CLOUD ACCOUNTABILITY PROJECT

D:A-2.2 Project Horizons report (final)

Deliverable Number	D12.2
Work Package	WP 12
Version	Final
Deliverable Lead Organisation	HP
Dissemination Level	PU
Contractual Date of Delivery (release)	28/02/2016
Date of Delivery	11/04/2016

Editor

Nick Wainwright (HP)

Contributors

Vasilis Tontopoulos (ATC), Mario Sudholt (EMN), Tobias Pulls (KAU), Martin Jaatun (SINTEF), Michela D'Errico (HP), Frederic Gittler (HP), Jean-Claude Royer (EMN), Christoph Reich (HFU), Anderson Santana de Oliveira (SAP), Rehab Alnemr (HP)

Reviewers

Anderson Santana de Oliveira (SAP), Frederic Gittler (HP)

Executive Summary

This is A4Cloud's Final Project Horizons report. A4Cloud is a European Framework Programme 7 research project developing principles, models and tools for achieving accountability in the cloud. This Report discusses the opportunities for take-up and commercialisation of the techniques and tools developed as during the A4Cloud project.

With the date for implementation of the new General Data Protection Regulation fixed (it will come into force across the EU in 2018) we consider business motivation for adoption of A4Cloud architecture and tools – the "stick" of the big penalties that may be applied for non-compliance, and the 'carrot' of greater efficiency and reduced cost of implantation. Two strategies for exploitation are considered – Top Down and Bottom Up. The top down strategy is most appropriate for large organisations who have the capacity to directly engage with strategies to meet obligations under the new General Data Protection Regulations. An alternative bottom up strategy for commercialisation discussed how tools may be separately commercialised building towards a greater degree of accountability with increased availability of tools.

A4Cloud has provided a way for organisations to think about accountability for data in cloud services by providing both a conceptual framework and an architecture that can help them move towards being accountable. This is a significant step – before the work of the project, accountability for personal data in cloud computing was not defined and no blueprint existed for how to get to become an accountable organization. A4cloud has published significant results on both topics and these will both continue to be available on the A4Cloud legacy website.

At the conclusion of the funded project an analysis of the potential for exploitation of the tools and technologies produced was completed by all parties who have been involved in developing the tools (section 3) to evaluate the path forward for each of the tools. To summarise, the A4Cloud tools and technologies are at technology readiness levels 3, 4 and 5, and significant number of these tools are available in open source form making them available past the end of the project for further development and exploitation.

Delays in the data protection regulation to 2018 has also delayed the demand for tools to address the issues that organisations will face as a consequence of its implementation. However we are beginning to see demand from large organisations to be informed as to what they need to do to get on top of the accountability requirements of the regulation ahead of time.

Note: this report does not address commercial plans of A4Cloud partners which remain confidential to those partners.

Table of Contents

1 Introduction

This report explores how the results developed by The Cloud Accountability Project may be exploited and taken up by the various stakeholders in the provision, use and governance of cloud computing. It is the second report from the Project Horizons work package and looks at the potential for exploitation and impact of the project results.

The main focus of this report is a tool-by-tool summary of readiness for each of the tools developed by the project along with a discussion of transfer technology to the market. A discussion regarding the adoption and use of the cloud accountability reference architecture is also included.

Project Horizons has been a continual activity throughout the life of the project. The first report looked at the landscape in which the project would be operating anticipating key external factors that would be significant during the lifetime of the project and would influence the outcome of the project. Additional activities carried out during the project have addressed some important cross-project topics including requirements for accountably from across the project and how to demonstrate accountability and the technical results effectively.

At the time we started our project, the general data protection regulation (GDPR) was drafted and proposed to the EU Parliament. We have used the draft as a guideline in our research and the creation of our compliance tools. The GDPR has been approved by the European Parliament and Council on December 2015 and expected to take effect on mid-2018, leaving two years period of transition. Cloud actors in the EU or providing services into EU member states continue to operate under the national regulations implemented under the 1995 Data Protection Directive, however they are expected to fully adopt the GDPR by 2018.

The external context for the projects work remains largely as discussed in the first A4Cloud Project Horizons report. Cloud continues to be a growing part of ICT provision and use and we see continuing interplay between public, private and hybrid cloud use as organisations debate the cost, security, and resilience issues of shifting to cloud. Privacy and data protection has continued to be a hot topic for the technology industries.

During the lifetime of the project, three notable events have kept privacy and data protection in the headlines: the Snowden revelations[1] regarding access by intelligence agencies to data stored in cloud services and passing through communication systems, the Right to be Forgotten ruling of the Court of Justice of the EU[2] requiring search providers to certain entries from search results, and the recent Court of Justice declaration that the Commission's US Safe Harbor Decision is invalid[3]. After several negotiations, on February 29th 2016, the European commission and the US department of Commerce have reached an agreement for the transatlantic data flow stating that the agreement has new strong safeguards. The new EU-US agreement, called Privacy Shield, has a new reformed EU data protection rules (which apply to all companies providing services on the EU market), high data protection standards for transferring data for law enforcement agencies, and a renewed framework for commercial data exchange.

As was noted in the first Project Horizons report, technology developments continue to play a huge part in the treatment of personal data in cloud-based IT services. Whether we are dis-cussing mechanisms designed to provide greater protections against unauthorized access such as novel encryption schemes such as Homomorphic encryption, or awareness of the ever growing capacity of those who seek to gain access, be they criminal or state actors, the discussion of security of cloud services continues to be at the forefront of consideration of the shift to cloud services.

As the research efforts of the A4Cloud project draw to a close one particular fact outside the control of the project has a significant influence on the potential impact of the project: specifically the delay to the EU Data Protection Regulation. At time the project proposal was developed, submitted and approved

[1] http://www.theguardian.com/us-news/the-nsa-files
[2] http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf
[3] http://curia.europa.eu/jcms/upload/docs/application/pdf/2015-10/cp150117en.pdf

for funding this was expected to have been concluded by the end of the term of the previous commission under President Manuel Barroso which ended in October 2014. With that timetable, those involved in cloud provision, use and governance would be looking at how to implement the requirements of the legislation during the second half of the project. Feedback from those involved in cloud provision, use and governance with whom we have interacted throughout the project indicates that the value of the A4Cloud results to effective implementation of the obligations that providers will face as they come to implement measures to comply with the new regulation.

It should be noted that this final horizons report is a public deliverable of the project. It addresses the exploitation and take up of the project results in general terms with respect to the readiness of the technology, market and regulatory context. However, the commercial exploitation of A4Cloud results by any of the partners is the responsibility of partners individually and remains confidential to those partners and is not covered in this report.

2 Towards Accountability for Cloud Computing

The aim of the A4Cloud project is to increase accountability as a property of the whole cloud ecosystem. The objectives of the project are to provide tools for cloud actors (providers, customer and those involved in governance) and an architecture that integrates these towards increasing accountability. The project has research socio-economic, legal and technical factors arising from data protection obligations, developed a conceptual framework, architecture and tools that can assist cloud organisations in being accountable, and integrated these into a demonstrator that shows how these can work together to increase accountability.

In this section we discuss the A4Cloud results and the readiness of the ecosystem and explore what should happen next. This comprises an overview of the reference architecture and tools[4], grouped by category, a discussion of incentives for adoption, and finally recommended next steps. A subsequent section included detailed information regarding each of the tools.

2.1 Overview of tool status

Table 1 summarises the status of each of the tools and the reference architecture, and gives technology readiness level and open source status where appropriate. A4 Cloud tools are proofs of concept and technical feasibility, at technology readiness level 3, 4 or 5 according to the scale used by the European Commission[5]. As such, further development of the tools towards market ready solutions requires investment towards productisation.

