



About ENISA

The European Network and Information Security Agency (ENISA) is an EU agency created to advance the functioning of the internal market. ENISA is a centre of expertise for the European Member States and European institutions in network and information security, giving advice and recommendations and acting as a switchboard of information for good practices. Moreover, the agency facilitates contacts between the European institutions, the Member States and private business and industry actors.

This work takes place in the context of ENISA's Emerging and Future Risk programme.

This report has been edited by Mr. Daniele Catteddu

Contact details:

Daniele.catteddu@enisa.europa.eu

Internet: <http://www.enisa.europa.eu/>

Legal notice

Notice must be taken that this publication represents the views and interpretations of the authors and editors, unless stated otherwise. This publication should not be construed to be an action of ENISA or the ENISA bodies unless adopted pursuant to the ENISA Regulation (EC) No 460/2004. This publication does not necessarily represent state-of-the-art and it might be updated from time to time.

Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication.

This publication is intended for educational and information purposes only. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication.

Reproduction is authorised provided the source is acknowledged.

© European Network and Information Security Agency (ENISA), 2010

List of contributors

This paper was produced by an ENISA editor using input and comments from a group selected for their expertise in the subject area, including industry, academic and government experts.

The views expressed in this publication are those of the editor, unless stated otherwise, and do not necessarily reflect the opinions of the participating experts.

Expert group members in alphabetical order:

Amanda Goodger, CESG, UK

Andrea Glorioso, European Commission (observer)

Prof. Antonio Lioy, Politecnico di Torino, Italy

Ben Katsumi, IPA, Japan

Daniele Catteddu, ENISA (chair)

David Wright, Trilateral Research & Consulting LLP, UK

Dennis Heinson, LL.M. (UCLA), Universität Kassel (provet)/Center of Advanced Security Research Darmstadt (CASED), Germany

Dr. Giles Hogben, ENISA

Prof. Fabrizio Baiardi, Dipartimento di Informatica, Università di Pisa, Italy

Jim Reavis, Cloud Security Alliance, USA

Liam Lynch, eBay, USA

Marcos Gomez, INTECO, National Institute for Communication Technologies, Spain

Prof. Dr. Milan Petkovic, Philips Research and Technical University Eindhoven, Netherlands

Dr. Paolo Balboni, Balboni Law Firm, Tilburg University, European Privacy Association, Italy

Dr. Peter Dickman, Google, Switzerland

Philippe Massonet, CETIC – Reservoir Project, Belgium

Raj Samani, McAfee, EMEA

Rui Barros, ELANET (CEMR) - European Network for eGovernment and Information Society (supported by the Council of European Municipalities and Regions)

Dr. Srijith Nair, British Telecom, UK

Dr. Theo Dimitrakos, British Telecom, UK

Steffen Schreiner, CASED, Germany / CERN, Switzerland

Table of Contents

Executive summary	6
1. Introduction.....	11
1.1. Structure of the report and how to read it.....	12
1.2. An introductory scenario: Making a decision	14
2. Objectives and analysis	22
2.1. Target audience	23
2.2. Analysis method.....	23
3. Model for decision-makers	25
3.1. Security and resilience parameters.....	27
3.2. Business and operational variables.....	36
3.3. Legal and regulatory framework.....	39
3.4. Architectural options.....	45
4. SWOT analysis	48
4.1. Public cloud.....	49
4.2. Private cloud	53
4.3. Community cloud	56
5. Example scenarios	59
5.1. Service description	59
5.2. Parameters and requirements.....	62
5.3. Comparative risk assessment	68
5.4. Selection of the solution and identification of threats and weaknesses	78
6. Preparing a request for proposal.....	80
7. Conclusions and recommendations.....	84
7.1. Recommendations to governments and public bodies.....	86
8. Glossary.....	90

9. References.....	97
Annex I – Legal analysis.....	99
Annex II – Scenarios.....	123
Annex III - Reservoir architecture description.....	137
Annex IV – List of threats	144

Executive summary

Cloud computing offers a host of potential benefits to public bodies, including scalability, elasticity, high performance, resilience and security together with cost efficiency. Understanding and managing risks related to the adoption and integration of cloud computing capabilities into public bodies is a key challenge. Effectively managing the security and resilience issues related to cloud computing capabilities is prompting many public bodies to innovate, and some cases to rethink, their processes for assessing risk and making informed decisions related to this new service delivering model.

This report identifies a decision-making model that can be used by senior management to determine how operational, legal and information security requirements, can drive the identification of the architectural solution that best suits the needs of their organisation. The main objectives of the report are:

- to highlight the pros and cons, with regard to information security and resilience, of community, private and public cloud computing delivery models;
- to guide public bodies in the definition of their requirements for information security and resilience when evaluating cloud computing service delivery models;

Moreover this report wants to indirectly to support European Union Member States in the definition of their national cloud strategy with regards to security and resilience.

The proposed decision-making guide helps the reader to compare community, private and public clouds, and to decide on the most suitable IT service deployment model, the controls to apply and the key questions to ask of a service provider in order to reduce the risks involved in migrating to the cloud to a level that is in accordance with their appetite for risk.

The analysis is based on three possible cloud usage scenarios: healthcare, local public administration and publicly-owned cloud infrastructure as a business incubator, as we have assumed that these use-cases are of particular interest for EU Member States.

The tool used in this report to compare the pros and cons of the security and resilience of community, private and public cloud models is a SWOT analysis which, for an informed risk-based decision, has to be used in conjunction with the security assessment described in the ENISA report *Cloud Computing: benefits, risks and recommendations for information security*. Public bodies should always undertake a thorough risk analysis of their specific applications in the context of the cloud model, and this report should be considered a supporting document and guide.

As a result of our analysis, we have concluded that the cloud computing service delivery model satisfies the most of the needs of public administrations, on the one hand, since it offers scalability,

elasticity, high performance, resilience and security. However, many public bodies have not yet built a model for assessing their organizational risks related to security and resilience. Managing security and resilience in traditional IT environments is very challenging for public bodies. Cloud computing presents some additional challenges. For example, understanding the shift in the balance of responsibility and accountability for key functions such as governance and control over data and IT operations, ensuring compliance with laws and regulations, and, in some instances, the poor quality of internet connectivity in some areas of the European Union (1).

This shift to indirect governance and control over data and IT infrastructure appears to be an inherent challenge in migration to the cloud model (especially with regards to public clouds and SaaS deployments) even though, as already stated by ENISA (eg, in its 2009 report), the situation can be improved by achieving transparency in the market and negotiating appropriate terms and conditions in contracts.

National laws and regulations in the Member States of the European Union currently impose some restrictions on the movement of data outside national territory; moreover, a problem exists in the determination of the applicable body of law (governing laws) when data is being stored and processed outside the European Union or by a non-EU service provider. The main questions that each public organization, and more generally each EU central government must address are:

- whether current legal frameworks can be changed to facilitate the communication, treatment and storage of data outside national territory without exposing the security and privacy of citizens and national security and economy to unacceptable risks;
- if so, whether moving citizens' data outside the national territory is a risk that may be undertaken;
- whether the trade-off between the risks of losing control over data and the beneficial effects of geo-distribution is positive for them.

These considerations apply in general to all cloud deployment models (ie, public, private, community and hybrid), but the impact of those weaknesses and threats varies depending on the specific internal and external environment of public organizations in different Member States and the deployment and delivery model considered.

In terms of architecture, for sensitive applications private and community clouds appear to be the solution that currently best fits the needs of public administrations since they offer the highest level of governance, control and visibility, even though when planning a private or community cloud, special regard should be given to the scale of the infrastructure. If a private cloud infrastructure does not

reach the necessary critical mass, most of the resilience and security benefits of the cloud model will not be realised.

Of particular interest seems to us the case of the community cloud model since it shows the potential to conjugate data and IT solution governance and controls with an high level of resilience, especially in the case of a distribute and federated infrastructure ([see Annex III](#)).

The public cloud option is already able to provide a very resilient service with an associated satisfactory level of data assurance and is the most cost effective. Moreover public cloud offers potentially the highest level of service availability, but due to the current regulatory complexity of intra-EU and extra-EU trans-border data transfer, its adoption should be limited to non-sensitive or non critical applications and in the context of a defined strategy for cloud adoption which should include a clear exit strategy. At the same time a number of emerging initiatives, including CSA Guidance, Control Matrix, and Consensus Assessment as well as the work of the Common Assurance Maturity Model (CAMP) (2) consortium are pushing the yardstick on providing the transparency and assurance that will allow using public cloud model in more sensitive applications.

Recommendations to governments and public bodies¹

- Governments are recommended to adopt a staged approach in integrating cloud computing into their operations because the complexity of the cloud environment that introduces a number of unknown variables for which Public administrators (PAs) will need to build new approaches to assessing and managing risks.

Public administrators (PAs) at any level should consider system interconnection and interdependencies (most of which may be unknown) especially when simultaneously moving multiple services to a cloud system(s). PAs should consider this caveat in the context of a dynamically changing environment and a currently incomplete understanding of vulnerability and attack mechanisms, and the complexity of related controls. PAs should not assume that the successful deployment of an application in a cloud environment is automatically a positive indication for proceeding with many other deployments; the security and resilience requirements of each application should be examined carefully and individually and compared to the available cloud architectures and security controls. In this perspective, the ability of backtracking from the adoption of a cloud solution should be planned before moving to the cloud.

¹ The full list of recommendation can be found in Chapter 7

- National governments should prepare, in the context of a wider EU approach, a strategy on cloud computing that takes into account the implications for security and resilience that such service delivery models will have in the context of their national economies and services to citizens over the next 10 years. The early adopters in each Member State should be seen as possible test beds, but it will be essential to have, at least at a national level, a coherent and harmonized approach to cloud computing in order to avoid: 1) the proliferation of incompatible platform and data formats (lack of service interoperability), 2) an inconsistent approach to security and resilience, including an inconsistent and inefficient approach to risk management, and 3) a lack of critical mass.
- We recommend governments to study the role that cloud computing will play in the context of protecting critical information infrastructures. It is not unrealistic to assume that cloud computing, in all its possible implementations, will serve, in the near future, a significant portion of European Union citizens, small and medium-sized enterprises and public administrations, and therefore the cloud infrastructures from which services are provided should be protected as such. In other words, a national strategy for cloud computing should aim to understand and address, among other issues, the effects of national and supra-national clouds interoperability and interdependencies, and assess the impact of possible cascade failures, evaluate the opportunity to include cloud providers in the scope of already announced reporting schemes (in particular we refer to the reporting mechanism introduced in articles 4 and 13 of the newly adopted Telco Package (3)) and to be prepared for crisis management in the event of large-scale incidents of this nature.
- We recommend national governments and European Union institutions to further investigate the concept of a European Governmental cloud as a supra national virtual space where a consistent and harmonized set of rules could be applied, both in terms of legislation and security policy and where interoperability and standardization could be fostered. Moreover such a European Union wide infrastructure could be used in the context of a pan European mutual aid and assistance plan for emergencies.

As public bodies evaluate the benefits and risks of adopting cloud computing, they should:

- Assess their risks and define their requirements (possibly using as a support those suggested in this report) in order to identify which cloud solution matches their needs. PAs should also consider human factors (eg, security and resilience awareness, resistance to new security policy models) and legal frameworks.
- Review their existing information security management policies and processes and assess how these would be addressed or supported in various cloud models.

- Define acceptable levels of service (a benchmark to evaluate parameters such as availability, response time, etc)) for their requirements. They will use the benchmark(s) to measure the performance of their services. Identify the set of controls and their degree of specificity needed to reach a minimum acceptable level of data assurance and services resilience.
- Make sure that all the essential security, resilience and legal requirements are detailed in their service level requirements and specified in their service level agreements.
- Have tools, methodologies and governance structures to, for example, assure due diligence.
- Ensure that satisfactory telecommunication connections, critical dependencies (eg, electricity), processing power and storage capacity are guaranteed and maintained.

Check the priority for the resumption of third-party communications and cloud services in the event of a disruption.

- Test the business continuity plan along the whole services supply chain.

Finally, cloud providers and independent service vendors should consider the recommendations included in this report as a possible source of information when aligning their business offers and values proposition with users' needs and requirements.

1. Introduction

Many ministries, governmental agencies and public administrations (PAs) outside the European Union, eg, in the USA, Japan², Singapore (4), and many other countries, are now approaching the cloud.

The main reasons for this choice are very well summarised in the document *State of Public Sector Cloud Computing* from the US Federal Chief Information Officer that '... cloud computing has the potential to greatly reduce waste, increase data centre efficiency and utilization rates, and lower operating costs...'. In the European Union (5), some countries such as the UK, Denmark, and the Netherlands, as well as the European Commission (6) (7), are analysing the cloud model and working on the definition of their strategies.

