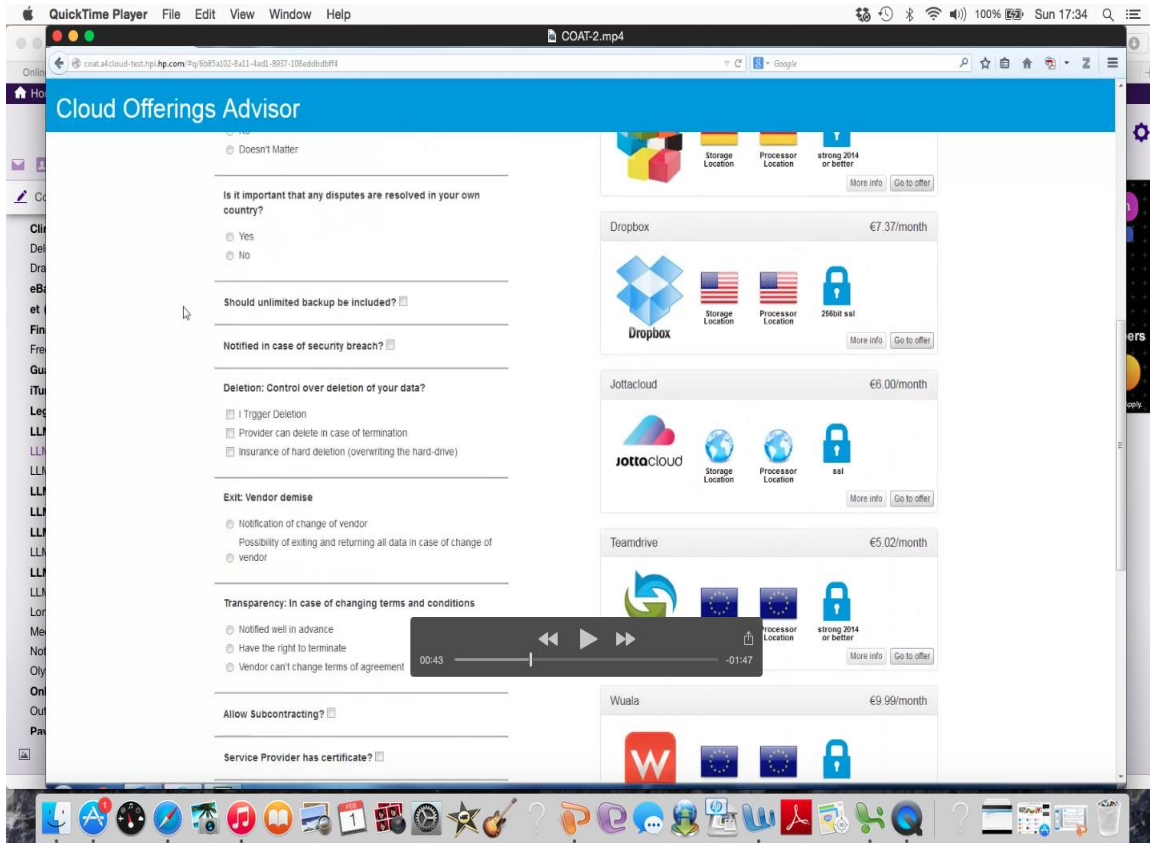


Figure 2: COAT questions

If a user of the COAT tool is unfamiliar with the relevance of any question and would like further information about what is being asked, they can hover over the question. For example, doing this for **Acceptable Storage Location including Backup** opens a pop up box which provides a detailed explanation about what the question is actually concerned with. See figure 3 below by way of illustration.