[4] *Cloud Accountability Reference Architecture*, and *Cloud Accountability Reference Architecture - The A4Cloud Toolkit* can both be downloaded from the A4Cloud website www.a4cloud.eu

[5] Technology Readiness Level definition from European Commission https://en.wikipedia.org/wiki/Technology_readiness_level#European_Commission_definition

Table 1 Summary of the tools status

Tool Grouping	Tool Name	Summary of tool status	TRL	Open Source
Framework	Cloud Accountability Reference Architecture	Practical documentation available from the A4Cloud website.	N/A	N/A (docs Creative Commons)
Contract and Risk Management	Cloud Offering Advisory Tool	Demonstrator system plus documentation covering contractual arrangements	5	No (docs Creative Commons licence)
	Data Protection Impact Assessment Tool	Demonstrator system plus documentation enabling impact assessment	5	No (docs Creative Commons licence)
Policy Definition and Enforcement	Data Protection Policies Tool	Prototype implementation of tools mapping privacy provisions to machine readable policy	3	No
	Accountability Lab	Prototype of tool for checking and validation of policies	4	Yes
	Accountable Primelife Policy Engine	Accountability Extensions to open source policy engine	5	No[6]
Evidence & Validation	Audit Agent System	Proof of concept prototype evaluated in A4Cloud demonstrator	4	Yes
	Data Transfer Monitoring Tool	Prototype implemented on a single IaaS platform (OpenStack)	4	No
	Assurance Tool (SPACE)	Framework for integrating data and app with conceptual demonstrator	3	No
	Assertion Tool	Technology that can be applied in the longer term	3	Yes
Data Subject Controls	Data Track Tool	Open source system with connectors to A4Clould demonstrator and to Google cloud	6	Yes
	Transparency Log	Open source system available for use	4	Yes
Incident mgmnt & Remediation	Response & Remediation Tool	Proof of concept prototype evaluated in A4Cloud demonstrator	4	Yes
	Incident Management Tool	Proof of concept prototype evaluated in A4Cloud demonstrator	4	Yes

[6] The extensions to this engine that have been done by the A4Cloud project are not open source, however it is built on the open source Primelife Policy Engine.

2.2 Incentives for adoption

In this section we discuss the high level framing for adoption of the A4Cloud tools and architecture.

2.2.1 Stick and Carrot: Risks and benefits of taking accountability based approach

As has been noted in other A4Cloud reports, the General Data Protection Regulation comes with a huge stick in terms of potential penalties for personal data breaches that cloud service providers and cloud service customers cannot ignore when thinking about business risk.

Whilst this has been on the horizon for some years (it's at least three years later than first anticipated when this research project was first conceived in 2011) the implementation of this regulation is now set for 2018 and this will be something that businesses cannot ignore.

From the viewpoint of the IT sector that serves business customers, we are beginning to see awareness of this regulation from our customers and a desire to be informed as to the potential impact on their business. This is motivated by awareness of the implementation of the regulation in two years' time.

Our expectation at the time we started the project was that this desire would move from those intimately involved with data protection (the data protection officers of large corporates for example) to the security and risk management organisations of those organisations, and eventually to those responsible for implementation and deployment of IT systems – the IT specialists – during the lifetime of the project. A consequence of the delays in the regulation is that demand has not yet moved from the governance to the operational side of the business.

On the other hand, the penalties for non-compliance with GDPR are not the only incentives for adoption of an accountability based approach and the practices and tools that support it. There are other benefits to businesses of implementing good business data governance practices which include for example:

- Operational efficiency – clear and well documented practices for personal information management supported by effective tools and automation of elements such as policy implementation can drive savings particularly for large enterprises.
- Supporting cloud contract management for the myriad cloud contracts in force can be useful for a company. Tools such as the Cloud Offerings Advisory Tool (COAT) support the supply chain management needs of businesses.
- Brand and reputation is a significant assets for businesses. Companies typically take considerable effort to communicate the steps they take to ensure their customers' assets are protected and that they are not exposing them to increased risk. Taking an accountability based approach can enhance the reputation of a company with its customers.
- Corporate Social Responsibility is an important factor in business management. Accountability and good stewardship of personal data are an increasingly significant factor in companies' values[7].

In summary – those involved with data protection are telling business operations that the EU is going to be wielding a much bigger stick than it currently has under current data protection regulation, but that message has not yet been actioned by the operational side of the business. A4Cloud's conceptual framework and reference architecture provide a way for organisations to engage with the need to be accountable for implementing new obligations that they will have to address as a result of the

[7] The recent dispute between Apple and the FBI regarding the former's unwillingness to unlock an iPhone used by a terrorism suspect is an example of the lengths companies will go to to protect their brand and reputations with their customers (*Beyond surveillance: what could happen if Apple loses to the FBI*
https://www.theguardian.com/technology/2016/mar/21/apple-fbi-encryption-battle)

regulation. A4Cloud's contribution is to provide a way for cloud actors who are being threatened with the big stick to understand their obligations and a route for them to implement the processes and practices that support it.

2.2.2 Top down vs. Bottom Up: Where to start?

A4Cloud tools are generally Technology Readiness Level 3-5. Further investment in developing the product is required to bring them to market. A consideration of the 'Stick' vs 'Carrot' discussion leads to two different strategies for exploitation which we call 'Top Down' and 'Bottom Up'.

A Top Down strategy is motivated by the needs of organisations to address the GDPR 'stick'. Over the next two years there will be increasing awareness of the risks to business that this will impose. Large organisations are well positioned to implement accountability governance within their organisations, and the reference architecture from A4cloud provides a blueprint for action. On the other hand, implementing a framework of technological controls within large organisations is complex, risky and costly. A top down approach would focus on implementing accountability governance and seek to develop experience and practice within the cloud supply chain of how to do this. A key part of this would be to understand the technical and organisational barriers to implementing tools support for accountability, and thus provide clear opportunities for tool vendors to target their business. In a top-down strategy, accountability is the goal and a coordinated effort is required of actors in the company. Technology and tools follow as organisations seek to automate and scale the process.

A Bottom Up strategy focuses on the individual tools in the A4Cloud toolset, considering the individual benefits to the customer of deploying one or more. In a bottom up strategy, accountability is a side effect of a focussing on the individual benefits of each of the tools. Given that A4 Cloud tools are proofs of concept and technical feasibility (TRL 3/4/5), it is worth considering what the necessary conditions are for progression of each of these tools or technologies.

One way to think about this is to consider technical barriers vs non-technical barriers. Technical barriers arise from the difficulty of deploying the tools or technology that has been created. Non-technical barriers arise from the need for clear incentives for adoption and the need to coordinate and get alignment between actors in the cloud ecosystem. High technical barriers arise for example from the complexity of deploying technology in cloud service. High non-technical barriers arise from the difficulty of getting agreement to implement common approaches, define standardised interfaces, or establish common specifications (Figure 1)