In May 2010, the European Commission published its *Digital Agenda for Europe* (8) which states that '... the Commission will ensure sufficient financial support to joint ICT research infrastructures and innovation clusters, develop further eInfrastructures and establish an EU strategy for cloud computing notably for government and science'.

At the same time, in the private sector, the number of companies using the cloud continues to grow rapidly, and the maturity of offers is increasing with the introduction of new services.

According to Gartner, worldwide cloud services revenue is forecast to reach \$68.3 billion in 2010, a 16.6 percent increase from 2009 revenue of \$58.6 billion. The industry is poised for strong growth through 2014, when worldwide cloud services revenue is projected to reach \$148.8 billion

Given the above policy and business context, ENISA believes it is important to provide guidance on the security and resilience factors influencing the choice for (or decision against) cloud computing solutions for public bodies and organizations. For this reason we decided to support public bodies by mean of a comparative assessment of different approaches to cloud computing.

This report represents a follow-up to the 2009 ENISA report *Cloud computing: benefits, risk and recommendations for information security*, in which a risk assessment of cloud computing business models and technologies was conducted. The result is an in-depth and independent analysis that outlines some of the information security benefits and key security risks of cloud computing. The

² In Japan, the Ministry of Internal Affairs and Communication (MIC) is building the Kasumigaseki Cloud in order to optimise operations in central governments. The Ministry of Economy, Trade and Industry (METI) has built the J-SaaS and the e-METI Idea box. There are several existing or planned cloud projects in the government and business sectors such as finance, airlines, communications, water, and a number of other projects.

report also provides a set of practical recommendations.

Both reports, *Governmental Cloud: making an informed decision* and *Cloud Computing: benefits, risks and recommendations for information security* were produced in the context of the Emerging and Future Risk Programme.

For other ENISA work in the area of resilience, see (9)

1.1. Structure of the report and how to read it

The report is structured as follows:

Chapter 2 describes the objectives of the report, the method of analysis and the target audience.

Chapter 3 presents a simple model for decision-makers, and describes the first three steps in the process, namely the identification of the security and resilience parameters (Step 2), the identification of the operational and legal parameters (Step 1) and, finally, the available architectural options for IT services (Step 3). It should be noted that the steps are not presented in their logical sequence; we first introduce our audience to Step 2 and then Step 1 of the process, as security and resilience are the main focus of this report.

Chapter 4 describes the fourth step of the model, the comparative assessment, and a general SWOT analysis for community, private and public clouds is provided.

Chapter 5 offers a demonstration on how to apply the first five steps of the simple model for decision-making by considering four sample services taken from the three scenarios considered in the report.

Chapter 6 describes the actions to be taken and the controls to be considered with respect to information security and service resilience when preparing a request for a service proposal.

Chapter 7 proposes a set of recommendations in the area of information security and service resilience for national governments and public bodies evaluating cloud options.

Finally, the following documents are included as annexes:

- Annex I: Legal analysis
- Annex II: Scenarios
- Annex III: Reservoir architecture
- Annex IV: List of threats to be used as a supportive document for performing an in-depth risk assessment

As this report is targeted at several diverse audiences, the reader should consider that:

- Essential information can be found in the executive summary and key recommendations at the beginning of the report
- Information for non-experts can be found in a story telling format in the introductory scenario.
- Detailed analysis can be found in the main body of report
- In depth analysis can be found in the annexes

1.2. An introductory scenario: Making a decision

The Minister of Communications and Technology was frowning and tapping his fingers impatiently on his polished desk when his assistant opened the door and brought in the task force. The assistant introduced them: Paulo and Hardizon from the private sector; Apik, a privacy advocate; Hitch, head of the Ministry's IT department; Luther, the Ministry's General Counsel; Fudge, the Ministry's Chief Financial Officer; Veeraswami, an independent auditor.

"I'm not happy about this situation, gentlemen," he paused and nodded toward Veeraswami, "and lady. I'm disappointed that you were unable to reach a consensus on whether I should recommend to the Prime Minister that we should move all government computing services to the cloud or not."

Hitch delicately put a 300-page report on the Minister's desk and slid it toward the Minister. "All of the task force's considerations, the pros and cons, are here."

The Minister's assistant, Ference, leaned toward Hitch and whispered. "You know the Minister will not read anything longer than two pages."

"Well, then," said Hitch. "I can sum up the entire report, Minister. I regret to say that the task force was evenly divided. You will need to make the decision."

The Minister looked at his watch. "This is really a nuisance. Okay, then, give me the key points. First of all, by how much will we cut our IT costs if we were to move our services to the cloud?"

Fudge spoke up. "A lot Minister, about 90 per cent. There is a lot of duplication across the government. Each public institution has its own IT department, staff performing the same functions. Plus their own servers. Sometimes with proprietary services, sometimes using publicly available services. If we consolidate all of our storage and services on the cloud, there are operational efficiency gains to be had. We wouldn't have to license the same software many times over for each government department. We could downsize the IT departments. We could cut our IT costs by about 30 billion euro per year (10). We could pay for the service we actually use, rather than paying for something whether we use it or not. Plus we could mark it down as an operating expense, rather than a capital expenditure, so it would make the government's budget figures look better."

"The IT department would also enjoy benefits," followed Hitch. "Using cloud as a development and testing bed, we could significantly reduce the time and cost of new service development. We wouldn't have to wait until new machines are delivered, nor would we need to produce estimates of peak capacity loads because the cloud is scalable by its nature. The cloud might just give our staff agility in preparing new services."

“Excellent” said the Minister, breaking into a satisfied smile. “So where’s the problem? Why didn’t you reach a consensus?”

Veeraswami smiled serenely. “There’s no doubt about the considerable cost savings, but it’s not simply a matter of cost savings. The benefit could be significant. There are some, shall we say, hidden costs and issues that must also be factored in...”

“Such as?” asked the Minister.

“Modifications to existing applications, disaster recovery, liability, lost of immediate control, insurance for data losses... and perhaps most importantly by using non-European cloud providers, we forego the opportunity to develop national capabilities, so there is an opportunity cost too.”

The Minister grumbled, “Opportunity costs and intangible issues are difficult to convey to voters.”

“It’s not just a question of costs,” said Apik, interrupting. “The government would lose control of all its data, of all its citizens’ data. You would have no idea where the data is. It could be anywhere in the world. It could be stored in a country where they don’t abide by the Data Protection Directive. Imagine how the press would react if they discovered that an unsafe third country was mining all of your data, your state secrets, your personal data. The scandal could sink your government.”

“Come, come,” said Hardizon. “There’s no need to be melodramatic. We’ll only use our facilities in ways that the law allows, and EU law allows some flexibility to us. It isn’t in any ones interest for us to put your data somewhere inappropriate, and there’s a limited list of countries where we have appropriate facilities anyway.”

“I quite agree, Minister,” said Luther. “It’s a contractual issue. The government could have a contract which specifies where the data could be stored. It would have to be somewhere in the European Union.”

“Sure, sure,” said Apik. “But you would have no way of knowing whether they were adhering to the contract or not.”

“That’s true,” said Veeraswami. “From our study, it’s impossible to say where the data is or where it goes.”

Hardizon snorted. “That’s not true. We can agree a service level agreement that the data would be backed up only in Europe.”

“But the point is, there’s no way of auditing the proper performance of the contract, of knowing where the data is at any given time or who access it or the security measures in place to protect it.”

“You’ll have some knowledge, but not the total insight. That’s because our systems are proprietary,” said Paulo. “My company, for example, has a record in the industry that is the envy of our competitors. There’s no way we could give away our competitive advantage.”

“That may well be,” said Hitch, “But if there were a disruption of some sort in Cloud A or we found their services unsatisfactory, and we wanted to switch to Cloud B, it would be virtually impossible because their systems are proprietary – which is another way of saying they are not interoperable. That’s my main concern. We would be locked into a particular supplier.”

“Yeah,” said Apik. “How would that go down with the voters?”

“Do you mean to say,” asked the Minister, “that if we chose one cloud provider, and we weren’t satisfied with their service or if its service was disrupted, we would not be able to take our business elsewhere?”

“Not really, Minister. Most cloud providers have unique application programming interfaces or APIs as they are known, which basically means that applications are not easily portable across clouds.”³

“Hmm, I don’t like the sound of that,” mused the Minister, rubbing his chin. “Perhaps we need to bring in some new regulation, to force providers to standardise these, what did you call them, APIs?”

“The APIs may not be the same, but as long as they are open it’s feasible to insert simple libraries to get portability. If you try to make us standardize too soon, you’ll kill innovation by new entrants, including our upcoming EU competitors. Besides, your current bespoke systems have far more serious restrictions, which is why you have so many problems. But let’s be honest,” exhorted Paulo, attempting to take a new tack in the discussion. “It’s a question of cost and availability. Our services and your data would be available more or less all the time. We offer 99.5 per cent availability whereas I’ve seen some numbers that suggest government services are nowhere near that good!”

“Right, so have I,” said Hardizon, not to be outdone by his competitor. “Governments often undertake big IT projects that run into trouble. Schedule and cost often overruns. If you were to switch to our cloud facilities, you could shift that risk to our shoulders.” Hardizon straightened, pulled back his shoulders as if to illustrate the point he was making.

³ “Currently no standard is available, nor is there any concerted effort by CSPs to develop a ubiquitous and consistent API across clouds – and that makes porting of an application across PaaS clouds a monumental task.” Mather et al., p. 57.

“My company provides not just a back-up, but several back-ups,” said Paulo, “and in different countries, for example, in Europe and the United States. Governments don’t do that. If there were a massive power failure, like the one that started in Germany a few years ago⁴, and cascaded right the way down to Spain and Portugal, we could deal with that. So we don’t put all your eggs in one basket, we put them in several baskets...”

“If I may intervene, Minister,” said Luther. “While having geographically disparate locations will help ensure resilience in the event of a disruption in power supplies like the one in Germany, they also raise some jurisdictional problems. We may not have Safe harbour agreements in place with some of those third countries. Thus, our government data would be subject to the laws and regulations of other countries.”

“Exactly,” said Apik. “Cloud providers operate across or in many jurisdictions. There’s nothing to stop them shopping around for the most favourable regime.”

Luther didn’t take kindly to being interrupted, but as Apik was supporting him, he smiled wanly at the privacy advocate and went on to his next point: “And there might be nothing to prevent law enforcement authorities getting access to the data....”

“Or the cloud providers themselves might mine the data,” said Apik. “Imagine what a treasure trove all that government data would be.”

“That’s ridiculous,” said Paulo. “It’s even insulting.”

"My apologies," responded Apik, "I realise that your reputation is extremely valuable and there's no benefit to you in ruining it by mining citizens' data, but an unscrupulous provider might consider that."

"Perhaps, but there's the same risk from the companies you pay to build, write software for and manage your existing data centres. Indeed more of a risk as you're undermining their business model by even considering real cloud systems. I agree you need to consider the risk, but why would we destroy a valuable business by behaving so stupidly?"

"But there's another concern," parried Apik. “Isn’t it true that, as a result of the USA PATRIOT Act, the Canadian government instructed departments not to use computers operating within US borders,

⁴ Graham, Dave, and Allan Hall, “Power cuts in Germany spark wave of blackouts across Europe”, The Scotsman, 6 Nov 2006. <http://news.scotsman.com/international.cfm?id=1640182006>

because it had concerns about the confidentiality and privacy of Canadian data stored on those computers?”⁵

Ference, the Minister’s assistant, seeing that task force members were getting hot under the collar again, hoped to defuse things a bit: “Perhaps you should brief the Minister on how resilient a move to the cloud might be.”

“Good idea,” said Hardizon. “Minister, resilience is not only a matter of having widely separated data centres. The fact is we have some of the top security experts in the world working to ensure that only authorised personnel have access to our sites and systems. It’s virtually impossible for attackers to penetrate our security, both in physical and cyberspace.”

“Except that it has happened...” said Hitch. “In any event, we have to be concerned with the security of data from the time it is generated and while it is in transit to the cloud and having access to it 24 hours a day, seven days a week, 52 weeks a year. No security is perfect. By centralising storage and services, the risk is that you simply provide a bigger target for attackers.”

“I agree no security is perfect, but government systems are attacked and penetrated even more often. We may make a bigger target, but never forget that we can also then build stronger and deeper defences than individual IT departments can ever hope for, as ENISA pointed out last year, size gives us major security advantages.” said Paulo.

“Maybe, but cloud providers must contend with criminal elements, rogue customers and insider threats just like any organisation,” said Hitch.