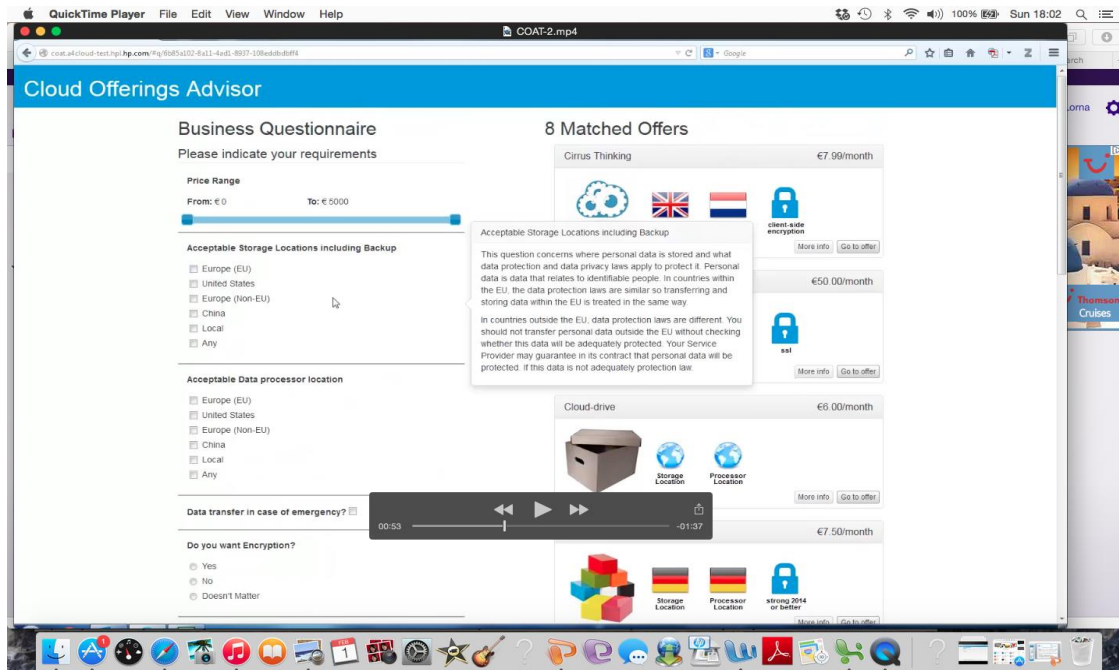


Figure 3: COAT pop-up explanation boxes

Once a user has answered all the relevant questions and selected their choices they will be shown a selection of cloud service providers on the right hand side of the screen. Each suggested match highlights the actual cloud provider, storage location, processing location, level of encryption and price per month. A user can also select to see *more info* about the cloud provider or *go to offer*.

On selecting the *more info* option a user will be shown additional information about the cloud service provider. At the bottom of the page the user has the option to select *view contract* to see the actual contract text for that cloud service provider. This facility may require review, as it is increasingly common for providers to spread the terms across multiple documents (commonly terms of service, service level agreement and privacy policy).

COAT educates users about the cloud service provider they choose. For example, they will be made aware if the cloud service provider uses a sub-contractor, in which case the user is informed about who owns the resources their cloud service provider uses. Thus if a user opts for Dropbox they will be told that Dropbox uses Amazon’s servers. Again this might need review, as the identity of the person who controls and operates the resources is more important for data protection purposes than mere ownership, and indeed, since the current version of COAT was constructed Dropbox no longer states that it uses only Amazon’s cloud services.

Once a user has selected its cloud service provider and the *Go to offer* option is selected, the user leaves the COAT platform and is redirected to the website of the chosen cloud service provider.

7.1.2 Legal and regulatory issues

From the user’s perspective, COAT appears to be offering individually tailored advice about how to select a cloud service provider, and in particular advice that selecting that provider will meet the user’s data protection obligations. A human who offered such advice would, in the absence of anything else, owe a duty of care to the user in tort, contract or both (depending on whether use of the tool is free or as part of a contractual relationship). We think the courts would be likely to treat provision of the tool

as provision of advice, as occurred in the *Market Wizard* case (see Section 5.1.2), though the producer (if different from the provider) might not owe a tortious duty of care in common law countries because the user-base is potentially the whole world and it has no individual relationship with the user.

There is a real risk that such advice, if incorrect, would be in breach of the duty of care for a number of reasons:

- The tool does not elicit information about every fact or preference which might be relevant for data protection purposes;
- The legal and regulatory issues covered by the tool are only a subset of data protection law, and are based on the Directive, not on the user's national law, which are likely to have some differences (of regulatory approach and practice, even if national law adopts the wording of the Directive);
- The legal explanations given in the pop up boxes are gross simplifications of the law, in order to be understandable by the user;
- The assessment of a service provider by the tool is inevitably over-simplified in order to make the tool usable. As an example, if the tool reports that data are encrypted, it does not explain which parts of the data are encrypted, and at what stages in the provider's processing.

For the COAT tool it is therefore critical that these limitations are explained, in such a way that the user realises that the tool is merely providing some assistance to the user in making its *own* assessment of data protection compliance.

There are also issues with the initial accuracy of the information in COAT's database, whether extracted by human analysis as at present or through automated text recognition if that is implemented. More importantly, attention will need to be given to maintaining accuracy, and this will be a challenge because cloud service providers regularly change their terms, the underlying technology and how their services are constructed.¹²⁶ Information about the limits of this information should therefore also be provided to the user.

Where users are SMEs, rather than consumers, it will be sensible to back up these explanations with terms limiting the tool provider's liability. As explained in Section 5.1.4, liability to consumers is almost impossible to exclude or limit.

From a user compliance perspective, using COAT will not be sufficient to ensure data protection compliance when selecting a service provider. However, if the tool is developed to a high degree of reliability, using it might convince a DPA to exercise any enforcement discretion it has if the choice turns out to have been an incorrect one.

7.2 Data Track

7.2.1 How Data Track works

Data Track (DT) is a tool aimed mainly at individual users in their role as data subjects (DS). It aims to enable them to discover who (which data controllers) process which data about them and for which purposes, and also to give DS control over those processes by enabling them to request that their data should be modified or deleted.