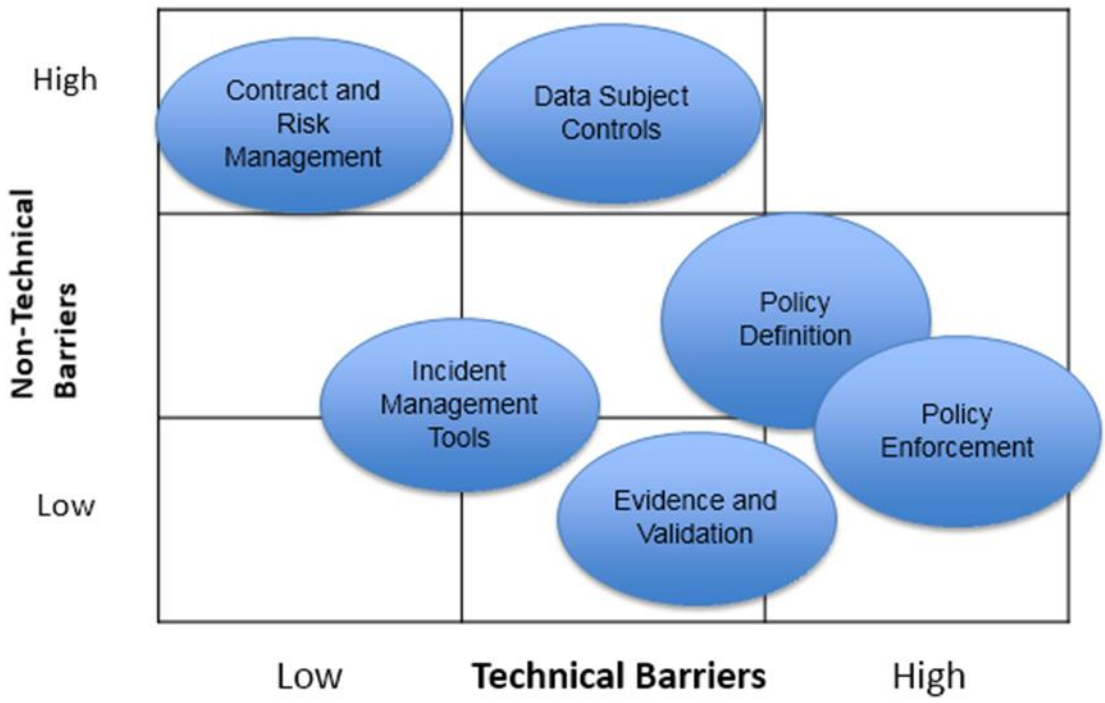

Figure 1 Technical and Non-Technical Barriers

- **Contracts and risk management** tools provide greater transparency to cloud users, but require greater transparency from cloud providers in sharing contractual information either in machine readable form or in a way that can be transcribed into a machine readable form. This would need to be available from a critical mass of cloud providers to give a useful degree of transparency to the cloud service customer. The technical difficulty in developing and using these tools is not high, the real barriers are in getting cloud service providers to be transparent about contract terms and to do so in a machine readable form, and to accept common ways to assess and compare risk. This is likely to be motivated through industry grouping such as Cloud for Europe, Cloud S Alliance, or Cloud28+[8] Technical Barriers: Low, Non-Technical barriers: High.

- **Policy definition & enforcement** tools can programmed to represent legal and regulatory obligations. However translating regulations and policies into machine processable forms, thus enabling automated enforcement, is tricky, requiring technical and legal/policy knowledge. With regard to policy definition, A4cloud has developed specific language that enable expression of accountability policies, and also prototyped easy to use approaches that enable the IT user in a cloud service providers/customer to define policies. Also establishing enforcement points where data is stored in databases requires deep technical integration with the database systems, which, given the number of big data NoSQL databases currently available and in use is a significant barrier to a general solution. Technical barriers: High. Non-Technical barriers: medium.

- **Evidence and validation** captures data and supports compliance and assurance. Capturing information from operational cloud systems requires significant integration with cloud systems. For

[8] http://www.cloud28plus.eu/

example A4Cloud's Data Transfer Monitoring Tool monitors for specific events in an OpenStack cloud infrastructure environment. Validation against is technologically complex. Technical barriers: medium/high, Non-Technical barriers: Medium/low

- **Data Subject Controls** provide the end user a view that any single provider cannot give, and anticipate greater access by data subjects to personal data held by cloud service providers as a consequence of implementation of the new GDPR. In the absence of any standardised and implemented approaches to extracting personal data from cloud service, tools of this kind have to struggle with extracting data from many different sources. Obligations on cloud service providers who hold personal data, acting as data processors further back in the supply chain may be visible to the end user making it difficult to see the full picture. Technical barriers: Medium. Non-Technical barriers: High

- **Incident management tools** - provide support to IT managers and individuals. Adoption improves transparency and accountability. The technical complexity of incident management tools is not great, requiring technical integration with cloud systems to ensure that incident notifications are connected to the incident management systems. Technical complexity: Medium/Low. Non-Technical complexity Medium/Low

2.3 Open Source

A significant number of the tools have been made available open source. This represents the best way to ensure that IP is available for exploitation especially given the make-up of the A4Cloud consortium which includes two large and one small enterprise, 9 universities and research organisations and the cloud security alliance.

2.4 Reflections on Socio-Economic Impact Assessment

A4Cloud carried out an assessment of the socio economic impact of the project results which was completed contemporaneously with this report[9]. Figure 2 is extracted from the executive summary and summarises the final recommendations from this study.

Reflecting on the conclusions of this study it is apparent that actions will need to be taken by European actors to support increasing accountability in how cloud services manage personal data. Actions such as pilots and trials of accountability based approaches would be well received by larger organisations. At the same time further standardization efforts and certification schemes are required to ensure that the framework is widely usable and available. This kind of activity can help develop the business case for technical mechanisms and tools that can be deployed such as those proved by A4Cloud.

The socio-economic impact analysis supports the idea of a "Top Down" as the next steps after the end of the project, focusing on large organisations and especially those working in the public sector. Specifically addressing the last point in these conclusions, the development of a business model encouraging greater uptake and use, our analysis of technical and non-technical barriers to adoption suggests that incident management and evidence management tools could be the best place to further develop business models to encourage greater update and use.

[9] A4Cloud DA4.1 Socio Economic Impact Assessment

1. Provide a stronger legal base for and enforcement of data protection and accountable behaviour
2. Increase public awareness of the need for accountability
3. Facilitate Independent Auditing of responsible data stewardship, especially in Small-to-Medium Enterprises
4. Focus on larger enterprises working in the public sector first, as these can serve as an example for other types of businesses.
5. Balance existing information asymmetries via partnerships
6. Demonstrate how A4Cloud tools and mechanisms can be turned into a business model in order to encourage greater uptake and use.

Figure 2 Summary recommendations of Socio-Economic Impact Assessment.

3 Analysis of exploitation of technical results

3.1 Methodology

In this section we discuss exploitation and uptake of each of the prototype tools and systems developed by the A4Cloud project. The approach taken to considering exploitation of the A4Cloud technical results is to analyse these tools along the dimensions described below:

Function
What is the main functionality of the tool, described as were being sold to a customer? Recalling that A4Cloud Tools are prototypes produced in a research project, it is necessary to anticipate the customers role in the cloud ecosystem, and also what type of entity they are, for example Large Enterprise, Small or Medium Enterprise, or Consumer. Here for example we might note that a Large Enterprise typically might have specialist staff assigned to address issues arising from Data Protection, whilst for an SME, these functions are probably carried out by someone who is not a Data Protection Specialist, but has to take on that responsibility in addition to many other responsibilities.

Innovation
What are the main innovation points in the tool? What makes it different than any other (commercial) tool? Innovation in A4Cloud tools may be technical innovation bringing new techniques to bear on a problem, or it may be through applying current technology to a clear insight into customer needs.

Impact
What impact the tool will have on the market if it becomes available? Here we seek to outline what effect a particular tool would have on the ecosystem – how does it change the way cloud actors operate? What does it bring to customers and users? As has been noted in the previous section, A4Cloud operates in the context of highly complex and changing regulation in Data Protection. Incentives to productise A4Cloud tools arise from Push Factors, primarily risks management associated with the need to comply with data protection regulations, and Pull Factors, the desire to improve one's position in the market either as a cloud provider or a cloud user, through offering a compelling proposition around data protection and data governance.

Commercial Readiness
Who is involved to taking the product to the market? And how? In this report we do not discuss the plans of any individual partner with respect to the technology developed within the project. That is of course a matter of commercial confidentiality. However it is useful to explore the steps involved in moving from current state through to commercial offering.

TRL (technology readiness level)

A scale from 1-9 to estimate tools maturity (Using European Commission scale[10]).

Improvement
If you have the time and resources, how would you improve your tool and why? A4Cloud tools have been designed and prototyped as part of a research activity. During the project we have explored the landscape of data governance in cloud ecosystems as it applies to data protection issues, and shared and demonstrated tool concepts with many interested parties.

Next Steps
Finally, we consider what next steps should be taken to ensure successful exploitation of the tool. Again it is important to note that we cannot discuss specific business or research plans of any of the parties involved in the project in this document, due to commercial confidentiality issues.

The cloud accountability Project has prototyped a large number of tools and has developed an architecture for cloud accountability. Figure 3 provides an overview of the major tools categorised according to their relevance to Cloud Actors and their function in terms of implementing preventative, detective or corrective controls.