“True,” said Paulo, “but we also scrutinise our prospective employees more rigorously than the government. And, because of the way our systems work, far fewer people are involved in managing our systems than have access to the data in systems you manage for yourselves.

“But who scrutinises you?” asked the Minister.

Apik, relishing the moment, jumped in again. “A good question, Minister. There’s a distinct lack of transparency about cloud providers’ security measures. They expect clients to trust them with their valuable, and sometimes critical data, yet no one knows what measures they take to protect that data. A lack of transparency is a recipe for a lack of trust, at least, so it seems to me.”

⁵ Mather, Tim, Subra Kumaraswamy and Shahed Latif, *Cloud Security and Privacy*, O’Reilly Media, Sebastopol, CA, 2009, p. 33.

“Well, what about it, Messrs Harizon and Paulo? How do you respond to that point?”

“It’s quite simple, Minister,” said Paulo, “if you’ll forgive me for saying so. We don’t want attackers to know what security measures we have in place. We don’t say what security measures we take in order to protect you and our other clients.”⁶

“Independent, third party audits are the answer,” said Veeraswami, the independent, third-party auditor.

“Audits are important, I agree,” said Luther, “but service level agreements (SLAs), rock solid contracts, are what we need to get that right.”

“They are,” agreed Apik. “But commodity cloud providers usually only offer a vanilla-flavour contract. You take it or leave it. Opportunities of negotiating individual clauses are pretty minimal.”

The Minister grumbled again. “If they want our business, then we negotiate a contract that is satisfactory for us, that appears satisfactory to the public... and we can always back that up with regulation and standards. If you want to operate in our country, then you meet our standards and you comply with our regulations.”

Hardizon nodded his head. “Of course, Minister, of course.”

“I don’t usually agree with my friend Hardizon, Minister, but on this, we are as one. However, I would like to remind the Minister of his government’s commitment is to better regulation, to less regulation and to freeing enterprise.”

“Minister, you do not have to make a recommendation to the Prime Minister to migrate all government data and services to a cloud – or nothing at all. You could recommend a phased approach,” suggested Hitch.

“Do continue.”

“What I mean is that we could migrate some data and services, but not everything all at once. That would be very risky indeed. A better way would be to migrate some data and services. Non-essential stuff, I mean, non-critical data. Say tourism and public works. By gaining experience through such a

⁶ “Until now, CSPs have been reluctant to share information pertaining to platform security using the argument that such security information could provide an advantage for hackers. However, enterprise customers should demand transparency from CSPs and seek information necessary to perform risk assessment and ongoing security management.” Mather et al., p. 56.

first phase, we could then make a determination whether we should proceed to the next phase, where we might migrate more sensitive data, for example, social services and health-care data.”

“Good, I like it.” He looked around the room. “Consensus? Yes? Done.”

The Minister was about to wave the task force away – he had a heavy schedule – but before he could, Hitch spoke up: “Minister, I think you’ve achieved a good result here today, but there are just one or two little things that we should be clear about.”

The Minister, who had been pleased with his banging heads together, frowned again: “Yes, what may they be?”

“Well, I think that industry should assure us that certain security mechanism are in place, such as encryption, digital signatures, hashing, etc., in order to achieve a satisfactory level of confidentiality, integrity, availability and non-repudiation.”

The Minister was unsure what all this meant, but it sounded reasonable.

“I agree, Minister,” said Luther, the General Counsel. “We need a contract that specifies both the security of data and resilience of service, but we also need an agreement that clearly establishes which legal jurisdiction will apply and under which circumstances we can access which type of data.”

“Surely, our jurisdiction will apply,” said the Minister.

“Yes, of course,” said Luther. “But it’s not quite so simple. If we agree to the cloud provider’s backing up our data in Iceland or Canada or wherever, then we should have an enforceable agreement that specifies the types of data and which services can be transferred outside our country. We should establish a level of assurance for each category of data and service. In concrete terms, that would mean that sensitive data could be transferred outside our country only if certain conditions are met.”

The Minister scratched his head. “Yes, that sounds correct.”

Luther carried on. “If we agree to the transfer of our citizens’ data to a third country, then that would imply that a strong relationship of trust exists between our governments. We might need more than that. We might need an international treaty that provides for full control over data location and jurisdiction. We would need a bilateral or multilateral agreement between our government, the State where the data is backed up and the government of the country where the cloud provider has its primary legal base. That agreement should contain provisions for the regulation of subpoena and e-discovery.”

Apik added. “Minister this is not just concerning technical or legal factor, but there must be also a consideration of the human factors as well!”

“Hmm,” said the Minister. “It seems there a lot of factors to be considered...”

Hitch interjected. “That’s true, Minister, there are. Fortunately, just this morning, I received a report from ENISA....”

“ENISA?”

“Yes, Minister, you know, the European Network and Information Security Agency,” explained his assistant.

“Yes, of course. And...?”

“And,” said Hitch. “It addresses these very issues.”

“Excellent,” said the Minister. “Review it carefully and then let’s get on with it.”

The task force nodded their heads and collectively chimed, “Yes, Minister.”

2. Objectives and analysis

This report has a twofold objective: 1) to guide public bodies in the definition of their information security and resilience profiles, and in the evaluation of the NIS strengths, weaknesses, opportunities and threats of cloud computing service delivery models, and 2) to indirectly support Member States in the definition of their national cloud strategies with respect to information security and service resilience.

In this report public bodies will find ideas and tools which are meant to facilitate their answer the following questions:

- What is the value of a governmental cloud solution in terms of resilience and reliability?
- Can the cloud service delivery model offer at least the same level of security and resilience that public organizations (local and regional public authorities and healthcare authorities) currently have?
- Which deployment model (private, public, hybrid, or community), if any, is best suited for a specific public administration?
- Which is the best match (if any) between service models (IaaS, PaaS, SaaS) and services (eg, online collection of medical files, online tax payment, and other less critical services, such as, back-end, HR, payroll, and e-learning)?
- How can public administrations ensure effective controls over security and resilience? What forms of audits, SLAs, financial penalties or incentives, etc, will work best in providing adequate assurance?
- Who should be liable for what aspects of policies related to security and resilience in a typical government cloud deployment?
- Which rules and regulations have to be observed? Which duties and obligations have to be fulfilled?
- Is it realistic for governments to plan and deploy governmental clouds using currently available technology? What are the main open issues that have to be addressed in terms of security and resilience before government clouds can be deployed operationally?

2.1. Target audience

The target audiences of this report are:

- CEOs, CTOs, CISOs and other ICT staff members in EU Member States evaluating the information security, resilience and reliability of a governmental cloud;
- public bodies in the EU (local and regional public administrations, agencies, local healthcare authorities, etc) evaluating the costs and benefits for a public administration considering migrating to a cloud;
- European Union policymakers deciding on appropriate policy and economic incentives, legislative measures, awareness-raising initiatives, etc, vis-à-vis cloud-computing technologies for governments and public administrations;
- cloud providers and cloud VAS (value-added services – including security) providers trying to achieve an early understanding of the needs and requirements of central governments, public administrations and citizens.

2.2. Analysis method

This report contains three fictitious use-cases or scenarios that describe:

- A local healthcare authority implementing electronic healthcare records and other e-services. This scenario aims to capture the requirements of services dealing with more sensitive data and stricter needs for resilience.
- A local public administration rolling out new services for citizens and evolving existing ones, while consolidating its internal IT infrastructures and platforms.
- A central government planning the creation of a governmental cloud as a subsidised platform for stimulating business innovation.

The data to build and refine the scenarios were collected from:

- five local healthcare authorities (Italy);
- a national healthcare authority (The Netherlands);
- a local public administration (Spain);
- IPA – Information Technologies Promotion Agency, Japan;

- ELANET (CEMR) – European Network for eGovernment and Information Society (supported by the Council of European Municipalities and Regions);
- a data protection authority (Greece);
- a questionnaire distributed to qualified sources in public administrations;
- an open online consultation.

The definition phase for the scenarios (please note that the three scenarios can be found in [Annex II](#)) was followed by the analysis which included the following steps:

- definition of a simple model for decision-makers;
- identification of the business, operational and legal requirements and constraints;
- alignment of information security and resilience requirements with business, operational and legal requirements;
- description of the IT architectural options available;
- analysis of the strengths, weakness, opportunities and threats of cloud delivery models based on security and resilience parameters;
- identification of specific security, resilience and compliance requirements for the four (4) sample services described in the three (3) scenarios;
- scenario-specific comparative assessments (based on SWOT) of cloud deployment models;
- definition of recommendations, which include a set of controls or questions that should be used both in the design phase of a service and for monitoring compliance with the service level agreement.

3. Model for decision-makers

This chapter proposes a simple model to support government decision-makers in their approach to a cloud computing service delivery model. The idea is to guide public administrations:

- the identification and collection of their business, security and legal requirements;
- the definition of their service level specifications and service level agreements;
- the identification of the solution that best addresses their needs;
- preparing a proposal for a request-for-service and establishing their mitigation plan.

In the description of the simple decision-makers model, we stress the importance of the requirements collection phase which is a key factor for taking a final informed decision.

In general terms we can state that the implementation of new governmental services and the evolution of existing ones are conditioned by:

Internal environment

- mission and business requirements
- financial constraints
- status quo

External factors

- available technology options
- expectations of users (citizens, private companies, patients, etc) and public opinion
- existing laws and regulations at both the national and European Union levels

The variables mentioned should be considered when aligning an information security strategy with the business goals of a public institution. They should be the main drivers in the definition of a risk profile for a public organization and therefore the main drivers in the determination of the level of maturity in information security and resilience the organisation demands for service provisioning.

It is important to note that the security and resilience objectives and needs of an organization should be identified clearly (eg, total services availability = 99.9 %) on the base of quantifiable metrics (eg, total availability each month), defined through an SLAs and monitored on a constant basis.

A fundamental part of the decision-makers process is to perform a comparative risk assessment (at least a SWOT analysis) in order to achieve a sound and informed decision that takes information security and resilience into account from the planning phase of a project.

Organizations in their approach to service provisioning, will eventually use a decision-making model similar to the one described in the figure below.




FIGURE 1: DECISION PROCESS

Figure 1 shows how operational, legal and information security requirements, as well as budget and time constraints, drive the identification of the architectural solution that best suits the needs of a public administration, agency or healthcare authority (Steps 1, 2 and 3).

By architectural solution, in this report, we mean either: 1) public cloud, 2) private cloud, or 3) community cloud. Each solution can support one of the delivery models: IaaS, PaaS, or SaaS. The hybrid cloud architectural solution was not considered as it represents, in our view, a second step in the cloud approach, since it combines the use of different cloud models. The hybrid cloud architecture as a solution shall thus be neither qualified nor disqualified. Yet we see the distinction between public, private, and community clouds as the key criteria in identifying and deriving lower limits regarding the differing security aspects. In the second phase of architectural design, a hybrid approach might be taken while respecting the outcome of the foregoing analysis.


The most appropriate model, as far as security and resilience are concerned, is identified by performing a comparative assessment based on specific security and resilience criteria which are derived, directly or indirectly, from the essential requirements of a service (Step 4).

Assuming that the risk assessment in step 4 confirms that a cloud solution can be considered and once the architectural solution has been identified, the next steps are: the identification of the specific threats and weaknesses of the selected IT service model (step 5), and the preparation of a request for a proposal in order to select a business partner, service and/or product provider (Step 6). A safe and cautious approach to this selection step would be the identification of a control checklist that should be used to compare and assess the proposed services and solutions.

3.1. Security and resilience parameters

In this section we offer some possible variables to consider for understanding the requirements of a given service.

As already mentioned earlier in this report, we first present Step 2 as information security and resilience are focal points of our analysis.



In paragraph 3.1.2, decision-makers will find a set of security and resilience variables which are likely to be considered when defining their requirements.

End-to-end secure and resilient service

Resilience is the ability of a system (network, service, infrastructure, etc) to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation.

Security is the ability to *protect information and information systems from unauthorized access, use, disclosure, disruption, modification or destruction and to respond and recover in case of a fault or incident* (12) .

In this report we assume that data security and service resilience are considered when defining the acceptable level of service for each organisation. Hence, a service can be considered end-to-end secure and resilient when it performs as described in the service level specification (SLS).

In the context of this study that means that a service should provide:

- a level of data confidentiality, integrity and availability according to specified requirements;
- a level of service availability and reliability according to specified requirements;
- compliance with the applicable law.

The lack of one or more of these requirements will render the service unsuitable to meet the service level requirements and to satisfy the users' expectations.

When considering the technical aspects of end-to-end security and resilience it is necessary to take into account the organization of the architectural components of the entire supply chain: clients,

network (eg, LAN, WAN), data centre, public services, systems management, and security services, as well as the solutions adopted at the infrastructure, platform, application and data levels.