DT will be installed as a plug-in on the data subject's devices that may collect his or her data (i.e. on his computer, phone or wearables), mainly in his or her internet browser.

¹²⁶ For example, if Dropbox were to switch to operating its own storage service, that change would need to be reflected in COAT. Similarly, the recent change by Dropbox to using additional providers to Amazon needs to be captured.

DT collects information about personal data disclosed to cloud entities in two ways:

- 1) **Self-tracking** - It logs all data that are being disclosed to different entities while the DS is using the relevant app (such as Facebook) with an installed DT plug-in. The DT plug-in will log what types of data, for which purposes and by which entity, are collected based on the DS's online activity. Therefore, self-tracking only enables the user to track his data from the moment of installation of the plug-in, never backwards.
- 2) **Data imports** – DT will help the DS to make requests (in electronic form) to discover what personal data various data controllers process about him. This way of obtaining data via DT presumes that the DS knows who has his data, and also depends on the willingness of data controllers to provide them to DSs (they may not disclose all the data they are handling on a particular DS). DT developers are still exploring options on how to make these data imports work, since not all data controllers enable electronic takeout of data. However, DT works well with technologically helpful data controllers like Google or Facebook (in which case DS just clicks on the *import* button in DT). In the case of cloud service providers who have the A4Cloud tools installed, such requests would be received and handled by the A-PPL Engine (including requests for deletion and correction of data).

DT can track not only explicit, but also implicit personal data disclosures. This category of data includes “secret” data disclosures (eg capture of browsing history), but might also encompass data derived from the data explicitly provided. Derived data might only be accessible to DT from those providers who offer a data import function.

DT runs locally on a DS's device. It collects and then visualizes personal data and connections between them on demand of the DS. DT then stores all the personal data locally on the DS's device. This creates a potential problem, because it is very privacy invasive to collect and store all personal data on one device. Malware might therefore attack DT first to obtain the DS's personal data. Since data must be available for the tool unencrypted, at least every time the user opens DT (since it is installed only locally), malware can also obtain access to unencrypted data. One possible solution is to program DT so all the data will be destroyed after closing the personal data visualisations, but this has not been done yet.

A DS may also choose to store his or her self-tracked data (via a plugin) with a centralised cloud provider, and to synchronize the DS's devices and their installed DTs to eg visualise personal data using only one of them, in which case data may be stored also in the cloud with the possibility to locally encrypt them.

DT does not produce outputs to other A4Cloud tools. Its only outputs are the personal data itself, visualised for DS. No service provider has access to those outputs or their stored version unless the DS chooses to store the DT data in the cloud, and this data is encrypted locally by the tool prior to storage at a cloud provider.

Relations with other A4Cloud tools

A-PPL Engine generates notifications (i.e. about data breaches) and sends them to the DS's DT via the Transparency Log. These notifications are encrypted (hiding the relationship between data controller and DS) and are stored in the Transparency Log (since the DS may not be online) and then requested by the DT app. They will be displayed for the DS in an inbox.

If DT gets notification from A-PPL Engine about a data breach, a DS will see the notification and can then decide what action to take. Work is under way on a **Plugin for Assessing Policy Violations**, which assesses how severe the policy violation notified by A-PPL Engine is, informing the DS if the notification is worth his attention (if the data breach is serious) or if he can look at it later. It will probably be visualised in an inbox with the alert box for more important notifications. The researchers are currently investigating whether it is possible to link DT with IMT/RRT so that the DS can request

that tool to take action in respect of the breach, but this work may not prove achievable in the project lifetime.

7.2.2 Legal and regulatory issues

The main functionality of Data Track raises few legal and regulatory issues. Assuming it is released as a paid app, or included in a contractually-supplied suite of accountability tools or services, its provider will owe a duty of care to the user, but:

- If the limitations of the tool, such as a list of the providers it works with and any unexpected limitations (eg retrieving incomplete data from that provider), are disclosed, it is hard to envisage there being a breach of duty; and
- In any event, because users are individual data subjects, any losses they might suffer as a result are likely to be so small that a claim is not worthwhile.

DT enables a DS to have access to its own personal data, and this does not raise any data protection compliance issues.

However, the Plugin for Assessing Policy Violations is, potentially, in danger of crossing the line from information provision to advice, and thus raises the same issues as discussed in relation to COAT (Section 7.1.2). These risks are low in magnitude for the reasons given above, and can be controlled by giving clear and accurate information to the user about the tool's limitations.

7.3 Data Protection Impact Assessment Tool (DPIAT)

The DPIAT is targeted at SMEs and provides them with a report about the level of risk they may be exposed to by storing and processing their data with the Cloud Service Provider (CSP) they have selected.