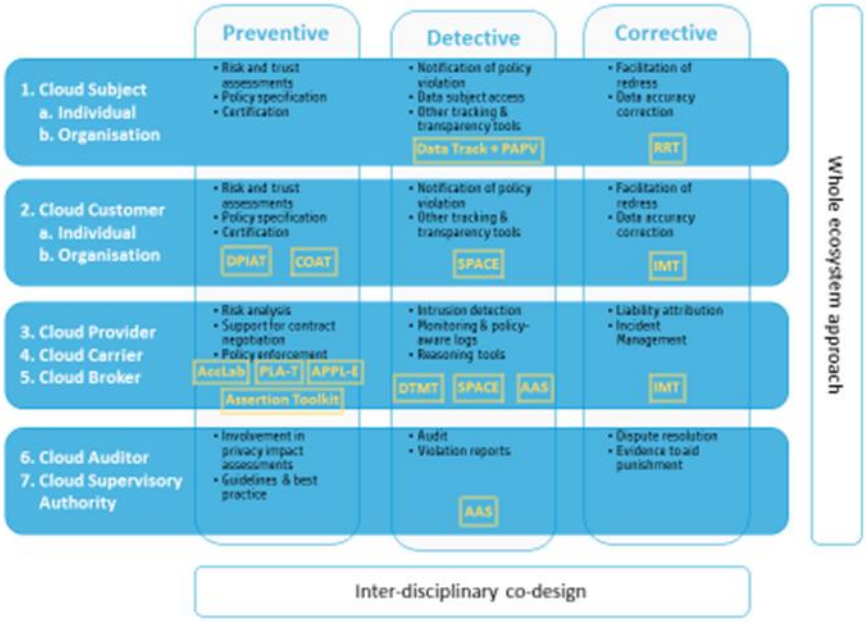

Figure 3 Overview of A4Cloud Tools

For the purposes of this analysis we consider the results in a more functional categorisation which maps more closely to where tools would be positioned in the market.

[10] https://en.wikipedia.org/wiki/Technology_readiness_level#European_Commission_definition

3.2 Cloud Accountability Reference Architecture

The Cloud Accountability Reference Architecture has been developed to provide a means for organisations to move towards being accountable. To understand the significance of having a reference architecture for accountability, one should put oneself in the place of a CEO, CIO, or COO of a company that is providing, or using, cloud services who has just become aware of the implications of the new GDPR on the organisation and of the obligations and potential risks that result from it. She/he will immediately be asking – what do I need to put in place in my business/organisation/team to address this? This will be relevant for large and small enterprises, public and private sector organisations alike. The reference architecture is the starting point for an organisation that has understood that it must be more accountable.

Table 2 Cloud Accountability Reference Architecture

	Cloud Accountability Reference Architecture
Function	The Cloud Accountability Reference Architecture (CARA) provides an abstract but powerful model for designing accountability in modern cloud and future Internet ecosystems. This model addresses both the needs of organisations which seek to behave in an accountable manner, and the exchange of accountability-relevant information between stakeholders in the context of the cloud. To this effect, CARA describes: • A set of action principles, a control framework, and a collection of best practices to be adopted by accountable organisations • A maturity model associated with the above • An extensive analysis and practical guidance of the account, the main instrument used to demonstrate accountability • A lifecycle view of the actions required to create, operate, and evaluate accountable solutions • An inventory of the accountability artefacts exchanged between cloud actors, along with the list of associated services and processes • The concepts developed in CARA can be directly used by two groups of stakeholders: • For the action principles, control framework, best practices, maturity model, and the account: organisations (including SME) which provide or use cloud services and aim to improve their level of accountability • For the lifecycle, processes, artefacts and services: researchers, architects and developers which are building cloud-related frameworks and aim to support accountable solutions
Innovation	CARA is the only known IT-focused reference architecture which addresses accountability in a holistic manner and remains agnostic to the object of accountability. CARA leverages prior art addressing accountability in a focused manner; e.g. guidance on selecting cloud providers for the processing of protected data. CARA is also unique in addressing both governance and technology for accountability. This reflects the fact that accountability is primarily a governance concern, but addresses the trustworthiness characteristics of the services provided across a value chain.
Impact	Organisations which provide or use cloud services and aim to improve their level of accountability will find practical guidance in CARA to realise their objectives. This is particularly true for SME, which do not have the means to adopt complex and expensive frameworks. Behaving in an accountable manner will be within reach of a larger number of organisations. Enterprise architects and developers concerned with accountability will be able to leverage the analysis and building-blocks identified in CARA to integrate accountability in the solution architecture they instantiate.

	Cloud Accountability Reference Architecture
Commercial Readiness	CARA is disseminated as a public document, under copyright of the A4Cloud consortium partners. No patents have been filed by any of the partners in regards to the content. The methodologies described in CARA are available for anyone to apply and deploy. CARA will be featured on the project website (http://www.a4cloud.eu) and be available as both a web-based document and a PDF file.
TRL	An unaudited self-assessment, based on the descriptive text associated with each level, leads to the following maturity assessment: • SME guidance (aka action principles) – TRL2 (these are new and have never been applied) • Simplified Control Framework – TRL3 (these have been validated against CCM) • Best Practices – TRL4 (these have been abstracted from actual best practices) • Maturity Model – TRL3 (these have been developed from existing practice but never applied) • Account – TRL8 (this describes actual practices in the field) • For the lifecycle, processes, artefacts and services: TRL 2 It must be here noted that CARA does not provide technology in the traditional sense of the word; we could sustain that assigning a TRL to CARA is not a valid exercise. The TRL for the action principles, control framework, best practices, maturity model, and the account is provided as a range of TRL 3 to TRL 9 as it has been assessed based on the individual components of the area, which are not uniform.
Improvements	A training could be developed and delivered in the context of relevant privacy- or security-focused events in order to train potential adopters of the action principles, control framework, best practices, maturity model, and the account. This represents a business opportunity for training providers such as those who might be members of the Cloud Security Alliance ecosystem. Parts of CARA could be submitted as contributions to standards or to certification schemes (such as CSA's Cloud Control Matrix and the associated STAR programme).
Next steps	HPE will continue to work with CSA to have some of the controls adopted in CCM. Overall, it would be appropriate to trial accountability architecture in a pilot or trial with operational cloud services.

3.3 Contract and Risk Management Tools

The Cloud Accountability Project developed Contract and Risk Management tools with the aim of ensuring that cloud customers can be better informed as to the data protection risks and how these are addressed by cloud service providers. In addition, the DPIAT tool addresses a specific need that SME cloud users will have in being able to carry out the data protection impact assessment that will be an obligation under the new data protection regulations.

These tools fall into the category of Governance Risk and Compliance (GRC) tools aimed at cloud providers and users. The tools developed by A4Cloud are broadly applicable across the ecosystem and increase transparency in cloud ecosystems. However, increased transparency has to been seen by all actors as a business case that has benefits for cloud providers and customers alike. Sharing information regarding service offerings in forms that can be used by tools such as these is critical to the effectiveness of the tools.

This topic is being explored in the both the Cloud Security Alliance community which looks at certification and assurance, and through the Cloud28+ community that has been established. Cloud28+ is “Cloud28+ is an open community of Cloud Service Providers, Cloud Resellers, ISVs, Systems Integrators and government entities dedicated to accelerating enterprise cloud adoption across Europe, Middle East and Africa”[11]. Cloud28+ was founded by Hewlett Packard Enterprise.