In other words, each organization should consider how the overall service delivery supply chain might be built out of a mix of internal infrastructure and services provided by external suppliers. Therefore it is necessary to pay attention to all the components and their interconnections along the supply chain, such as communications between user client and application, between application and database, between networks (LAN to LAN, LAN to WAN, etc), as well as hardware components, chips etc.

Legal compliance is a requirement that is as important as the technical and organizational requirements for security and resilience. In fact, a lack of it could give rise to legal disputes with citizens, disputes between local or regional administrations and governments, conflicts and disputes with national regulatory authorities (NRAs) for the protection of telecommunications and data, and conflicts with law enforcement agencies. Ultimately it could prevent public administrations from offering their services.

Security and resilience selection parameters

In this paragraph we suggest, based on the ENISA report *Metrics for resilience*⁷, a set of security and resilience parameters that should be considered when evaluating possible deployment and delivery models for IT services.

This section presents some considerations that government agencies and public administration organizations evaluating cloud services should take into account when defining their service requirements.

These qualitative and quantitative parameters we propose are mainly based on the ENISA report *Metrics for resilience*. We grouped set of parameters into four categories which describe most of the requirements that should be considered when planning for an end-to-end secure and resilient service. The four categories are:

1. preparedness: including the parameters and criteria used to understand the level of preparedness of an organization to efficiently maintain an acceptable level of service while protecting the confidentiality and integrity of data both during daily operations and in case of an incident.

⁷ <http://www.enisa.europa.eu/act/res/other-areas/metrics>

2. service delivery: including the criteria used to assess the capability of the systems to offer a level of service in the line with the requirements expressed in the service level agreement;
3. response and recovery: including the criteria to measure the capacity of the system to react in cases of incidents or faults;
4. legal and regulatory compliance: including the criteria for assessing the level of legal compliance.

Most of the suggested parameters are, or can be turned into, metrics and parameters to monitor the sound execution of the operations in the cloud as well as to understand whether SLAs are fulfilled.

It should be noted that the level of security governance in each organization will have an impact on the way the controls underlying the suggested parameters can be implemented and therefore it will strongly influence the security and resilience of the service itself. A higher level of governance implies a higher degree of control over the parameters suggested below.

In general terms we can say that the SaaS delivery model is clearly the solution that offers the customer less direct control over the security and resilience parameters, but placing more control and responsibility in the hands of the CSP, while IaaS is the one that guarantees more direct control but also leaves the customer fully responsible for the implementation of technical and procedural security and resilience measures (for more details, please check *Division of Responsibilities* in ENISA's cloud report 2009).

The following sections present some selected parameters that must be understood by the organization in developing requirements for a cloud service migration. These requirements will need to be addressed by all parts of the end-to-end solution including the organization itself, the CSP, as well as network and telecom providers involved in delivering the service.

A. Preparedness

These parameters describe the level of preparedness required from a system in order to continue in the face of faults and incidents. Preparedness parameters include all the actions and measures taken to prevent an incident from happening or to minimise its impact.

A1. Risk analysis and assessment

In this area we suggest some metrics which cover the adequacy of risk analysis and assessment practices.

- Risk analysis and assessment frequency
- Vulnerability assessment coverage

- Vulnerability assessment frequency
- Security testing (eg, penetration testing) frequency

As a general consideration, we can say that private clouds should offer a higher degree of customization of risk analyses and assessment practises, so a public administration may more easily define the frequency and coverage of tests and analyses according to their specific requirements.

A2. Prevention and detection

In this area we include parameters that cover the extent to which a public body requires the service to be monitored in real-time as well as whether the resource capping mechanics in place are suitable for guaranteeing a controllable use of resources.

In the category of real-time security monitoring we include network performance and integrity, operating system performance and integrity, baseline comparison and unauthorized attempts at access, as well as security monitoring (collection, analysis and triage of security events generated from firewalls, intrusion prevention and detection systems (IPS-IDS), proxies, antivirus, application firewalls, and any other network and security components).

- Reporting frequency
- Resource capping mechanisms in place

A3. Patch management

We propose the use of the following measures to verify the effectiveness of patch management.

- Mean time to patch
- Patch management coverage

A4. Access control and accountability

The parameters included here focus on the collection of evidence (logs) to prove the soundness of the processes and mechanisms to control authentication, authorization and the accountability of users that are in place.

- Level of availability of logs
- Visibility of logs

A5. Supply chain

The more control is maintained over the service delivery supply chain, the better the security and resilience that will be achieved. Bearing this in mind, we suggest an ‘auditability’ parameter as a way to understand the possible level of transparency and control of the supply chain; more specifically:

- the type of audit that can be performed (internal, third-party independent, self assessment, etc);
- the scope of the audit (which link(s) in the chain can be audited), the methodology used, etc.

B. Service delivery

This set of parameters is included to evaluate the requirements for the service architecture to maintain an acceptable level of service in the face of unexpected events, random faults, performance degradations, or targeted attacks. In public or community clouds, some of these problems may arise because of the users who share the cloud. Hence, these problems may be more critical whenever an organization cannot control the platform or infrastructure. In the case of SaaS, the service delivery parameters fully depend on the internal software architecture that is controlled by the provider. More control over these parameters is possible in the case of PaaS and IaaS, but it should be clear that significant expertise is required to properly use those control parameters.

B1. Availability and reliability

In IaaS and PaaS, the overall system may be designed and deployed to achieve better values in the tolerance of faults and malicious attacks. At the same time the availability of a cloud service is often dependent on the network used to access it; therefore some of the measure will apply to the ISPs as well. Parameters that should be used to measure the availability and reliability of services are:

- mean time to failure
- mean time between failures
- total monthly (or daily) availability
- incident rate
- tolerance to malicious attacks
- redundancy
- replication.

Other parameters that could be used, especially with regards to data integrity, could for instance be:

- percentage of systems with automatic virus definition updates and automatic virus scanning
- the percentage of systems that perform password policy verification
- length of encryption keys

- use of integrity and non-repudiation controls, eg, checksum functions, hash functions, fingerprints, and cryptographic hash functions.

Moreover, the response time to the user is influenced by the quality of the network connections between the user and the cloud as well as within the cloud. Even if the cloud is properly designed and deployed, a low performance connection to the cloud may reduce the final performance available to the users. Important parameters include:

- throughput (bandwidth)
- latency (average round trip time)
- packet loss
- jitter (packet delay variation).

As availability and reliability are two of the core parameters in the evaluation of service resilience, it is of paramount importance that the measures and metrics used are consistent and that the object of measurement is the same. In this report we always refer to the availability and reliability of the service for the end-users.

B2. Scalability and elasticity

Capacity management is the process responsible for ensuring that the capacity of IT services and IT infrastructure is able to deliver the agreed service level target in a cost effective and timely manner. Capacity management considers all resources required to deliver the IT service, and plans for short, medium and long-term business requirements (13).

The ability to manage changes in the resources demanded (storage, CPU time, memory, web service requests and virtual machine instances, etc) and to scale up and eventually down is a crucial factor for a service to be effective and efficient. Therefore, when considering capacity and demand management, two important criteria should be taken into account, namely:

- capacity fluctuations: unpredictability in traffic load variations, link capacity fluctuations, node failures or other types of intentional misbehaviour that may lead the network into overload conditions;
- long-term scalability (up and down): the ability of the system to increase or decrease its capacity to provide the requested resources in a timeframe suitable for meeting the service level requirements.

The following parameters relate to the ability of an application to fully exploit cloud resources to react to changes in the load imposed on an application. Important parameters include:

- Load tolerance: this can be calculated as a ratio of the maximum load a system can handle compared to the normal expected load, eg, the percentage of the normal load by which a system can temporary upscale (bandwidth, processing power, etc). The unit is relative; it indicates the allowed variation in the system's load without impacting the performance as a whole. It should be noted also that the load tolerance for two service providers of the same size, with the same ratio of maximum load to normal load, can result in different levels of tolerance if one serves thousands of small customers and the other a handful of very large ones.
- Traffic tolerance (including anti-DoS/DDoS provisions, eg, filtering, firewalling, rerouting, shut-off of clients producing excessive traffic, etc): the ability of a system to tolerate unpredictable offered load without a significant drop in carried load (including congestion collapse), as well as to isolate the effects from cross-traffic, other flows and other nodes. The traffic can either be unexpected but legitimate, such as from a flash crowd, or malicious, such as a DDoS attack.
- Load variability of the services (the difference between peak and mean demand).

Other parameters to be considered are those related to the provisioning of the services and hardware components, for instance:

- time to procure new hardware components
- time to service fulfilment (service deployment and provisioning).

It should be noted that these parameters are of particular relevance when comparing the cloud computing model to a non-cloud IT solution.

C. Response and recovery

These parameters relate to the capability of an organization to adequately respond to and effectively recover from incidents. The required RTO (recovery time objective – how long) and RPO (recovery point objective) should be identified. In this stage an organization would need to consider when the resilience plan needs to be invoked, who has to be informed and the channels to be used. The organization has to make sure it has the capability (analyst team) to promptly understand the root cause of the incident and its impact (understand what happened). The organization has to make sure the incident is tracked during its lifecycle (lesson identified), and that the event is adequately communicated to the external world. Finally, the response and recovery plans need to be tested.

In order to measure the effectiveness and efficiency of the response and recovery strategy in place, the following metrics can be used:

- mean time to incident discovery (delay): the time that it takes from the time an incident occurs to when the incident is discovered;

- time to invoke: the time it takes to realise that the recovery-response phase should be invoked (mean time to invoke);
- time to repair (mean time to repair): the time it takes to bring the service back to an acceptable level;
- mean time to incident recovery.

It should be noted that there is a relationship between some parameters, such as the time to recovery and the frequency and architecture of the backup systems used.

D. Legal and regulatory compliance

These parameters generally relate to requirements for the cloud service provider's SLAs and contractual provisions (eg, system execution states).

D1. Forensics

- Requirements for the extraction of evidence contained in cloud services (eg, e-discovery, data retention)


D2. Data retention and track back

- Minimum and maximum data retention periods
- Minimum and maximum log retention periods
- Data storage modality
- Log storage modality
- Time to transfer back

D3. Confidentiality

The degree of confidentiality required will depend on national legislation, eg, social healthcare, social security data, and tax records. Possible implications of this requirement are the encryption solutions required of the provider, eg, key length.

3.2. Business and operational variables



In this section we suggest a number of criteria likely to be considered when defining the business, operational, legal and regulatory requirements for public organizations.

Some of the criteria and parameters suggested are explained and described while others are only enounced.

Types of data

One of the most important criteria to take into account, when considering the implementation of a cloud solution, is the type of data that will be processed and stored by the service provider.

According to what is mentioned in the three scenarios considered in this report, the types of data to be considered are:

- **Personal data:** names, addresses, occupations, contacts details, etc.
- **Sensitive data:** intellectual property, business confidential and financial transaction data, and health records.
- **Classified information.**
- **Aggregated data:** information that can be inferred from data that has been aggregated, by allowing the inference of information or simply co-locating data that should not be related because of its sensitivity. Note that aggregations of data are considered under the EU Data Protection Directive as the perusal of data.

User profile

The analysis of the user profile represents a very important criterion for consideration, especially in community and private clouds (for instance, see scenario ‘Gov cloud as a business incubator’). Depending on the type of potential users and their geographical spread, the other business requirements will be identified and the specification designed.

In principle, three important characteristics to consider are:

- user communities (citizens, companies, and other PAs)
- geographic distribution
- level of ICT literacy and security awareness.

Scalability and capacity management

Capacity management is the process responsible for ensuring that the capacity of IT services and IT infrastructure is able to deliver the agreed service level target in a cost effective and timely manner. Capacity management considers all resources required to deliver the IT service, and plans for short, medium and long-term business requirements (13).

The ability to manage changes in the resources demanded (storage, CPU time, memory, web service requests and virtual machine instances, etc) and to scale up and eventually down is a crucial factor for a service to be effective and efficient. Therefore, when considering capacity and demand management, two important criteria should be taken into account, namely:

- capacity fluctuations: unpredictability in traffic load variations, link capacity fluctuations, node failures or other types of intentional misbehaviour that may lead the network into overload conditions;
- long-term scalability (up and down): the ability of the system to increase or decrease its capacity to provide the requested resources in a timeframe suitable for meeting the service level requirements.

Interface interoperability

Interoperability is the ability of two or more systems or components to exchange information and to use the information that has been exchanged (14).

In this study we consider the following attributes to describe the interoperability needs of services:

- interface interoperability / interface complexity

- data format exchange capabilities
- means of transfer / exchange
- identity system
- policy interoperability.