Before a user begins the main data impact assessment questionnaire they undertake a screening. The screening is a preliminary assessment to determine whether the user is handling personal data, whether that data is sensitive, what the purpose of the processing is, the extent to which third parties are involved as well as how the information will be used. The results of this screening will show whether or not it is necessary to continue to the full questionnaire.

The full questionnaire is a series of 56 questions (the **Questionnaire**) which a user completes to understand the level of risk that their company faces when using a particular CSP. Before a user begins to answer the Questionnaire they select from a database of CSPs the one which they intend to use. If they do not select a CSP the report which they receive will be of limited value. Currently the development team has input data into a database regarding 46 different CSPs who voluntarily provided their terms and conditions.

The Questionnaire has been designed to be future proof¹²⁷ and thus reflects the present DPD and the proposed GDPR as at the time of the European Parliament's first reading.¹²⁸ In addition to the legal and regulatory framework the questionnaire also covers security and operational issues, using input from experts in those fields. The 56 questions¹²⁹ cover five areas: the type of project; the collection and use of data; the project's storage and security policies; data transfers; and cloud specific issues.

A user of the DPIAT does not need to complete the Questionnaire in one go. It may be that the user needs to leave the Questionnaire in order to find out some further information. The user's answers to

¹²⁷ See D:C-6.2 Prototype for the data protection impact assessment tool, 20/10/2014, 22.

¹²⁸ Ibid, 24.

¹²⁹ Ibid, 64.

date will be saved and once the user has found the appropriate information he can resume answering the Questionnaire by logging back into it via his ID and password.

At present the login details, the answers provided and the generated report are stored, but the user's response to the questions is not retained. No decision has yet been made about what, if any, uses that stored data might be put to.

There is currently no constraint on who can use the DPIAT tool. Thus the user who completes the Questionnaire might not be the main data protection officer at the company; it could be someone from projects or an employee from elsewhere in the business. This means that the person may not have expertise in data protection (although the questions have been designed to be in an understandable format for non-specialists), and of course the accuracy and utility of the report is dependent on the quality of the answers that are input by the person completing the Questionnaire.

As with the COAT tool there is also some education of the user about data protection law and CSPs. An explanation is given about each question asked and then responses are given to the answer the user provides. The user is then instructed about which question to continue with next.

The final output of the tool is a report which indicates the level of risk an SME will be exposed to using the CSP it has chosen. A score is provided: the lower the score, the lower level of risk. If the score is particularly high the tool does not recommend the user does not use the chosen CSP, rather it just highlights the areas of risk.

The score is based on a formula which integrates the weightings given to the user's answers. The risk is assessed for two main areas: the risk of using the selected cloud provider based on a) privacy, b) security, and c) service provisioning risks; and a privacy impact assessment of the user's intended processing activities. If no CSP is selected, the report only consists of the privacy impact assessment.

Weightings have been assigned by the research team on a comparative, rather than an empirical basis. For example, if third parties are used by the CSP for some aspects of data processing and storage this carries a score of 3, whereas in situations where no third parties are used the score is 0. The team has attempted to use a non-biased (as between CSPs) methodology for the weighting of the answers, but it should be noted that there are no external standards against which the methodology can be checked.

The final score given in the report ranges from 0-300 and is split into 5 categories (very low → very high). The scale is arbitrary and has no pyramid value, and is thus designed to compare between CSPs rather than to provide an objective risk measure.

7.3.1 Legal and regulatory issues

In our view, making the DPIAT tool available to users is likely to amount to the provision of advice unless very careful steps are taken to avoid this. This is because:

- The report is individually tailored to the user; and
- The risk weightings and scores are based on the judgment of the tool producer. There is no possibility of arguing, as might be done for the COAT tool, that DPIAT simply collects information and re-presents it to the user in a more comprehensible format.

It is clear, though, that if the report *is* advice it is by no means comprehensive or reliable. This is because, just as for COAT:

- The tool does not elicit information about every fact or preference which might be relevant for data protection purposes;

- The legal and regulatory issues covered by the tool are only a subset of data protection law, and are based on the Directive and the proposed Regulation, not on the user's national law, which is likely to have some differences (of regulatory approach and practice, even if national law adopts the wording of the Directive);
- The legal explanations given in the pop up boxes are gross simplifications of the law, in order to be understandable by the user.

In addition:

- Ensuring the information about CSPs and their systems is accurate and up to date is a challenging task, and it is likely that at any one time the database will contain inaccuracies. Extracting this information from changes published by CSPs is difficult because CSPs do not provide updates to their terms and policies in a uniform fashion, and because updates are in free text this requires judgment to interpret and resists automation. There is an inevitable margin of error and a danger of misclassifying the answers.
- It is difficult for users to check for errors in the report because the answers which have been given to the Questionnaire are not included. We think that this information needs to be included, perhaps accompanied by some explanation of how each answer affects the final risk assessment.