[11] http://www.cloud28plus.eu/

Table 3 Data Protection Impact Assessment Tool

	Data Protection Impact Assessment Tool
Function	The Data Protection Impact Assessment Tool (DPIAT) is a decision support tool that identify the main risks of a project with respect to the rights of data subjects concerning their personal data. It is a systematic process to elicit threats to the privacy of individuals, identify the procedures and practices in place to mitigate these threats, and document how the risks were addressed in order to minimize harm to users. These were tailored to satisfy the needs of Small and Medium Enterprises (SMEs) that intend to process personal data in the cloud. The approach is based on legal and socio-economic analysis of privacy issues for cloud deployments and takes into consideration the new requirements for DPIAs put forward in the European Union (EU) General Data Protection Regulation (GDPR).
Innovation	The main innovation is the process of automating the data protection impact assessment with a questionnaire that is aligned with the new GDPR and identifies to what extent an SME complies with the new regulation. This, combined with a plugin to assess the security controls implemented by cloud service providers and determining risks associated with using specific providers, presents a novel approach in privacy impact assessments. The tool has a user centric design, facilitating understanding and also educating users about privacy risks.
Impact	The tool has potential to become the *de facto* standard for DPIAs for SMEs. Most SMEs are not aware of the obligation to perform DPIAs in the near future. As soon as the GDPR becomes effective, the market demand will be considerable, especially for “Cloud born” projects oftentimes brought by start-ups across the globe.
Commercial Readiness	The tool targets SMEs as its user. Whilst the IP in the DPIAT code belongs to HPE and SAP, however the questionnaires, which are the result of detailed research in this area, that are used in the tool are public. These have been made available so that the knowledge of how to conduct data protection impact assessments is widely available to the ecosystem.
TRL	TRL 5. The tool was validated with several potential customers.
Improvements	Currently the tool obtain the information about companies’ security controls statically from the CSA STAR Registry. With time and resources, potential improvement of the tool is to automate this process to automatically obtain up-to-date results.
Next steps	We have produced a document that that contains all the research and information gathered to build the tool. The document enables a smooth knowledge transfer to those who wish to build a similar tool – such as SMEs - or want to exploit the research done during the design and implementation of the tool. The document will be available in the A4Cloud legacy website.

Table 4 Cloud Offerings Advisory Tool

	Cloud Offerings Advisory Tool
Function	COAT is a cloud brokerage tool that matches user's non-functional and data protection requirements such as transparency, legal terms, court of choice, privacy and security, etc. with the contract terms in cloud providers' service offerings. The tool has a list of attributes (the particularities of a certain cloud provider) that can be used in this matching based on the analysis of standardized cloud contracts, service level agreements (SLAs), the EU Data Protection Directive (DPD), and the new EU Data Protection Regulation (GPDR). For cloud customers, it eases the comparison of alternative cloud offerings, makes the contract terms transparent for them, and it offers guidance in the meaning and importance of each attribute and contract term. For cloud service providers, the tool can increase their market exposure and offer them the chance to highlight a unique criteria in their offer.
Innovation	Some of the aforementioned non-functional requirements are in the contracts but they are not clearly categorized and used in offers-requirements matching. COAT uniquely focuses on achieving compliance by doing the matching based on well investigated data protection requirements while presenting the user with a familiar store-like interface to reduce complexity. One of the unique features of the tool is the guidance and information on the requirements that encourages understandability and enhances transparency.
Impact	So far most cloud brokerage tools in the market do not make the matching based on data protection requirements. With the new EU GDPR, users need to be more informed on their data protection practices especially the selection of a compliant service provider. As such, COAT will present a unique market value.
Commercial Readiness	A4Cloud has produced a prototype tool. Another option is, creating a brokerage SME which uses COAT as their main business idea and will be responsible for the development and maintenance of the tool. The right to COAT's code belongs to HPE however the information about the non-functional attributes used in the tool is public
TRL	TRL 5, the tool was validated with several potential customers
Improvements	With time and resources, potential improvement of the tool is to automate the reasoning and discovery of contract terms. Creating a legal ontology for the contract terms and using a machine-readable knowledge representation for the contract terms is also a potential research and development activity around the tool.
Next steps	We have produced a document that that contains all the research and information gathered to build the tool. The document enables a smooth knowledge transfer to those who wish to build a similar tool – such as SMEs- or want to exploit the research done during the design and implementation of the tool. The document will be available in the A4Cloud legacy website.

3.4 Policy Definition and Enforcement

Policy Definition and enforcement for the purposes of data protection is a very broad topic A4Cloud results have developed technical advances that have demonstrated that complex polices can be expressed by non-specialists and mapped to machinery that can implement the policies and detect violations. Specific machinery for implementing policy controls deep within the cloud service provider's infrastructure is specific to the database platform being used. In A4Cloud project, the specific data management engine is the APPL-Engine. Further, the A4Cloud system links policy specifications introduced in this way to audit, incident detection, and assurance.

Table 5 Data Protection Policies Tool

	Data Protection Policies Tool
Function	DPPT is a tool that Cloud Service Providers (CSPs) can use to create a machine readable privacy policy statements specification, automatically generate a technical policy file and deploy it to the component in charge of the enforcement of the policies. The tool offers a GUI which presents a set of data protection aspects such as personal data processing (personal data that will be collected), personal data retention, sending of notifications about data processing related events (notably data breach events), and access control over the data collected
Innovation	The automation of the technical policy authoring task, by the means of a template-based mechanism which maps a policy statement to the template to be instantiated. A second point of innovation is the creation of a machine readable policy statements specification that enables automation of the extraction and processing of information about the policy terms.
Impact	Technical policies authoring, when performed by humans, is prone to errors and cannot be used at scale. The DPPT tool facilitates this task, in that the technical policies are automatically created based on the policy statements selected and specified by the CSP through the graphical interface offered by the tool. The tool also generates a machine readable representation of the policy statements that can be used to extract information about the policies implemented by the CPS. This representation enables an efficient processing of policy terms that otherwise should be performed by humans, who are required to read and analyse language policy documents
Commercial Readiness	The tools is a proof of concept that machine readable policies can be created through user friendly GUIs
TRL	TRL 3 - The tool has been developed as a prototype and integrated in the A4Cloud demonstrator.
Improvements	The coverage of the policy aspects could be extended by taking into account policies to be enforced at the level of the infrastructure resources, which would require creating different language adapters able to translate a policy statement into actions to be performed on a specific infrastructure resource.
Next steps	The tool has been developed as a GUI-based standalone prototype. After the end of the project the backend of the tool, which provides the mechanism to bind policy statements to a technical policy template, should be provided as a set of extendable RESTful APIs. In this way, CSPs could enrich the set of policies covered based on the capabilities of their enforcement environment, and create their own customized GUI.

Table 6 Accountability Lab

	Accountability Lab
Function	AccLab is an environment to support accountability policy writing, checking and concrete enforcement, it is mainly dedicated to privacy officers. Currently there is no formal language to express accountability policies in an abstract and readable form. It is impossible to check correctness with natural language privacy statements thus we propose a formal approach. AccLab provides support for policy conflict detection and to formally check the compliance between policies
Innovation	AccLab helps to consider accountability from design time and provides the first formal domain specific language for accountability which is rooted in a rich formal temporal logic framework. Mainly it allows specifiers to check for policy consistency and compliance. The tool support is based on a web user-friendly interface with dedicated templates and a sophisticated editor with syntax highlighting and auto-completion. A component diagram editor is helping users in describing its system and generating accountability templates
Impact	There is no tool which provides technical accountability policies writing, and checking. Some tools provides matching of natural language statements (privacy, security) based on ontology or more ad-hoc technics. AccLab relies on a more flexible and rigorous technique and its efficiency is comparable with other state of the art tools (model-checkers and sat solvers).
Commercial Readiness	Most of the information behind AccLab is public and the code is open source. Our aim is to transfer it to a software editor in charge of building a real product. Of course our research team is able to provide assistance in this transfer.
TRL	TRL4: it has been checked on several use cases during our experiments.
Improvements	Our tool will improve the conflict detection mechanism, since it is a critical requirement for end users. Currently it can provide a concrete enforcement of policies based on a simple translation to XACML. But this standard is debatable and more rigorous proposals are expected. We targeted a more abstract approach and we have yet a runtime verification and monitoring system which can be integrated into AccLab
Next steps	Scientific challenges behind AccLab are parts of our ASCOLA team research activities thus we will continue the extension of AccLab after the end of A4Cloud. Regarding the software maintenance and future development, a PhD student is working on it until the end of 2016. This will provide public visibility via github and a complete documentation as well as a report thesis and several research papers. There is also a possible connection with the AssertionTool and we expect to make it explicit through a future research project