Collaboration

Collaboration between systems, platform and services needs to take into consideration:

the geographical dispersion of the entities (organizations, infrastructures)

the other requirements of the service

the level of heterogeneity in the ICT systems involved.

Cost and budget

Even though budget and financial implications are outside of the scope of this report, we would like to highlight the basic factors that are normally considered by CEOs, CTOs and CISOs when evaluating ICT investments, as a sound use of the available budget has an impact on the amount of resources that can be dedicated to information security.

For this purpose, the three most important variables to consider are:

- operational cost
- capital expenditure
- cost of migration.

Ownership

- Government owned and provided
- Government owned, operated by a third party
- Government sponsored
- Third party provided, referred to by the government
- Partnership
- Code of connection or statement of compliance

3.3. Legal and regulatory framework

General legal considerations

As the rule of law applies to governmental actors in all Member States, they are directly bound by their respective constitutions. This is in stark contrast to private actors who have full private autonomy ('liberalism') unless there are laws that constrain their actions. Sometimes this is not very apparent, mostly because many (sub-constitutional) laws apply to governmental actors and assure compliance with constitutional requirements. Sometimes these laws even apply similarly to both the government and the private sector. So, in most cases, discussing this sub-constitutional law will be wholly sufficient for the purposes of this analysis.

Governmental sovereignty and control over the information/data: law enforcement access, confidentiality and intellectual property issues

For governments and PAs in general, one of the main legal issues is sovereignty and control over the data that is being handled. A governmental body that is entitled to handle data retains responsibility for its proper handling and should ensure that its obligations to protect the data extend by contract to its third party providers. Where cloud infrastructure hosting extends beyond the local legal jurisdiction, the public body must consider the implications and related safeguards offered by their provider(s). If governmental data is being handled abroad by private parties in foreign jurisdiction, this creates the risk that foreign courts subpoena the private entity and thus reach into the government's data. Additionally, this may mean potential breaches of confidentiality and intellectual property laws related to the information, data, know-how, copyright or patent material they migrate to the cloud.⁸ These issues apply equally to all forms of outsourcing, including any current outsourcing arrangements as well as public, private and community cloud provision. A government body therefore should ensure that its outsourcing providers impose adequate security measures, and that procedures and mechanisms are in place so that only relevant data would ever be surrendered in response to legitimate demands by the judicial authorities. This includes checking whether the evidence is rightfully requested (by subpoena or during discovery).⁹

⁸ More on confidentiality and intellectual property issues can be found in ENISA (2009) *Cloud Computing Risk Assessment*, pp 97 et seq.

⁹ Cf Article 29 Group Working Document 1/2009 on pre-trial discovery for cross-border civil litigation, adopted on 11 February 2009, WP 158; available at: http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2009/wp158_en.pdf

Government procurement

Because private third parties will often be contracted to provide cloud services, the extensive EU regulations on public procurement will have to be observed.¹⁰ In this regard, there will not be any significant differences to procurement in other areas of governance, so that governments and PAs will be able to apply their existing knowledge and experience with the applicable laws and regulations. On the other hand, CSPs need to pre-qualify as suppliers according to EU regulations on public procurement and thus deal with all the regulations concerning government procurement.

Data protection and data security

General European data protection and data security issues related to cloud computing have been singled out already in: (i) recent European Commission Communications; (ii) documents adopted by the Article 29 Data Protection Working Party; and (iii) ENISA report: Cloud Computing Risk Assessment (15).

Here we will try to summarize the most relevant for the present analysis.

- *Restriction on applicability of the Directive 95/46/EC (Article 13(1)):*
- Pursuant to Article 13(1) of Directive 95/46/EC, Member States may restrict the application of certain provisions of Directive 95/46/EC for matters of national and public security or the prosecution and prevention of crime.¹¹ Thus, depending on local law in a Member State, in certain circumstances some data that municipalities handle may not be subject to all of the regulations under Directive 95/46/EC.
- *Data Controller – Data Processor (Directive 95/46/EC, Articles 2(d) and (e)):*

It is necessary to identify the controller, the processor, and their interactions in order to determine ‘who is responsible for compliance with data protection rules, how data subjects can exercise their rights, which is the applicable national law and how effective Data Protection Authorities can operate’ (16)¹². Directive 95/46/EC clearly distinguishes between

¹⁰ See: http://ec.europa.eu/internal_market/publicprocurement/legislation_en.htm.

¹¹ Restrictions are permitted regarding the obligations and rights provided for in Articles 6(1) (principles relating to data quality), 10 and 11(1) (information to be given to the data subject), 12 (right of access) and 21 (publicizing of processing operations).

¹² Article 29 Data Protection Working Party: Opinion 1/2010 on the Concepts of Controller and Processor’.

controller and processor. The *controller* is the individual or entity that determines the purposes of and means for processing of personal data. The *processor* is the individual or entity that processes personal data on behalf of the controller. However, applying such a definition to the cloud computing environment is quite challenging. At first glance, one might conclude that the PA/GOV is the controller and the CSP the processor.¹³ Nevertheless, CSPs often determine the means and sometimes also the purposes of the processing – thus falling within the definition of controller (17). To address this issue and provide some guidance on Article 29, the Data Protection Working Party issued an opinion on 16 February 2010 in which it adopted a viewpoint on interpreting such definitions in complex environments (16). However, the opinion did not shed much light on the specifics of the cloud computing environment, for which the roles of controller and processor still need to be determined on a case-by-case basis and in relation to the nature of the cloud services¹⁴ (18)

- *Prior checking (Directive 95/46/EC, Article 20)*

Pursuant to Article 20 and depending on national law, prior checking may be necessary for the processing. This depends on the type of service and types of data being processed.

- *Appropriate technical and organizational measures (Article 17): data integrity, identity management, and access control*

Data integrity and availability are essential elements in the provision of cloud computing services. According to Directive 95/46/EC, the controller and its processors must implement technical and organizational measures to protect personal data against accidental or unlawful destruction or accidental loss, alteration, unauthorized disclosure or access; having regard to the state of the art and the cost of their implementation, such measures must ensure a level of security appropriate to the risks represented by the processing and the nature of the data to be protected (article 17). The problem is that the concept of appropriate has been interpreted in different ways throughout EU Member States. Thus, although CSPs quite often implement widely recognized technical standards (e.g., ISO 27001) to secure customer data, these may not match perfectly to national requirements for appropriate measures. Further consistency and harmonization across the EU is required. In addition, the high level of data security

¹³ ENISA (2009) *Cloud Computing Risk Assessment*, pp 101 et seq.

¹⁴ European Data Protection Supervisor, Peter Hustinx, confirmed this approach in his speech on 'Data Protection and Cloud Computing under EU Law' on 13 April 2010, where he called for further guidance from the Working Party on the matter. Cloud Computing is on the Working Party Agenda for 2010 and 2011.

requested of a CSP in an e-health scenario is worth noting with special regard to identity management and access control.

- *Data breach and security incident notification (not mandatory, yet)*

The EU's revised framework for electronic communications clarifies the responsibilities of network operators and service providers, including their obligation to notify breaches of personal data security (articles 4 and 13). The recently launched review of the general data protection framework will include a possible extension of the obligation to notify data security breaches (19) (8). If European data protection regulations go in this direction, it will be necessary that they clearly identify the degree of a breach of data security that should be notified, to whom it must be notified (CSP client, competent data protection authority, data subjects), and the relevant modalities. An indeterminate obligation to notify any (including minor or irrelevant) breaches of data security may severely penalize CSPs and unnecessarily alarm governments, PAs and citizens in general.

- *Data transfer to countries outside the EEA (Articles 25-26)*

Cloud models entail that customer information and data may involve the transfer of data by the CSP from one data-centre in the EEA to another that can be located anywhere in the world. However, Directive 95/46/EC prohibits transfers of personal data from the EEA to countries which do not ensure an adequate level of protection within the meaning of article 25(2) – unless the data subject has previously given unambiguous consent to the proposed transfer or other procedures are in place in accordance with article 26 (eg, 'Model Contracts for the transfer of personal data to third countries', 'Safe Harbor Principles' (where the data is being transferred to the United States), or 'Binding Corporate Rules')¹⁵. There are challenges with each of these ways to legitimize a transfer, however: basing it on the consent of the data subject exposes the transfer to the uncertainties of possible withdrawals of that consent; the Safe Harbor Principles, which apply to data transferred to the United States, may fall short in a cloud environment, where data flows may concern non-EEA countries other than the United States; and Binding Corporate Rules have yet to be fully endorsed by large CSPs, mainly due to weaknesses in the application and approval process of the BCR regime. Cloud providers often must resort to using model contracts to support repeated or multiple data transfers, but these

¹⁵ Directive 95/46/EC allows personal data to be transferred outside the EEA only when the third country provides an 'adequate level of protection' for the data (article 25) or when the controller adduces that there are adequate safeguards with respect to the protection of privacy (article 26). Binding Corporate Rules (BCRs) are one of the ways in which such adequate safeguards (article 26) may be demonstrated by a group of companies in respect of intra-group transfers, although the BCRs are not a tool expressly listed and set forth in the Directive. See Article 29 in the Data Protection Working Party Opinions 74, 133, 153, 154, and 155; all available at: <http://ec.europa.eu/justice_home/fsj/privacy/workinggroup/wpdocs>.

can be burdensome to implement especially where national regulators impose additional administrative requirements (such as a duty to seek regulatory approval of a contract). In light of these challenges and as part of its review of the EU data protection framework, the European Commission is seeking to improve mechanisms to transfer personal data outside of the EEA. The Commission also is encouraging self-regulatory initiatives such as codes of conduct or codes of practice.¹⁶ However, in cloud services provided to governments and PAs, all the arguments concerning government sovereignty presented above will be a consideration with regard to data transfers. Last but not least, it is worth mentioning that there are quite some issues concerning the transfer of patient data, which are detailed in the section dedicated to the e-health scenario.

- *Data subject's right of access to data (Directive 95/46/EC, Article 12)*

The Controller has the obligation of guaranteeing the data subject the rights laid down in article 12; eg, to obtain confirmation as to whether or not data relating to the data subject is being processed, to obtain information on the purposes of the processing, the categories of data concerned, the recipient or categories of the recipients to whom the data are disclosed, to rectify, erase or block the data processed in a way which is not compliant with the provision of the Directive, etc. It is extremely important, especially when the CSP falls under the definition of processor, that the CSP engages in very close cooperation with their customers (ie, governments and PAs) to ensure that the latter, in their capacity as controllers, are in a position to fulfil their data protection obligations towards the data subjects. It is advisable to specify the terms of such cooperation between the parties in the relevant contract. Specific issues in this respect emerge in the e-health scenario, where it is a fact that not only has Directive 95/46/EC been implemented in an inconsistent way but also that patients' right are defined and implemented differently under various national laws.

Interoperability / Transfer back / 'Vendor lock-in' provisions

A cloud solution should be interoperable, enabling governments and PAs to migrate cloud services from one CSP to another without technical or contractual restrictions or substantial switching costs. Furthermore, interoperability will be a necessary condition in the e-health scenario. Moreover, the timing and modalities of information and data transfer back should be defined in contracts. It is

¹⁶ The speech by Neelie Kroes, European Commission Vice-President for the Digital Agenda, on Cloud computing and data protection at Les Assises du Numérique conference, Université Paris-Dauphine, 25 November 2010; available at: <http://europa.eu/rapid/pressReleasesAction.do?reference=SPEECH/10/686&format=HTML&aged=0&language=EN&uiLanguage=en>.

extremely important for governments and PAs to avoid any form of ‘vendor lock-in’, as any (temporary) unavailability and/or inefficiency of services may lead to significant liabilities for governments and PAs (one can think about the damage and liability that can occur in the e-health scenario).


CSP professional negligence

By migrating to cloud services, governments and PAs become very dependent on the adequacy of a CSP’s performance. CSP failures or shortfalls in the provision of the cloud services will most likely have a very negative impact on the services offered by governments and PAs to citizens. This may translate not only into economic losses for governments and PAs but also into damage to their image (thus political damage). The liability and indemnity clauses in SLAs will play a fundamental role in this matter. Detailed SLAs, in which CSP levels of performance are accurately spelled out, coupled with contractual clauses that clearly allocate, on the one side, the general duties and obligations of parties, and, on the other side, the parties’ liabilities and responsibilities will be crucial interests of governments and PAs. They should request CSPs to be vigilant in order to avoid mistakes and assure this through contractual clauses that set significant penalties for shortfalls in CSP services.