There is an additional danger of claims from CSPs if DPIAT uses out of date security data and other data. If users are, for example, informed that the CSP's security is risky, when in actual fact it has been updated and rigorously tested since the database was created, this may create liability under the tort of defamation.

One way to reduce the risk of inaccuracy, which is under consideration, is providing the user with a facility to update the database by importing information direct from CSPs, but this would require CSPs to produce that information in a standardised format.

Competition law is also an issue. CSPs not included in the database may suffer a market disadvantage as a result. Certainly a CSP which is newly added to the DPIAT database might obtain a commercial advantage over CSPs already included if their information is outdated, as the new CSP might appear to users to be using more advanced technology. These issues are outside the scope of this deliverable, but will need assessing before the tool is placed on the market.

Provision of detailed information about the limitations of DPIAT will reduce the risk that reliance on the output of the tool causes loss to a user, but may not be sufficient on its own. It may be necessary to market the tool under a contract which effectively limits the provider's liability, and that in itself is no easy task (see Section 5.1.4).

Despite these challenges, though, it is clear that DPIAT has real potential to assist data controllers in meeting their legal and regulatory obligations. It is unlikely that DPAs will accept use of DPIAT alone as compliance with obligations to conduct a data protection impact assessment, but it will be seen as a useful tool to complement (more expensive) professional advice.

7.4 AccLab

7.4.1 How AccLab works

AccLab's function is to assist cloud customers to code policies about data use. If those policies are expressed in A-PPL (Accountable PrimeLife Policy Language) code, they can be used as input to the A-PPL Engine which controls data usage in accordance with those policies (see Section 7.5).

The customer begins with human-originated text statements of data policies (eg "this data item must be deleted after two years"). AccLab then provides a graphic interface which enables the tool operator

to create representations of actors and to assign them their roles (data subject, controller, etc). The tool operator can then link these representations with arrows which describe their relationship.

AccLab is in theory usable by a non-technically trained person, such as a customer's data protection compliance officer. In practice, some technical expertise is useful, and it is more likely that the data protection/privacy expert will need to work with a technician to use AccLab effectively.

Once the graphic representation is complete and has been checked by the human operators, the tool automatically converts this graphic representation into AAL (Advanced Accountability Language) code. AAL is an abstract and formal language with semantics based on first-order temporal logic; it is expected to be more easily readable by humans than pure logic.

Once all policies are coded in AAL, the tool then parses the complete code suite so as to validate the policies against each other. This process uses an external tool, the TSPASS prover,¹³⁰ to analyse the logic. Logical conflicts are automatically identified where policies do not agree (eg data subject's policy is to delete data after two years, but data controller's policy is to keep data for 5 years).

Resolution of logical conflicts is achieved by manually editing the AAL code and then re-parsing the code suite to check that the conflict has been eliminated. The tool can provide some assistance in identifying the causes of conflicts, but this is not always achievable. Where the cause is identified by the tool, in some cases it can suggest corrections to the code.

The final step in the process is to translate each AAL policy automatically into A-PPL. Work on this is in its early stages, but hand coding from AAL to A-PPL is a much easier task than working direct from human policy statements.

7.4.2 Legal and regulatory issues

We think that there are two potential sources of liability risk for AccLab:

- Failure to translate a graphical representation of relationships and obligations accurately into AAL; and
- Failure to identify logical inconsistencies in policies, or suggesting incorrect amendments.

In both cases the consequence would be that the resulting policy suite does not accurately represent the human input.

These risks are similar to those of any software which produces incorrect results, and we think that they can be dealt with by adequate explanations to users about the responsibilities of the tool operators and the limitations on the AccLab's analysis and translation capabilities.

If the provider of a policy writing service used AccLab and produced defective policies, there is no doubt that the service provider would owe a duty of care, probably contractual, to its customer. Whether use of AccLab discharged that duty of care would depend on how reasonable it was to rely on its output, applying the principles explained in Section 5 above.

So far as DPAs are concerned, we think they would have little interest in ACCLab per se. It is the responsibility of data controllers to implement adequate controls on data for compliance purposes, in this case through its data policies. The mechanism used to produce those policies is irrelevant to compliance, though if AccLab is widely accepted as producing accurate results then a compliance failure as a result might influence the DPA in exercising its enforcement discretion.

¹³⁰ Michel Ludwig and Ullrich Hustadt, "Implementing a fair monadic temporal logic prover" (2010) 23 AI Communications 69-96.

7.5 A-PPL Engine

The A-PPL Engine is a tool which aims to ensure that the cloud system on which it is running actually applies A-PPL policies, and thus controls the use of data in accordance with those policies.

There are two ways in which the A-PPL Engine can be implemented as part of a cloud service:

- The cloud service provider runs the tool, and the customer (data controller or processor) uses that tool as a service to constrain its processing of personal data. In effect this offers the tool as SaaS; or
- The customer can run the tool directly and integrate it with its other applications. This is IaaS or PaaS use.