Table 7 Accountable Primelife Policy Engine

	Accountable Primelife Policy Engine
Function	In the cloud, mechanisms to automatically enforce organizational and data governance policies are fundamental for compliance management. Using the A4Cloud policy enforcement tool, Cloud providers can offer more transparency about the data handling, and help to prevent privacy breaches. An independent party, providing assurance about the data handling, can audit the deployment and configuration of the tool. Then, policy enforcement will happen in a predictable manner, satisfying the data controller needs and obligations, as determined by the data controller, giving transparency to the (cloud) data subject – since policies have clear semantics. The enforcement engine works in cooperation with further A4Cloud tools to reinforce the assurance about the correct policy execution.
Innovation	It is based on previous standards for access control, i.e. XACML, which was extended with usage control and accountability features. It automatically enforces obligations based on time-based triggers, or a series of other events. It is easily auditable, producing traces to the data controller (or yet to the Data Protection Authority), but also towards the data subject.
Impact	It will drastically reduce the risk of non-compliance and improve consumer trust in the services using it. A-PPLE is a powerful accountability mechanism guaranteeing that if the data controller uses it as a gateway to personal data, there is adequacy between data collection and purpose of the processing.
Commercial Readiness	Many consortium partners have participated to the design, implementation and integration of the tool (For example, Eurecom, SAP, and ATC, among others). There is already adoption of its core components for enforcing automated actions in an Internet of Things prototype where Karlstad University is involved. SAP has also used part of its implementation for prototyping declarative authorization enforcement in some products
TRL	TRL 5- technology demonstrated in relevant environment (industrially relevant environment in the case of key enabling technologies): The tool has been tested in a cloud environment similar to commercial offers.
Improvements	It would be interesting to improve performance for faster data storage (as data is stored together with rules and obligations governing access to it), and retrieval (for the same reason).
Next steps	All elements are at hand, since the new regulatory framework (EU GDPR) puts accountability as a central concern. The documentation of the tool is up-to-date. The code is available to all partners. If anything else should happen is to create services embedding the tool. The main motivation would be customer demand. We believe there will be.

3.5 Evidence and Validation

A4Cloud's evidence and validation tools span a range of functions from monitoring for specific incidents that may result in data protection incidents, automated collection of evidence (Audit Agent System) to providing assurance to cloud users and providers over security and privacy properties of the system (Assurance Tool) and the ability to make assertions about the validity of the tool chain deployed (Assertion Tool).

Table 8 Audit Agent System

	Audit Agent System
Function	AAS is a cloud auditing and monitoring tool. It enables the automated collection of evidence and execution of technical audits in multi-tenant and multi-layer cloud applications. Furthermore, it enables the assessment of privacy and security rules within a cloud provider's service. The tool provides a set of audit tasks that combine evidence collection and evaluation for a pre-defined rule. For cloud auditors, it automates many aspects of technical audits such as the continuous collection of logs and their analysis in order to uncover potential policy violations
Innovation	Conventional audits rely extensively on manual processes such as documentation reviews, log analysis, configuration setups, etc.. Furthermore, intervals in which audits are performed are usually quite long, which leaves room for undetected violations happening during that time.
Impact	Continuous monitoring tools are already quite sophisticated. However, our tool provides the ability to continuously assess compliance of cloud operations based on policies that are actually agreed-upon between cloud providers and their users or imposed by regulation. It is highly scalable and follows the principle of data minimization by collecting only evidence needed to detect policy violation
Commercial Readiness	Cloud providers can extend the prototype implementation of the tool and integrate it in their cloud infrastructures in order to enable continuous automated technical auditing alongside their monitoring solutions. Software companies developing software for auditors can use the prototype to assist auditors to do their audit of Cloud infrastructures. The prototype implementation of the tool will be made available to the community under a liberal open source license to be further developed
TRL	TRL 4: Technology validated in lab
Improvements	With time and resources, the tool can be extended with more audit tasks in order to cover more rules and machine-readable policy languages as well as more evidence sources in an audit. Also, automated reaction on incidents detected by our tool could be investigated. Improvements on the audit report user interface are desirable
Next steps	We have produced several documents that contain information from a research and an implementation perspective. Additionally, we have developed a prototype implementation that can be used by those who want to build a similar tool. The documents and the prototype sources will be made available on the A4Cloud website

Table 9 Data Transfer Monitoring Tool

	Data Transfer Monitoring Tool
Function	A detective tool that tracks operations over the cloud infrastructure layer as to identify if intentionally or accidentally the processing is being carried out by authorized processors in allowed geographical reasons. For that, it intercepts

	events at the network level and infers, using rules and policies if the event signals irregularities. The tool operates on an OpenStack Infrstructure.
Innovation	It is the only tool doing such thing, according to our knowledge. Intrusion detection systems can be perhaps configured to offer a similar function, but that would be cumbersome.
Impact	By adopting the tool cloud providers can be more transparent and accountable towards their customers. They may win some market share.
Commercial Readiness	SAP developed the prototype. The documentation and source code are available to the partners. The whole consortium may help to promote it
TRL	TRL 4 - The tool has been tested in a cloud environment similar to commercial offers
Improvements	Increase stability – some component depend on native libraries that do not behave equally in all environments. Expand the scope, detecting other potential compliance problems in cloud environments
Next steps	The Safe Harbour agreement between the EU and the US is being replaced by a new arrangement, Privacy Shield, awareness about the location of the data processing will be more important to data controllers. The DTMT tool address this question very well, in particular if associated to the Transparency Log tool. These changes can represent a market opportunity. The best that could happen to the tool exploitation is that the new deal imposes continuous monitoring in addition to certificates

Table 10 Assurance Tool

	Security & Privacy Assurance Case Environment (SPACE)
Function	Cloud Service Providers and Cloud Service Customers seek increase confidence that security and privacy controls are in place and operational across the entire cloud service supply chain. SPACE provides a dashboard that provides continual assurance over operational cloud services
Innovation	SPACE provides systemic support for evidence-driven continuous assurance. Combining assurance and operations, SPACE provides the link between organisation's security and privacy policies expressed as assurance claims, and the underlying systems and controls that are in place.
Impact	Cloud providers and customers can be more confident that security and privacy controls are in place, not just at audit time but continuously.
Commercial Readiness	SPACE is a proof of concept prototype that has been developed to integrate evidence across the A4Cloud demonstrator and toolset and demonstrate feasibility of the approach through an assurance system and dashboard. Further technical research and pilot trials are required before this can be ready for deployment. In addition to readiness of the SPACE framework, the connectors to feed evidence of controls and incidents to the SPACE repositories need to be deployed which requires investment by cloud software developers and cloud providers.
TRL	TRL 3 - experimental proof of concept
Improvements	Much further work is required, this will be the subject of future work.
Next steps	A full assurance system needs to be developed before this can be used in trials and pilots.