Subcontracting of cloud services and CSP change of control

Given the highly dependent relationship, it is likely that governments and PAs will carefully select CSPs. Situations in which a CSP subcontracts the relevant services to a third party should be avoided or, at least, representations and warranties on possible sub-contractors should be included in the service agreement. Similarly, changes of control should be promptly notified by the CSP to the government or PA, which may want to negotiate the right to terminate the contract should such an event occur.

3.4. Architectural options



In this paragraph we have proposed short definitions for cloud models.

Non-cloud

- *Fully-owned and managed*: the IT services are provided through an infrastructure and platform that is fully owned and managed by the same entity that uses the services.
- *Outsourced*: the IT services are contracted out to a third party. The services might be provided from an infrastructure or platform owned by the same entity that uses the services (eg, an internal customer) or one that owned by the service provider itself. The service provisioning is regulated by a contract in which the terms, conditions, penalties and duration are clearly defined.

For a more in-depth description of a typical non-cloud IT service architectural option, please refer to existing literature (eg, ITIL).

Cloud

For the definitions of the various forms of cloud computing, we mostly refer to NIST (20)

Deployment models and ownership

- *Private*: the cloud infrastructure is operated solely for a particular organization. It may be managed by the organization itself or by a third party, and may exist on the premises or off the premises.

- *Public*: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization that sells cloud services.
- *Community*: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (eg, mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on the premises or off the premises (20). The 'cloud infrastructure' could be either a solely-owned data centre or a network (federation or community) of (smaller) data centres (21).

For an example of a federated or community cloud, please see Annex III, [Reservoir architecture description](#).

- *Hybrid*: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (eg, cloud bursting for load-balancing between clouds).


A possible configuration for a hybrid cloud is represented by a private cloud that scales out into a public cloud. On the basis that the hybrid model requires the combination of two clouds, we assume that a hybrid cloud represents a second step in a cloud approach. Given the short-term time horizon of this report, we will exclude the hybrid cloud from our analysis.

Computing and delivery models

- *Cloud Software as a Service (SaaS)*: the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (eg, web-based email). The consumer does not manage or control the underlying cloud infrastructure which includes the network, servers, operating systems and storage, or even the individual capabilities of the application, with the possible exception of limited user-specific application configuration settings.
- *Cloud Platform as a Service (PaaS)*: the capability provided to the consumer is to deploy, onto the cloud infrastructure, consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure, which includes the network, servers, operating systems and storage, but does have control over the applications deployed and, possibly, the configurations of the application hosting environment.
- *Cloud Infrastructure as a Service (IaaS)*: the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the

consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over the operating systems, storage and applications deployed and, possibly, limited control over select networking components (eg, host firewalls).

4. SWOT analysis



This chapter presents the results of a comparative assessment of public, private and community cloud deployment models based on the security, resilience and compliance parameters set out in section [3.1.2](#).

The analysis identifies the strengths, weaknesses, opportunities and threats for each cloud model.

We use SWOT analysis as a tool for comparison, but more exhaustive methods, such as risk assessment, could be used instead. As a matter of fact, an analysis of strengths, weaknesses, opportunities, and threats should be considered as the initial (and minimum) action to be undertaken, while a more detailed risk assessment has to be carried out to support a more precise identification and assessment of the risks affecting a particular organization.

In order to support public administrations in carrying out their own risk assessments, we have included a list of threats in [Annex IV](#).

The analysis does not take into account any specific requirement and it is meant to be read as a first general assessment of different possible candidates as cloud deployment models. By reading this chapter the audience should have the information needed to identify the most suitable cloud model for its circumstances. A more detailed assessment of the adopted model may then be implemented.

In chapter 4, where we consider the specific requirements for four sample services (electronic healthcare services, electronic administrative procedures, email, and human resources applications), we have provided more concrete advice.

4.1. Public cloud

Strengths

The public cloud model appears to be best positioned to offer strong resilience, in particular with respect to performance and reliability figures. In more detail, we may say that:

- Availability and reliability: the very large pool of resources simplifies the handling and the masking of faults in hardware resources and results in better reliability figures for the service to be implemented.
- Tolerance and elasticity: the very large resource pool simplifies the handling of performance loss due to peaks in demand as the service of interest may exploit available resources to avoid a drop in performance for the final users. However, this requires the proper design and implementation of the applications and proper application monitoring to detect losses in the performance of the user application. Monitoring is fundamental for exploiting cloud resources and achieving the target performance.
- Patch management: SaaS can guarantee the best performance in mean time between patches. In SaaS, users have less responsibility, but proper controls have to be introduced to guarantee that patches are actually applied. In PaaS and IaaS, users have a higher degree of responsibility¹⁷.
- Response time: performances in terms of the reliability and elasticity that may be achieved by properly reconfiguring an application make it possible to best exploit available resources so that the response time for the final user may be always kept at a predefined interval.
- Business continuity: the resources of a public cloud implemented by a large provider may be geographically distributed. This simplifies the definition of business continuity and disaster recovery strategies.

¹⁷ The attribution of responsibility has been analysed on pages 64-65 of the ENISA report 2009.

- Physical security: very strong security may be achieved at each site of the provider because no one is allowed to perform an audit on site and strong control over physical access can be implemented.
- Intrusions prevention and detection: given the large quantity of resources in the cloud, some of these can be devoted to integrity monitoring and intrusion detection to uncover malicious attacks without decreasing the final performance to the user.
- The strong physical security measures may allow the provider to delay possible subpoenas and e-discovery from law enforcement agencies of other countries.

Most of the benefits are due to the very large size of the pool of available resources and to their geographical distribution (as already mentioned in the [report](#) on security benefits and risks published by ENISA in 2009). The homogeneity of the resources used to build the cloud can strengthen and simplify the design and the management of the overall system. This implies that the strengths are directly proportional to the scale of the cloud provider. It is not unrealistic to imagine that, in the near future, public cloud providers will offer more governmental-specific private cloud services than they are doing now.

Weaknesses

The major weaknesses of a public cloud solution for governmental organizations are related to the lack of governance, the large number of tenants (users) in the cloud and to the strong negotiating power of the cloud provider in the definition of the contract. In more detail, some of the major weaknesses of the public cloud model from the perspective of a public body are as follows:

- The lack of legal and regulatory compliance (data retention, forensics, reporting): these threats are exacerbated in the SaaS and PaaS deployment model where the user has less governance, control and visibility over the infrastructure.
- Lack of control over the supply chain: it should be noted that in the case of IaaS the control on the service provisioning supply chain is higher than in PaaS and SaaS, but this potential benefit should always be compared with the extra costs generated by the platform and the management of software security and resilience.
- Logging capabilities: public cloud providers normally do not offer a sufficiently detailed logging capability on cloud operations and administration and, perhaps most importantly, there is a lack of information on incident response and forensics.

- Difficulties in accessing forensic data to determine data linkability and accountability in cases where illegal activities are performed.
- Lack of the necessary bargaining power of certain public organizations when negotiating terms and conditions and requesting an adequate degree of transparency from the provider(s).
- Specific legal and regulatory requirements that, in some countries, force public organizations to keep data within the national territory and reduce the degree of business continuity that may be achieved.
- Degraded performance (IaaS, PaaS and SaaS) due to poor quality in the structuring of an application and its ability to dynamically exploit available resources in order to handle faults and/or peaks in demand.
- Degraded performances (IaaS, PaaS and SaaS) due to poor quality in the connectivity (eg, rural areas, especially in the southern and eastern countries in the EU). This applies only to cases where the customer is situated in specific areas. For more distributed customers this is not an issue.
- Limited local distribution of the data centres in EU territory which can have an impact on the performance of the service. These considerations appear to be especially true in a situation where a public authority is located in a remote place (eg, ENISA located in Crete) that can most likely suffer from degraded performance due to poor quality in the connectivity.
- Difficulties in transferring data back to the user or on to the chosen alternative CSP. These difficulties appear to be a serious problem especially for healthcare services in which a failure or a delay in transferring healthcare related information can represent a serious threat for the healthcare authority as well as for patients.

Opportunities

Compared to in-house solutions, public clouds could provide opportunities to improve the current practices of potential governmental users, in the areas of preparedness and legal compliance and, more so, in particular in:

- risk analysis and assessment
- security testing
- real-time security monitoring

- forensics (please note that the apparent contradiction between having ‘forensics’ both as a weakness and an opportunity is due to the fact that at present such services are not offered by cloud providers, but they could become a factor in the differentiation of offers in the near future; thus we see as an opportunity (15).

This is due to the following reasons:

- It is difficult to rely on internal specialized staff to carry out, on a periodical basis, risk analyses and assessments, and security tests.
- The resources needed to build an in-house security operations centre to perform real-time security monitoring, or to buy those services in the market, are costly.
- Market or competitor pressures that will force public cloud providers to offer security features will deliver added-value for customers.
- Compliance pressure.

In order for a public cloud to take advantage of these opportunities the following measure should be in place:

- full control over asset inventory;
- detailed classification of physical assets, information and services;
- integration between risk analysis/assessment and real-time security monitoring processes;
- effective screening of the provider’s employees.

Threats

Various threats apply to the public cloud model and most of them have already been identified by ENISA as well as other organizations (eg, the Cloud Security Alliance, and LinkedIn interest group).

The biggest threats that public authorities selecting a public cloud solution will face are:

- A large public cloud is an attractive target for threat attacks due to the large quantity of information attackers can access after successful attacks. The size of this information justifies even a large investment in time and resources to implement an attack.
- The impact of attacks from insider threats may be rather large due to the amount of information stored in the cloud. Detailed logs of insider activities should be preserved, job rotation policies should be adopted by the provider, and need-to-know policies should be adopted.

- Isolation failure (15) can open the door to information leakage (illegal monitoring) as well as operational problems due to a lack of isolation from other tenants' resources. In this case, loss of information may be the result of an attack against another user of the public cloud.
- Poor definition of requirements and of the classification of assets may result in the exposure of the assets to other users of the cloud.
- Multiple jurisdictions may apply when the sites of the provider are distributed across several nations.
- A change in the control of the provider may result in the adoption of distinct security strategies as well as distinct marketing strategies that result in a lower quality of service.
- In SaaS or PaaS solutions, a proprietary format may be adopted to store data in the cloud. Moving to another provider can be almost impossible if there is no tool to automatically translate data into the new format

A detailed list of threats, which apply to all cloud models, can be found in [ANNEX IV](#).

4.2. Private cloud

Strengths

In a private cloud the owner-user has, in principle, full control (subject to economic constraints) over the feature set of the cloud implementation; however there are costs (which cannot be shared with other customers) associated with this increase in control.

The following list contains the most important features (concerning security and resilience) that can be defined in a private cloud:

- Risk assessment practices: it is possible to select the methodologies, scales, metrics, etc.
- Patching: it is possible to schedule patching when required, and also adjust the regime.
- Access control: finer granularity of access management and policies to prevent data leaks.
- Logging: it is possible to control what is logged, where it is stored, how storage is protected and how long it is stored for.
- Auditing: it is possible to establish and regulate the right to audit.

- Control over availability, reliability, scalability and elasticity: it is possible for the customer to specify the system and define SLAs for the private cloud to match, within technical constraints, the required service performance.
- Availability of the management interface: premium services can be negotiated more easily with ISPs to obtain a better network and connection (eg, priority in service resumption).
- Business continuity plan: the plan can be defined and all its components tested.
- Legal compliance: full transparency and control over legal requirements such as data location.

Weaknesses

- The beneficial effect of the economies of scale in private clouds is likely to be much less compared to public clouds (at least the large-scale ones, currently present in the market) or even communities.
- The possible lack of an adequate scale also represents a weakness in the purchase and implementation of security mechanisms.
- There is potentially less tolerance of malicious attacks than in a public cloud, on the assumption that available resources (especially in terms of computing capacity) may be less adequate than those of a public cloud. In some cases, also, the internal expertise of the provider may not be adequate.
- There is less flexibility for meeting unanticipated peak demands, due to the paucity of resources. This requires capacity planning and some benchmarking before moving to the cloud.
- Realistically, it is feasible for a private cloud to define a comprehensive redundancy regime; however it is highly unlikely that this will be equal to or better than the redundancy regime offered by the public cloud of a major cloud provider.
- Lack of geo-redundancy is a problem as far as business continuity is concerned. In general, the time to recover from a failure of the private cloud may be rather longer than a public cloud unless specific mechanisms and policies are implemented by the provider. An adequate SLA in this regard should be defined with the provider.
- Sensitivity of reputation: the reputation of governments and public bodies may be extremely sensitive to the leakage of information and any other security incidents, including the use of a government-owned infrastructure to launch malicious attacks.