With SaaS use, the service provider is clearly accepting some responsibility for the proper operation of the tool, as part of its suite of service provisions. For IaaS/PaaS use, more responsibility is assumed by the customer. But we should stress that the law is not concerned with how the use of the tool is classified between SaaS, PaaS and IaaS – the issue is about which party is responsible for the tool's performance in the particular factual situation.

The tool takes as input A-PPL code from any source, including AccLab. Each A-PPL policy represents a general case (a policy template) which is completed with specific information when personal data is collected. In these cases the tool fills in a policy template with the specific data for a given data subject. For instance his/her email for notification and the date/time the data was collected and consent given, etc. The tool interfaces with the customer's systems in such a way that the relevant information is extracted for insertion into the policy once it is entered by the customer into its systems.

Each set of personal data, after being provided by a cloud customer, is matched with the relevant policy governing all its further processing and then stored by A-PPL Engine (the data store is separate from the tool). Currently the cloud service provider controls the policies, so that the customer cannot substitute its own policies into the tool. However, if the customer ran the tool in an IaaS/PaaS relationship with the provider, the customer would control which policies are used by the tool.

Enforcement of policies is achieved by blocking those actions (eg data disclosure) which are not permitted by the policies and by automated obligation execution. External components report to the tool what has been done with data, or requests which have been received, and the A-PPL Engine decides whether these are in accordance with the policies. If a request (eg to disclose data) violates a policy the tool will respond to the request accordingly.

The tool logs its blocking actions, and also where it has permitted actions to take place, using the Transparency Log. It also logs when new personal data is acquired, when consent is given, when personal data is deleted, changed, etc.

The tool also performs certain actions requested by data controller, data subject (via Data Track) or a third party, including data retrieval (in case of data access requests), correction or deletion of personal data and retrieval or deletion of policies. It does this by accessing the relevant data store or policy store.

The A-PPL Engine can send notifications to anyone "declared" as a recipient in the policy (data subject, controller, regulator, etc). The declaration takes the form of an obligation, which is composed of a trigger (event) and an action. The tool will send notifications to the specified user as a response to any of the possible events. Events may be time based (periodic or otherwise), when data is accessed, when a violation is detected by another tool, etc. Notifications are in the form of a text string indicating what has occurred, and are currently sent by email (though SMS is potentially a possibility), and potentially via the Transparency Log. If the notification relates to a policy violation, the A-PPL Engine

sends it to IMT/RRT, where the customer may read it, write a human readable version and send that human-readable version back to A-PPL Engine to distribute it to relevant receivers.

In some cases where data is sent to a third party it is accompanied with the relevant A-PPL policy which covers the obligations of that third party (a “sticky” policy), so that if the third party is also running the A-PPL Engine the policy can act as input to the third party’s systems.

7.5.1 Legal and regulatory issues

The A-PPL Engine presents three potential liability scenarios:

- Where the cloud service provider enables its customer to use the A-PPL Engine as a SaaS offering, the service provider will owe a contractual duty of care to the customer. The contractual promise is to take reasonable care in providing the service, or more specifically to use reasonable care and skill to ensure that the system implements the customer’s A-PPL policies. The legal question is therefore whether the service provider’s reliance on the tool to achieve this end is justified, and whether a reasonable service provider would have taken additional precautions (such as periodic audits) to check proper functioning.
- Where the cloud service provider uses the A-PPL Engine to enforce its own policies, the courts would concentrate on the wider question whether the service, as a whole, was provided using reasonable care and skill. This again would raise questions about the reasonableness of relying on the tool.
- Where the customer runs the tool itself, under an IaaS or PaaS scenario, we have in effect a software supply. The most likely legal obligation to be relevant is that of fitness for purpose (see Section 5.1.1), and the description of how the tool works and its limitations (probably contained in the tool manual) will be relevant to deciding the question of fitness.

In the first two scenarios the cloud service provider will have procured the tool as software from its producer, and this will raise the same contractual fitness for purpose issues.

In all these scenarios the parties will be in contractual relationships with each other, and so those relationships can be used to identify the risks of using the A-PPL Engine and allocate them appropriately. However, we should remind readers that contractual terms alone may be ineffective (see 5.1.4), and that providing detailed information explaining the use, risks and limitations of the tool is potentially effective in reducing liability risks.

For data protection compliance purposes, the A-PPL Engine plays the most important role in the suite of A4Cloud accountability tools. In theory at least, if the tool user’s A-PPL policies fully and accurately represent the relevant obligations relating to personal data, the A-PPL Engine will ensure those obligations are fulfilled. It is, of course, research prototype software, so has not been tested to the level expected for commercial roll-out, but if it were constructed and tested to that level the risk of policy contravention should be very low. The main risk of contravention lies in the writing of the A-PPL policies, which is undertaken by data controllers/processors or their representatives and is not the responsibility of the A-PPL Engine producers or providers.