Table 11 Assertion Tool

	Assertion Tool
Function	The assertion tool (AT) enables the validation of tool chains that enforce accountability properties through runtime verification. It enables tool developers and cloud providers to validate the correct working of tool combinations by the execution of validation scenarios that include assertions of accountability properties. The AT may actively exert control over the accountability tools in order to verify validation scenarios or may passively observe the tools' executions and check their conformance to the validation scenarios
Innovation	Existing approaches for the testing and monitoring of accountability properties are limited to individual tools and fixed sets of typically small numbers of properties. In contrast, the assertion tool allows the definition of validation scenarios verifying properties involving an arbitrary number of tools. It also supports a rich set of policy-based or ad-hoc properties expressed in terms of general or tool-specific accountability assertions. The AT supports the validation process by automatic strategies for the collection of evidence (such as logs and notifications) as well as the implementation of high-level accountability process in terms of tool-specific ones
Impact	Currently, no tools (commercial or otherwise) exist for the validation of tool chains for accountability. Furthermore, the AT is unique in providing a flexible assertion language that allows the definition of a wide range of accountability properties, including transparency, responsiveness and responsibility properties. It can be used by several actors in Cloud ecosystems, tool developers and cloud providers principally. The AT tool can be adapted for use with other tools than the A4Cloud tools through a dedicated tool configuration interface.
Commercial Readiness	Most of the information concerning the AT is public and the code is open source. Our aim is to transfer it to a software editor in charge of building a real product. This should be done together with the other partners of the consortium
TRL	The AT tool's readiness level is TRL 3 (with some demonstration available). It has been validated and demonstrated in two OpenStack-based Cloud environments set up as part of the A4Cloud project
Improvements	The main point for improvement is the full distribution of the validation process. Currently, the AT tool provides web service functionality in order to be able to obtain information from the to-be-validated accountability tools at remote sites. The AT tool itself operators, however, as a central tool that interacts with the remote tools.
Next steps	We intend to use and develop the AT tool as part of two new French projects that will start in Q2 2016, one project on security and accountability properties that must be enforced on the virtualization level of OpenStack and another one on privacy and accountability properties for shared genetic data. As part of the A4Cloud project, two tools have been developed for the validation by runtime verification: the AT tool and the AccMon tool, an extension of the of the AccLab tool for the distributed monitoring of temporal formulae. We intend to integrate both in order to benefit from the general definition of properties provided by AccMon as well as the accountability predicates and third-party tool integration facilities of the AT tool.

3.6 Data Subject Controls

Table 12 Data Track Tool

	Data Track Tool
Function	The Data Track is a tool for end-users (data subjects) that provides answers to two pressing questions: “what information have I sent to which online service?” and “What information do online services have about me now?”. DT connects to online services, on behalf of a user, and requests access to all data stored in the service for the user. DT then provides a number of playful visualizations of all data. The user can also request that data is corrected or deleted services-side from within the DT tool.
Innovation	DT gives users insight into who knows what about them online through several easy-to-use interfaces. Users are also empowered online, by enabling them to exercise their rights online to access, correct, and request deletion of their personal data at online services. Today, users access these rights using analog tools often in the form of pen and paper. DT takes these rights online.
Impact	Primarily, DT may influence service providers to allow data subjects to exercise their rights with respect to their personal data online. The basic API needed by DT at service providers can serve as a foundation for future work on standardizing such an API for data subject access. Secondary effects include an increased awareness among users of online services of both the vast amount of personal data collected about them and that it’s not unreasonable to expect to influence the collection and processing of their personal data.
Commercial Readiness	Karlstad University (KAU) has open-sourced two versions of DT under a permissive license. We have no plants to commercialize DT.
TRL	For the A4Cloud-specific DT release, we reach TRL 3 due to integration in the A4Cloud toolset. For the standalone version of DT, we reach TRL 6 thanks to working with real Google Takeout data.
Improvements	The A4Cloud-specific DT version is tightly integrated with and limited by the dependencies from the A4Cloud toolset. The main issue is the A-PPLE DT API, which specifies both the API and the data format as was most convenient for A-PPLE. This API needs significant work on being more generic and to work how to tackle essential parts like user authentication. For the standalone version of DT, the primary area in need of improvement is support for more data sources. We currently only support a limited part of Google Takeout. I would focus on this area before working further on more visualizations
Next steps	Deliverables, published papers, and source code have been made available to the public for anyone who wishes to exploit the results surrounding DT. KAU plans to continue with some research relevant parts of the visualizations.

Table 13 Transparency Log

	Transparency Log
Function	TL is a tool that provides a secure and privacy-friendly one-way messaging system. Compared to. e.g., email or push notifications, TL protects both the content of the message and associated metadata (like who is the recipient of this message?) from third-parties. For distributed and dynamic settings, like in the cloud, TL also supports sending messages to potentially offline recipients anywhere in the cloud supply chain. This can be use, e.g., by a cloud service provider offering storage to another cloud provider to send messages to the data subjects whose personal data are being stored at the cloud storage provider. These messages can be used to send breach notifications, data handling descriptions, or other types of messages that increase the transparency of the service towards end-users.
Innovation	Both security and privacy have been taken into account from the start of the design of TL, resulting in strong cryptographically-based protections. TL is easy to integrate for both senders and recipients of messages. For end-users, TL provides a missing building block that can enable greater transparency and insight into what happens with their personal data in the cloud. Finally, TL provides publicly verifiable proofs of sender, recipient, content and time of selected messages. This increases the utility of all messages sent through TL, e.g., by enabling a cloud provider to prove that it provided timely breach notification to its customers or customers to prove the contrary
Impact	The technology underlying TL is closely related to Blockchain technology, which is currently being widely investigated in the market for everything ranging from banking to tracking intellectual property assets. The impact of TL would probably widen this investigation to also include transparency and accountability-related topics focused on end-users, especially given the upcoming EU GDPR.
Commercial Readiness	TL has been open sources by Karlstad University (KAU) under a permissive license and research around the core technologies of TL continues at KAU in other research projects. Tobias Pulls, at KAU, is working with research-colleagues in Belgium on a technology related to TL and a potential spin-off company.
TRL	TRL 4. TL has been widely used in A4Cloud by different tools but it is primarily a proof-of-concept implementation.
Improvements	Full rewrite of the code to take into account the latest design of Balloon (the authenticated data structure at the heart of TL). A proper storage back-end, probably using a mature database system, would also have to be properly investigated given the setting of TL. For the connections between the different components in TL, I would look at using something else than RESTful APIs over HTTPS, probably RPC or similar.
Next steps	Deliverables, published papers, and source code have been made available to the public for anyone who wishes to exploit the results surrounding TL. Part of the technology is, as mentioned, being spun-off at KAU and related technology is being investigated in other research projects at KAU (HITS and the EU H2020 project CREDENTIAL).

3.7 Incident Management and Remediation

Table 14 Response and Remediation Tool

	Response and Remediation Tool
Function	The Remediation and Redress Tool (RRT) targets individual data subjects who want to be made aware of any perceived incidents detected in the cloud environment that may have an impact on their personal data being collected and processed by provider operating in such environments. RRT facilitates provisioning of evidence to end users linking to incidents with an impact on their personal data and guides them through potential remediation actions in response to such incidents. RRT can be provided as part of the cloud service providers' set of tools to demonstrate their commitment to accountability. By enabling their end users to facilitate immediate incident response on their device, the cloud service providers will be well prepared in fulfilling their obligations towards the GDPR with respect to the notification of end users and their ability to request for a remedy.
Innovation	RRT fills in a gap among the existing tools in the market creating benefits for both individual end users and cloud service providers. RRT contributes to end users' empowerment by notifying them the incidents detected in the cloud with a direct impact on their privacy through a familiarised user interface and by supporting them in taking action in the occurrence of an incident in the cloud (simplifying filing complaints to DPAs, collection of evidence). Moreover, it assists cloud service providers in being more transparent regarding incidents, therefore, enabling them to comply with the obligations resulting from the GDPR. Through this way, the cloud service providers increase their transparency towards clients as they provide for incident notification and remedial options. Overall, RRT builds on the ex-ante effect of transparency and accountability, hence, fostering incident prevention in the cloud.
Impact	The unique proposition for RRT is that it increases end users' capability in responding to data breaches and other incidents detected in the cloud. RRT enables end users to exercise their rights by providing them with proof of what has happened in the cloud that can be used to support claims before courts. This empowers the confidence of the cloud users in the use of cloud services, while giving them the opportunity to play a key role in the formulation of the cloud service market with respect to the sustainability of providers with poor performance on accountability and low compliance with the applicable legal framework. Furthermore, RRT can influence the interaction between individuals and the competent supervisory authorities by facilitating filing of complaints. The tool strengthens the confidence of the cloud customers and their clients in operating their business with specific cloud providers, who are able to demonstrate a responsible behaviour for the protection of individual's privacy. The demonstration of the providers' capability to implement alternative controls to mitigate privacy risks, when an incident has detected, can increase the trust of the cloud customers and users due to the enhanced commitment to remediation.
Commercial Readiness	There are no commercial plans for RRT by ATC and TiU as we aim to offer the tool as open source (and under the Apache license) for anyone to improve and further share with the community. To this end, our aim is to approach cloud service providers and inform them about the tool benefits for them and their customers. We, also, aim at approaching cloud end user communities and raise their awareness on their rights about privacy and data protection in the cloud and communicating to them how RRT paves the ground towards this direction.