Opportunities

- **Monitoring:** in a private cloud, user and applications oriented monitoring mechanisms can be implemented making a quick adjustment of resources to meet peaks in demand possible. Furthermore, security events of interest can be fully monitored. As a counterpart, if the scale of the private cloud is not appropriate, handling peaks of demand for resources can be rather complex and no efficient solution may exist for unanticipated peaks. However, cloud resources should be exploited to improve the performance of the applications that are moved to the cloud.
- **Access control:** if required, access policies based upon non-discretionary access control systems (eg MAC (mandatory access control) or RBAC (role-based access control)) may be more easily adopted to further confine each user and minimize illegitimate flows among users.

Threats

A government or public body willing to build and use a private cloud should be prepared to face the following threats:

- **Politically motivated attacks:** while the quantity of information managed by the cloud may not be attractive per se, the defacing of a government site may be attractive for politically motivated reasons. (Obviously this threat is not unique to a private cloud, but a private governmental cloud could present a very high concentration of resources and therefore the incentive for a motivated attacker would be even higher).
- **Big brother effect:** the fact that the government will be collecting and managing information about citizens and eventually businesses (should the cloud be used as a business incubator for SMEs) it could be perceived, from the perspective of end-users, as a possible way to put a profiling and surveillance system in place.
- **High volatility in resource utilization and unanticipated peaks in requests** could force a private cloud to scale out into a public cloud (hybrid cloud), outside the realm of the defined security policy. In such a case, the control over the information in the cloud is partially lost anytime the security policy, which rules the information that may be exported, has not been defined.
- **Poor planning:** for example, the definition of requirements and classification of assets may result in a loss of security and integrity when scaling from a private cloud to a hybrid one.

- Inadequate definition of contracts with business partners (cloud operator, technology partners, hardware and software providers, etc) and lack of monitoring the execution of the contracts may be critical in relation to the size of the provider.

4.3. Community cloud

When analysing a community cloud, one should consider that, in principle, its strengths and weaknesses fall between those of a private cloud and those of a public one. In general, the pool of available resources is larger than in a private cloud with obvious benefits in terms of elasticity. However, the pool is not as large as that of a public cloud and this limits the elasticity offered by a community cloud. On the other hand, the number of users in a community cloud is much fewer than in a public cloud which has obvious benefits in terms of security.

Strengths

- Common requirements and constraints and risk profile: the users of a community cloud have similar requirements from a security and performance perspective. This makes the implementation of policies to satisfy these requirements more efficient and cost effective, even for the provider, resulting in a lower overall cost.
- Common requirements and risk profiles simplify the configuration of mechanisms and tools to protect the applications running on the cloud from internal and external attacks.
- Users have more bargaining power as a group (*vis-à-vis* the cloud provider) due to the larger number of users with similar requirements.
- Ability to set the entry criteria: memberships are issued according to the trustworthiness of potential members. This strongly reduces the risks due to attacks from another cloud user.
- Larger scale and better response to high peaks in resource demand (compared to a private cloud): the size of the resource pools may be noticeably larger than those in a private cloud and this simplifies the management of peaks in demand for resources.

Weaknesses

- There is more resource competition between the partners since they have common goals. Some of the benefits arising from a larger number of resources are lost because users in the

same community may exhibit similar patterns in accessing resources so that peaks in requests for resources by multiple users may arise in the same time window.

- Compared to a private cloud, a community is a more attractive target for motivated attackers due to the larger visibility achieved by successful attacks. Furthermore, the applications of other users may provide an avenue for attacks.
- Access control and authentication are weakened compared to a private cloud due to the larger number of users.
- Degraded performance (IaaS, PaaS and SaaS) due to poor quality in connectivity (eg, rural areas, especially in the southern and eastern countries of the EU) may reduce the quality of service for some users in the community (who are not in proximity to points of delivery) compared to a private cloud.

Opportunities

- Similar requirements across the community (see strengths) could allow improved security policies, baselines and standards as well as common practices for risk analysis and assessment, logging, and monitoring. This may result in highly efficient implementations that reduce the adoption cost for each user and result in a more resilient architecture.
- Common and shared incident management systems can simplify the adoption of mechanisms to store and manage forensics evidence.
- Information sharing among other community members (best practices in use, experience from past incidents, etc) may result in a larger diffusion of best practices, fine-tuned by the most expert community members.
- Stricter security may result because information about security policies, and the design and implementation of the cloud are shared only within the community. Compared to a public cloud, this increases the complexity, for an attacker, of acquiring information to implement its attacks.

Threats

- Lack of agreement on security baselines and security mechanisms: to exploit the opportunity to share mechanisms to protect and defend information, an agreement among all the community member has to be negotiated. A renegotiation that involves even just a few users may be rather complex and unsuccessful most of the time.

- Communities may grow either too quickly, which will eventually decrease the advantages of a community cloud in terms of flexibility compared to a private one, or grow too slowly, which will eventually affect dynamic scalability.
- Harder to predict resource usage (than in a private cloud): the larger number of users increases the complexity of anticipating resource requests from each user. Errors in the capacity planning of the community cloud are more likely.
- Failure of isolation mechanisms may result in the leaking of information which is more difficult to control because of the large number of users.
- It is difficult to identify the legal entity that is responsible for acting against a member of the community or the provider when super-national issues are involved.

5. Example scenarios

In this chapter we demonstrate the simple decision-making model with three scenarios. These scenarios are partly based on real-life concrete experiences.

The scenarios are based on the following use-cases:

- Healthcare cloud: the use of cloud computing in the implementation of an electronic health record service in national, regional and local healthcare authorities;
- Local and regional authorities;
- Governmental cloud as a business incubator.

For the sake of brevity we have included, in this chapter, only a description and analysis of four sample services which are representative of these scenarios: 1) Electronic Healthcare Record (EHR), 2) Electronic Administrative Procedure (EAP), 3) email, and 4) human resource applications.

For more details on the scenarios, please read [Annex II](#).

5.1. Service description

Electronic Health Record (EHR)

The *Electronic Health Record* (EHR) is a repository of information regarding the health status of a subject of care in computer processable form, stored and transmitted securely, and accessible by multiple authorised users. It has a standardised or commonly agreed logical information model which is independent of EHR systems. Its primary purpose is the support of continuing, efficient and quality integrated healthcare and it contains information which is retrospective, concurrent, and prospective (22).

- The EHR is a secure, real-time, point-of-care, patient centric information resource *for clinicians*.
- The EHR aids decision-making for clinicians by providing access to patient health record information where and when they need it and by incorporating evidence-based decision support.
- The EHR automates and streamlines the clinician's workflow, closing loops in communication and response that result in delays or gaps in care.

The EHR may also simplify the collection of data for uses other than direct clinical care, such as billing, quality management, outcomes reporting, resource planning, and public health disease surveillance and reporting.

The essential requirements of EHR (23) are to:

- provide secure, reliable, real-time access to patient health record information where and when it is needed to support care;
- guarantee that confidentiality and security of patient health information;
- be available and reliable 24/7;
- be responsive enough to integrate with the clinician's workflow;
- be accessible where needed – at inpatient and ambulatory care sites, with remote access.

Electronic Administrative Procedure (EAP)

Electronic Administrative Procedure (EAP) essentially concerns the tasks of electronic filing, and document services regarding a government's interaction with citizens, businesses, and other arms of government. Technically it describes the electronic management of administrative procedures relating to files and records.

As such, EAP comprises, for example, online requests regarding subsidies, aid, licenses, attestations, or forms and, to support these requests, may allow for:

- submission of requests;
- retrieval and access to status information about orders or processes;
- interactive access to pending orders or processes;
- information on required interactions or documents
- facilities to directly provide information or supply documents;
- notifications;
- retrieval of documents and forms;
- electronic payments.

A detailed description of these services, as well as others offered by public organizations considered in this report are published as [Annex II](#), together with the use-case scenarios.

Email

Electronic mail is the well-known means of communication used to exchange digital messages.

Email is considered in many organizations as a business critical service and often confidential information is exchanged via email.

Human resource applications

HR applications are those IT services that provide support to the management of HR. They consist of tracking existing employee data which traditionally includes personal histories, skills, capabilities, accomplishments and salary. Human resource management systems encompass:

- payroll
- work time
- administration of benefits
- HR management information systems
- recruitment
- training and learning management systems (LMS)
- performance records.

The payroll module automates the pay process by gathering data on employee time and attendance, calculating various deductions and taxes, and generating periodic pay cheques and employee tax reports. Data is generally fed from the human resources and time-keeping modules to calculate automatic deposit and manual cheque-writing capabilities. This module can encompass all employee-related transactions as well as integrate with existing financial management systems.

The work-time module gathers standardized time and work-related efforts. The most advanced modules provide broad flexibility in data collection methods, labour distribution capabilities and data analysis features. Cost analysis and efficiency metrics are the primary functions.

The benefits administration module provides a system for organizations to administer and track employee participation in benefit programmes. These typically encompass insurance, compensation, profit sharing and retirement.

The HR management module is a component that covers many other HR functions from job application to retirement. The system records basic demographic and address data, selection, training and development, capabilities and skills management, compensation planning records and other related activities. Leading edge systems provide the ability to 'read' job applications and enter relevant data into applicable database fields, notify employers and provide position management and position control. The human resource management function involves the recruitment, placement, evaluation, compensation and development of the employees of an organization.

Online recruiting has become one of the primary methods employed by HR departments to garner potential candidates for positions available within an organization.

The training module provides a system for organizations to administer and track employee training and development efforts. Sophisticated LMSs allow managers to approve training, budgets and calendars alongside performance management and appraisal metrics (24)

5.2. Parameters and requirements

The fictitious governmental organizations and local authorities considered in this report should take into account some key parameters when assessing the cost/benefit impact of a cloud approach on services security and resilience.



In the table ‘Service attributes’ included as Annex III, we suggest a few general parameters and indicate the associated requirement for the type of service considered. It should be noted that:

- Again, the parameters should be understood as just examples of possible constellations of the variables mentioned in [chapter 3](#) (Model for decision-makers). This is intended to be instructive to the reader and show him how to interpret the above variables.
- The requirements have been derived either from the direct responses to questionnaires answered by local and regional authorities and healthcare authorities or are based on the experience of the members of the expert group.

Service sample			
EHR	EAP	Email	HR Apps
	Electronic Administrative		

		Procedures		
Parameters	Requirements			
Data Sensitivity				
Type of data	Personal Data Sensitive Data	Personal Data Sensitive Data	Personal Data Business Data	Personal Data Sensitive Data
Information security and resilience Requirements	High Integrity High Confidentiality High Availability	High Integrity High Confidentiality Medium Availability	Medium Integrity Confidentiality (content specific) Medium Availability	High Integrity High Confidentiality Medium Availability
Scalability – Demand Management				
Volatility of the demand	High (for repository accessible from patients and research) Low (for EHR)	High	Medium	Medium
New services required	Yes	Yes	No	NO
Anticipated storage requirements for the next five years	Predictable	Predictable	Predictable	Predictable
Peak of concurrent users	High	High	High	Medium
Proportion of data in active use	Low	Medium	Low	Medium

Level of administrative (privileged user – IT dept) access required	Low	Low	Low	Low ¹⁸
Service reliability – Availability and Performability				
Availability Required	99.9% (High)	98% (Medium)	97% (Medium)	98% (Medium)
Unplanned downtime requirements	Very short – no longer than 1 hour	No longer than 4 hours	No longer than 2 hours	No longer than 4 hours
Real-time response	Low	Medium	Medium	Medium
Collaboration and Interoperability				
Other health authorities and public administrations need to access the service	Yes	Yes	No	Yes
Identity, Authentication and Access Management				
Identity management	Patient identities have to be managed internally	Citizens identities can be managed both internally or via an outsourcer (eg, cloud	User identities can be managed both internally or via an outsourcer (eg,	User identities can be managed both internally or via an outsourcer (eg,

¹⁸

On the assumption that the system has been properly architected

Making an informed decision

		provider)	cloud provider)	cloud provider)
Credentials and permissions provisioning for users	The process of credentials and permissions provisioning to the users (patients, doctors, admin personnel, etc) have to be managed internally	The process of credentials provisioning to the users (citizens, admin personnel, etc) can be managed both internally or via an outsourcer (eg, cloud provider)	The process of credentials provisioning to the users can be managed both internally or via an outsourcer (eg, cloud provider)	The process of credentials provisioning to the users can be managed both internally or via an outsourcer (eg, cloud provider)
RBAC	YES	YES	NO	YES
Strength of authentication	Strong Legal requirement	2-factor (optionally)	Medium	Pw (optional)
Federation required	Yes	Yes	No	No

Encryption

Encryption	YES in transit – recommended to encrypt the rest as per legal requirements	Optional	Optional	Recommended in transit (as per regulations)
Access to keys				
Credentials and permissions provisioning for admin access	Provider to provision authentication credentials for admin access	Provider to provision authentication credentials for admin access	Provider to provision authentication credentials for admin access	Provider to provision authentication credentials for admin access

Legal and Compliance

Data protection	Applicable	Applicable	Applicable	Applicable
Data location and legal jurisdiction	Both need to be specified (The law in some countries imposes a requirement that the data cannot leave the national territory.)	Both need to be specified (The law in some countries imposes a requirement that the data cannot leave the national territory.)	Both need to be specified (The law in some countries imposes a requirement that the data cannot leave the national territory.)	Both need to be specified (The law in some countries imposes a requirement that the data cannot leave the national territory.)
Access control	A combination of mandatory access control (MAC) ¹⁹ and role based access control (RBAC) systems should be in place.	A mandatory access control system (MAC) or a RBAC should be in place.	A mandatory access control system (MAC) or a RBAC should be in place.	A mandatory access control system (MAC) or a RBAC should be in place.
Accountability (court admissible logs)	YES	YES	YES	YES

¹⁹MAC has been suggested as a requirement for EHR for the following reasons: 1) The patient is the owner of his EHR and he should be able to decide a) who can access which type of data and b) to whom to delegate this decision. 2) Confidentiality and integrity is of the highest importance in EHR, and the patient must be able to enforce rules about any potential access beyond or at superior levels to the ownership model (eg, super-user or root in a UNIX system). Based on these assumptions, the MAC system appears to be more appropriate than, for example, a DAC, given the possibility that the former offers to set definitive access rules. For more information, see: (http://en.wikipedia.org/wiki/Discretionary_Access_Control), (http://en.wikipedia.org/wiki/Mandatory_Access_Control) and (http://en.wikipedia.org/wiki/Role-Based_Access_Control).