The most likely reasons for contravention of data protection obligations would be:

- The policy does not represent accurately the obligations in respect of that personal data; or
- Policies are contradictory, in which case the tool will execute the first policy it encounters and therefore never execute the second, contradictory policy. As a consequence the contradiction will not be identified by the A-PPL Engine.
- External components fail to act as they should in response to the outputs of the tool

These are aspects of how the wider accountability system is implemented, rather than issues relating to the A-PPL Engine specifically.

However, the increasing use of audits by DPAs means that the logs produced by the A-PPL Engine will be of particular interest during the auditing process. DPA audits often include software testing, to assure themselves that the controls which the system purportedly implements are in fact working effectively. Thus any features which can be included in the tool to make that process cheaper and easier are likely to be welcomed both by DPAs and those being audited.

As with all the other tools, use of this tool will not of itself necessarily achieve compliance with data protection law. As previously discussed, though, if the DPA accepts use of the tool as a reasonable way of attempting to comply, this is likely to influence its enforcement discretion.

7.6 Transparency Log

7.6.1 How the Transparency Log works

Together with Data Track (DT), the Transparency Log (TL) is one of two transparency-enhancing tools in the A4Cloud ecosystem. Its main objective is to assist the communication of data from data controllers (DC) to data subjects (DS). The TL is a “secure and privacy friendly replacement of email¹³¹ and plays the role of a privacy friendly transport¹³² service. The sender (DC) is fixed whereas the recipient (DS) needs to register with the DC. Both DT and TL are designed so that they are usable even if none of the other A4Cloud tools are implemented.

TL provides a secure channel to communicate notifications, such as a data breach notification, between the data controller and the data subject. The communication is passed by TL to the DS who is able to read about the notifications at their convenience. The communications are sent with stronger security and privacy protections than is provided by email or many other communications methods. TL is designed to hide metadata about the message and to protect its content and so, for example, it does not reveal the identity of DS or the details of what data has been breached to a third party. Equally the communication to the TL does not reveal any details about the relationship between the sender (DC) and the recipient (DS). This is a great advantage as the present systems of notifying a DS about a breach, such as via email, SMS or push notification, can expose the DS’s sensitive personal data and an otherwise unknown relationship with a DC. For example, if today there were a data breach at Google it would presumably send a notification by email to each affected DS. The email is likely to contain personal data that is sent unencrypted which could have a larger negative impact on the DS’s privacy than getting no notification. The fact that a particular service is being used by a DS can on its own be sensitive personal data, eg when it comes to the use of particular healthcare services.

Messages transmitted to the TL can be publicly verified, ie proof can be provided that the message has been sent at a particular time. This public verification of when the message was sent is very useful when there is a data breach. A DC will be able to establish that timely notifications were sent in accordance with its legal and regulatory obligations. Alternatively a DS will be able to confirm that a notification was not sent as per the DC’s obligations. The time of sending can be established via trusted timestamps provided by an external party.

A DS is made aware of a notification, for example, a breach notification or notification of data processing, by the TL. Most commonly the A-PPL Engine generates the notification which is sent to the DS via a TL plug-in to the DT. The notifications are stored, since the DS may not be online, until

¹³¹ D:D-5.2 User Centric Transparency Tools V1, 2.

¹³² Ibid, 2.

they are requested by the DS from his or her DT. The severity of any breach notification will be highlighted to the DS in the DT.

Notifications from the DC to the DS are stored via an intermediate untrusted server. The use of an untrusted party has been a design decision. It could be the DS's cloud service provider who stores the notifications until the DS is ready to read them. Due to the design of TL the notifications cannot be tampered with undetectably and if, for example, a third party attempts to delete the notification this interference will be detected by the TL. A DC who stores notifications cannot modify them in any way once stored.

The stored messages are released to a DS when the DT is started, which plugs into the TL and requests all notifications to be sent to the DT.

Where information is stored will therefore depend on where the party elects to store the data.

The verifiability of a notification's content (that it has not been tampered with) and the times it was sent can be proved cryptographically. The underlying technology has been peer reviewed in academic literature, and is as strong as the cryptographic signature and time stamping technologies recognised for legal purposes by the EU Regulation on electronic identification and trust services.¹³³

The TL can also be used as a trusted and unalterable log by other A4Cloud tools.

The tool uses a mix of some old and some state of the art design. The tool is however a proof of concept only, and would need further work to be taken to market.

7.6.2 Legal and regulatory issues

For liability purposes TL is a software application (if run on an IaaS or PaaS basis), a service (if offered as SaaS) or a component of a wider service. These have been analysed already in relation to the A-PPL Engine, and we think the same analysis will apply here. There are no special liability issues which cannot be dealt with through a combination of disclosure and appropriate terms.