	Response and Remediation Tool
TRL	RRT is currently in TRL 4. A proof of concept prototype tool has been developed and the underlying technology has been validated in a lab environment.
Improvements	RRT leaves a lot of space for improvement towards reaching a commercially strength level (product release). The optimization of the currently supported functionalities is a key priority and the following improvements should be considered: ▪ Enable cloud users in filling a new request for remediation, based on self-assessment on the compliance of a cloud provider with the exhibited accountability policies. ▪ Develop a standards-based communication of RRT with tools generating evidence. ▪ Develop a framework to support the reasoning behind the remediation, based on the risks associated with the detection of a data breach / incident and the controls to be implemented to mitigate these risks. ▪ Produce interfaces to enable cloud users invoking redress actions
Next steps	A critical success factor for RRT after the end of the project is the level of adoption of the accountability framework from the community. The Cloud Accountability Reference Architecture is publicly available and highlights the need to introduce tools like RRT in the cloud service environment. All A4Cloud partners and their contacts should share the effort to bring accountability on the fore front and underline the requirement for cloud users to be able to respond to the cloud anomalies on their own capacity.

Table 15 Incident Management TOOL

	Incident Management Tool
Function	The Incident Management Tool (IMT) is a tool targeted at organizations and teams that handle computer security incidents – in practice any organization that provides or consumes an internet service. A problem experienced by incident handlers in the context of cloud computing, is the lack of access to sufficient incident information throughout the cloud provider chain. A SaaS would not necessarily receive the needed information from their PaaS, nor the PaaS from their IaaS, etc. Furthermore, complicated cloud provider chains with multiple participants increase the need for more automated sharing of incident information – potentially allowing some response actions to be automated. The IMT interacts with other instances of IMT and other tools by a simple, extensible incident format and a publish-subscribe based API – exchanging incident information. The simplicity of the solution makes it usable for small companies as well as large. The integration with A4Cloud tools, allows for easy notification of end users. The solution supports incidents propagating through the Cloud Service Provision Chain while preserving traceability.
Innovation	IMT provides a simplified incident format and incident exchange model that makes the solution usable for small companies as well as large. Existing tool such as TAXII (using STIX) are generally considered to be too complex for all but the most dedicated organisations. The integration with A4Cloud tools allows for easy notification of end users, thus making it easier to comply with the GDPR requirements of notifying end users when personal information is affected by a breach. Furthermore, IMT supports incidents propagating through the Cloud Provision Chain while preserving traceability. This allows organisations to share incident information to customers, not revealing where the incident originated, but still allow an auditor to follow the incident trail.
Impact	Being simple in its nature, IMT will be easy to adopt by any organisation. The notification functionality will help organisations in complying with the GDPR – both notifying the supervisory authority and the end user by allowing the relevant information to propagate down to the provider which actually manages end users and are able to notify them. Information sharing will be improved, which could lead to better incident handling across organisations.
Commercial Readiness	Given how crucial sharing of incident information is for improved security and incident handling in the cloud, and following the example set by threat intelligence sharing solutions such as STIX and TAXII, IMT is provided free of charge under the Apache v2 License. Any interested party can contribute and use IMT.
TRL	IMT has been validated in a lab setting, and is hence classified **TRL 4**.
Improvements	The two currently most missed features in IMT are workflow support and logging. The workflow support would make it possible for large teams to define how to handle an incident, improve collaboration and make it more difficult to make mistakes. Introducing logging would make sure that the operators of IMT are held accountable for how an incident is handled.
Next steps	The most important next step would be to set up a large-scale test of IMT, involving at least two organisations exchanging security incidents with each other. This is needed in order to validate and improve the concept and implementation, as well as demonstrate that the solution is helpful to incident handlers and the overall organisation.

4 Conclusions

The A4Cloud project has developed tools that automate the processes of accountability for managing personal data in cloud services. It has created a conceptual framework and reference architecture that organisations can use to become accountable for how the manage it, and has shown how these can be linked together – both at a process level and a technical level.

A4Cloud's contribution is to create a path forward for organisations to implement their obligations under the data protection regulations. This will become increasingly relevant over the next two years as the deadline for implementation approaches. We have considered top-down process driven strategies vs bottom up tool driven strategies for moving the tools from research to deployment. Overall we think that the top-down approach:

1. Make accountability for how personal data is used in cloud service part of the culture of cloud service providers

   One key aspect that has apparent through the duration of the project is that neither technical tools nor a set of processes will by themselves increase accountability within cloud ecosystems. For those engaged in cloud service provision, a culture of accountability should be established that drives accountability at all levels. The incentives for this come from two directions – obligations and penalties arising from the new data protection regulations soon to be in force across Europe, along with the risks of actual harm to the business in the event of a data protection incident. The work of A4Cloud does not alter those obligations, risks and penalties – but the research that we that A4Cloud has published and publicised does help organisations understand them more clearly and the Accountability Maturity Model that has been developed goes some way towards giving them a way to evaluate where they are on the journey towards being accountable.

2. Implement process and practice for accountable cloud organisations

   The formal objectives of the project put tools and technology as the first three objectives and architecture as the fourth. However, for a cloud service provider or cloud service customer, the next concrete step towards 'being accountable' is not to deploy technology and tools but to put in place the processes and practices of accountability across the organization.

   A4Cloud's contribution is the Cloud Accountability Reference Architecture. The projects conceptual framework and architecture describe the elements of accountability and the relationships between them, and the processes and practices that would have to be adopted by cloud service providers and users in order that they become accountable organisations.

3. Trial tools and technology to support accountable practices

   As a result of our research, A4Cloud project has produced thirteen tool prototypes covering all phases of the accountability lifecycle and addressing preventative, detective and corrective strategies. These tools are prototypes, demonstrating technology works in the lab and exploring it's applicability to customer problems, i.e. Technology Readiness Level[12] 3-5.

   The landscape into which these tools must be eventually deployed is a complex web of business and technological issues. For example to provide good information on service provider capabilities for data protection requires information regarding service provider contracts and capabilities to be available in machine readable form to be ingested into the contract and risk assessment tools that A4cloud has produced. Other initiatives in Europe are addressing these topics simultaneously with the A4Cloud project. A4Cloud implements policy controls on data stored in the APPL-Engine. However the number of big-data platforms continues to grow (see for example the proliferation of NoSQL platforms, generally defined as

[12] Technology Readiness Level definition from European Commission https://en.wikipedia.org/wiki/Technology_readiness_level#European_Commission_definition

> "Next Generation Databases mostly addressing some of the points: being non-relational, distributed, open-source and horizontally scalable"[13]). Implementing policy controls likely requires implementing these controls in more than one of these platforms. Similarly, to make progress with monitoring and evidence, to ensure that organisations can provide 'an account' requires the deployment of many data gathering points within and across organisations. The willingness to share information between actors is not proven.
>
> Therefore we conclude that the best way to further develop the business model for the tools and technology is to use them first in a real operational pilot.

In summary, A4Cloud has completed the objectives set out at the start of the project – to prototype tools increase accountability for how data is used in cloud services – tools for cloud providers, cloud customers, and those involved in governance and compliance management.

However this is not the end of the journey. The A4Cloud techniques and tools need to be taken out of the lab and tested in a real environment. Pilots with real live services and real personal data would test tools effectively. The A4Cloud governance lifecycle should be piloted by a number of cloud providers and customers. These next steps can best be taken through a collaborative pilot or innovation action focusing on governance of personal data in clouds.

[13] http://nosql-database.org/

5 Index of figures

6 Index of tables