The role of TL in legal and regulatory compliance is a limited one. It is primarily a mechanism for logging and communicating information; where there are obligations to do this (eg breach notification) no particular mechanism is specified. However, the privacy-enhancing elements of communication via TL may well assist DCs to avoid further breaches during the communication process.

DPA's will find TL particularly interesting in relation to auditing. If TL is implemented as a logging mechanism for other accountability tools, it might become the primary source of information for much of the technical auditing. DPA's will be looking for ways to interface their auditing tools with TL, and this is something which might be considered as a potential marketing advantage if the tool is commercialised.

7.7 Experimental status of the A4Cloud tools

There is one important fact which readers of this deliverable must bear in mind. All these tools are research outputs designed as proof of concept. They are not intended for use in their current form, and substantial further work would be needed to take them to that level.

¹³³ Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC, OJ L 257/73, 28 August 2014.

The analysis above is therefore not a description of “defects” in any tool. Rather, it identifies the areas which will need further work, from a legal and regulatory perspective, in order to make these tools work more effectively and to take them to market.

8 Business models to take the tools to market

The analysis in this Deliverable has attempted to show that all the A4Cloud accountability tools are likely to be useful in assisting in achieving legal and regulatory compliance, but that they cannot on their own do this. Compliance is a mixed process, in which management, operations and oversight play an ineradicable part.

We have also shown that these tools will be of interest to DPAs, particularly in relation to the exercise of their audit functions. For this reason we have suggested ways in which some of the tools might be enhanced for these purposes when taken to market.

We can see no major obstacles in data protection law which should inhibit further development of the tools.

Liability, though, is a major concern for some of the tools. There is an important distinction to be made between those tools which appear to the user to be performing data processing/control and communication functions (Data Track, A-PPL Engine, Transparency Log) and those which might be seen as offering analysis and advice (COAT, DPIAT, possibly Acclab). The former group can probably safely be taken to market using normal software distribution and supply techniques, though we recommend a high degree of transparency about the safe methods of using the tools and about their limitations in order to ensure that users understand the risks of using these innovative tools.

The tools which appear to offer analysis and advice are more problematic. If they were made available only to sophisticated users, who would use them in conjunction with professional compliance advice, it could be safe to market them in the same way. But the likely market is SMEs, and their approach to the tools will be very different:

- SMEs are unlikely to research the functions and limitations of the tools in depth, but rather will treat their output at face value, ie as providing “the answer” to their compliance problems. Warnings and explanations, particularly if given online rather than in an operations manual, may be ignored in the same way that legal terms and conditions are rarely read.
- SMEs rarely have in-house professional compliance advisers, and find it too expensive to use external advice for everyday matters. Thus they are unlikely to receive advice which will help them to understand the limits on the role of the tools in legal and regulatory compliance.
- The contractual allocation of risk via contracts and licences may be ineffective where SMEs are involved. Many EU member states’ laws treat SMEs more like consumers than large businesses, and there is therefore a real risk that these contractual controls may be ineffective in some places.

We see two possible, though partial, strategies to address this dilemma.

The first is to take these tools to market solely via third party solutions or consultancy providers, or make them available only to established cloud service providers. This would have several advantages from a liability perspective:

- Purchasers of the tools would be sophisticated users, likely to assess and understand their limitations.
- Purchasers will have existing expertise in legal and regulatory compliance, which will reduce their factual reliance on the tool outputs and also make it less reasonable for them to rely on those tools alone.
- Contractual terms allocating risk and liability are likely to be effective.

However, this strategy inevitably reduces the market for the tools.

Targeting the wider market, including SMEs directly, is high risk for the reasons explained in Sections 5 and 0. Suggestions have been given there as to how that risk might be reduced, and this second strategy adds a further layer of protection. Our suggestion is that for the wider market, the risky tools (and perhaps all the tools) could be marketed as suitable for use only by trained persons. The aim would be that use by those with appropriate training would be far less likely to lead to liability situations, and also that users who elected not to be trained might be treated by the courts as accepting the risks voluntarily. This might assist in persuading courts that contractual limitations (for example, excluding liability completely if the tools are used by untrained staff) should be upheld.

This strategy would create a market for training, though for competition law reasons there should of course be no attempt to exclude non-consortium members from providing training. Adequacy of training might be assessed through third party certification, eg by a body such as the CSA.

How far that training should go is a matter which will require further thought as marketing plans are being devised. There are certainly two elements of training required: technical understanding of what the tools are (and importantly, are not) doing and how they work together; and training in privacy law to understand the role of the tools in compliance. This second element of training might be achieved in partnership with a relevant professional organisation, such as IAPP. It would also be worth investigating whether DPAs are, collectively, interested in helping devise or validate training.

Neither of these strategies can solve the liability questions completely, and professional advice on the specifics of marketing will of course be essential. But we think the suggestions here might assist in risk reduction, and if so work will need to begin on them immediately.