Making an informed decision

Access using digital IDs (citizen cards)	YES	YES	NO	NO
Digital signature	YES	Processes and documents must be validated with a digital signature (internally or by the citizens). Healthcare providers are signing the data.	Official emails can be digitally signed.	NO
SSO – single sign-on	Optional	Optional	Optional	Optional
Non-repudiation	YES	YES	YES	YES
Electronic time-stamping	Yes – needed for audit trail, etc, for medical investigations	Some documents (eg, permits, licenses, payments, etc), submitted or issued, must have a timestamp granted by a certified authority.	NO	NO
Single passwords or unique usernames	YES	YES	YES	YES
Enforcement of the need to know principle	YES	YES	YES	YES

(application)				
Transparency of the supply chain	Full transparency against third party providers is required.	Full transparency against third party providers is required.	Full transparency against third party providers is required.	Full transparency against third party providers is required.
Supply chain due Diligence	YES	YES	YES	YES


TABLE 1 – SERVICE ATTRIBUTES

5.3. Comparative risk assessment

In this section we carry out a comparative analysis of public, private and community clouds in order to find out which type of cloud could be, in the context of the proposed [scenarios](#), the most adequate solution to satisfy the requirements for service described in section [5.2](#).

As already explained at the beginning of this chapter, the service requirements on which the assessment is based represent only possible configurations of real-life situations and should not be taken ‘as is’ in a specific analysis. They demonstrate how the methodology in this paper may be used to make a decision.

Making an informed decision



As the four services considered present similar requirements, we have reported in the table that follows only a comparative assessment for the EHR service. But we have also included a reference to the other services, mentioning their specificities.

Electronic Healthcare Record
Parameter: Data Sensitivity and Criticality
Requirement
EHR: in providing electronic healthcare records, a healthcare authority manages personal and sensitive data. High integrity, high confidentiality, and high availability are the information security and resilience requirements for the EHR service. Next to that, non-repudiation and audit logging are increasingly important requirements when dealing with EHRs. EAP, EMAIL and HR APPS: high integrity, high confidentiality, and medium availability are the information security and resilience requirements for these services.

Community	Private	Public
EHR: The community cloud model has strengths similar to a private cloud in terms of control over data confidentiality and integrity, but if sharing by community members is required then weakening the security policy of multiple private clouds islands (or internal IT infrastructures) in order to accommodate sharing may lead to security problems and policy conflicts that create a more insecure environment than a community cloud where security policies and controls are tailored to the community. A high level of availability can be offered also, especially where the scale of the community's infrastructure(s) is adequate. A federation of clouds provide benefits for resilience in terms of availability, elasticity, and incident management. It should be noted that a federation of private cloud islands combined with some kind of virtual community management layer might be a better option than a hosted community when a balance between separate but interdependent private and shared services is required.	EHR: A private cloud seems to be the best solution to guarantee full control over data confidentiality and integrity. The cloud owner (eg, national, regional or local public body) is responsible for building, managing, maintaining, monitoring, and the evolution of IT services. Better performance regarding data availability than internal IT services can be achieved only if: - the scale of the cloud is adequate; - the cloud is properly managed, maintained and monitored. EAP, EMAIL and HR APPS: A private cloud seems to be the best solution to guarantee full control over data confidentiality and integrity. At the same time a private cloud could easily provide medium level data availability. In a private cloud it is easier to harmonize and enforce security policy and apply uniform risk	EHR: Public large-scale clouds seem to offer high data availability. The loss of control on infrastructures (IaaS), and eventually platforms (PaaS) and applications (SaaS), represents a serious threat to data integrity and confidentiality, and legal compliance. EAP, EMAIL and HR APPS: For integrity and confidentiality requirements, the same considerations as in the analysis of electronic healthcare services apply (ie, the loss of control is a serious threat that needs to be considered especially with regards to the difficulties of introducing the right to audit into contracts). The opportunity to have high availability guaranteed by a large-scale public cloud is not fully exploited in providing EAP since they have a medium

The lack of trust between members is a serious threat to the well functioning of a community cloud. By having transparency as a driving principle of the community, trusted relationships between members can be established and reinforced. EAP, EMAIL and HR APPS: See EHR analysis. It is a moot point whether local authorities would have adequate power to negotiate a contract with large cloud providers, even if they aggregate their needs (especially security needs). A community cloud would have less expertise for security and maintenance than a private cloud but would have a higher level of expertise with regard to the actual applications Specifically for the EMAIL service: Some public cloud providers provide security features such as content filtering, anti-phishing, anti-spam, or let the users apply their own security solutions.	assessment methods. A private cloud appears to be the solution that offers the most stable environment, a stakeholder's community, ownership (the owner will be a public organization), etc. The opportunity to have high availability, guaranteed by a large-scale public cloud, is not fully exploited in providing an email service since it requires a medium level of availability.	level requirement for availability. Specifically for the EMAIL service: Some public cloud providers provide security features such as content filtering, anti-phishing, anti-spam, or let the users apply their own security solutions.
---	---	---

Parameter: Scalability and Demand Management

Requirement

The IT infrastructure should be designed to cost-effectively manage:

- 1) low and very low workloads;

2) seasonal peaks (eg, influenza, payrolls, etc);

3) sudden and unexpected peaks due to changes in administrative procedures or the implementation of new laws and regulations;

4) implementation of new services;

5) demographic changes;

6) cyber attacks and ICT incidents.

Community

A community cloud can offer scalability and demand management capabilities that are somewhere in between private and public solutions. It can scale better than a private cloud (to meet demand) as, in principle, a larger infrastructure is available; however it cannot fully exploit economies of scale.

EAP: Scalability requirements are the same as for EHR. The same considerations apply.

Private

A private cloud, depending on its scale, is the least adequate option for managing unexpected events, cyber attacks (eg, DDoS) and ICT incidents.

EAP: Scalability requirements are the same as for EHR. The same considerations apply.

Public

Public clouds guarantee high level flexibility and effective demand management.

EAP: Scalability requirements are the same as for EHR. The same considerations apply.

EMAIL and HR APPS: The demand in terms of computational power and storage is predictable; therefore we assume that there is no significant difference in the value that the three cloud models could provide with regard to scalability and demand management.

Parameter: Service Reliability – Availability and Performance

Requirement

The service needs to be available 99.9% of the time and unplanned downtime should be less than 1 hour. High throughput and low latency performance are required.

EAP and HR APPS: The service needs to be available 98% of the time and planned and unplanned downtime should be less than 4 hours. High throughput and low latency performance are required.

EMAIL: The service needs to be available 97% of the time and unplanned downtime should be less

than 2 hours. Medium throughput values and low latency performance are required.		
Community	Private	Public
A community cloud could be unsuitable due to the need for high replication and performance.	A private cloud can offer high replication and performance if and only if the scale of the cloud is large enough.	Elasticity, flexibility, cost-efficiency, and total availability are the strengths of a large-scale public cloud.
EAP, EMAIL, and HR Apps: All the proposed solutions can satisfy the requirement for medium availability. A public cloud is the solution that, as already mentioned several times in this report, offers potentially the highest availability. Private and community clouds could offer higher service performance than a public cloud due to their closeness to the final users. A large-scale public cloud normally offers geo-distribution by default, but they concentrate their data centres in a few Member States.		
Parameter: Business Continuity		
Requirement		
EHR, EAP, EMAIL and HR APPS: The providers of these systems have to implement business continuity plans and be prepared for disaster recovery. EHR: A business continuity plan should take into account that any down-time could be longer than 1 hour. EAP and HR APPS: A business continuity plan should take into account that any down-time could be longer than 4 hours. EMAIL: A business continuity plan should take into account that any down-time could be longer than 2 hours.		
Community	Private	Public
EHR: A community cloud is a bit better than, but is still similar to, a	EHR (as well as all the other services): In a private cloud	EHR: A public cloud, which is implemented by a large

private cloud, unless it is implemented across different regions and countries. EAP: a community cloud is a bit better than, but is still similar to, a private cloud unless it is implemented across different regions and countries. The possibility of obtaining better business continuity levels depends on the number of the entities joining the community. The larger the number and the scale of the entities (public bodies) involved in the community, the higher the possibility of reaching a level of business continuity similar to the levels offered by public clouds.	the degree of business continuity is reduced compared to a public cloud.	provider, may be geographically distributed. Where this is in accordance with the legislation governing the protection of health data in a specific country, it simplifies business continuity. A public cloud, which is implemented by a large provider, may be geographically distributed. Should this be in accordance with the requirements for auditing, accountability and responsibility as well as the data protection legislation of a specific country, it would simplify business continuity. EMAIL: A public cloud, which is implemented by a large provider, may be geographically distributed. Should this be in accordance with the requirements for auditing, accountability and responsibility as well as the data protection legislation of a specific country, it would simplify business continuity.
Parameter: Collaboration and Interoperability		
Requirement		
EHR: Potentially all hospitals, clinics, etc, in national territory as well in the Member States of the		

European Union could have a need to access the service.		
EAP: Potentially all public administrations at any level (local, regional, and national) and law enforcement agencies could have a need to access the service.		
Community	Private	Public
EHR, EAP and HR APPS: In a community cloud, the level of interoperability intra- and extra-community is agreed based on the members' needs and requirements. EAP specific: Local authorities would need to federate some services in a community cloud, which would mean that local authorities would have to consider the availability of each node and interoperability. There would even be a need to ensure consistency of identification and the identification of authorities in a community cloud.	EHR, EAP and HR APPS: Private clouds enable the users to use certain setups, though interoperability needs to be introduced on top or from external sources. It is easier to build auxiliary services on top.	EHR, EAP and HR APPS: In a public cloud, interoperability is an intrinsic property and can allow for a systematic approach.
Parameter: Identity, Authentication and Access Management		
Requirement		
EHR: Patient identities have to be managed internally. The process of provisioning the credentials of and permissions for the users (patients, doctors, administrative personnel, etc) has to be managed internally. A combination of RBAC and MAC is used. Hospitals, clinics, etc, are often data owners. Healthcare providers define the access control policies on the basis of patient consent and national and organizational policies. EAP, EMAIL and HR APPS: The identities of citizens can be managed both internally and via an		

outsourcer or cloud provider).

The process of provisioning the credentials of the users (citizens, administrative personnel, etc) can be managed both internally and via an outsourcer or cloud provider.

A role-based access control system is required for EAP and HR APPS.

Community	Private	Public
EHR: In private clouds: • there is control over administrative access; • it is easier to provide access for patient data; • identity management is easier; • enforcement of MAC is simpler. EAP, EMAL and HR APPS: In private and community clouds, there is control over the identity and access control systems both for users and administrators. Government issued smartcards can be used. Governments can provide PKI.		EHR: A complex access management system, resulting from a combination of RBAC and MAC, is difficult to integrate as well as to enforce in a public cloud. EAP and HR APPS: A complex access management system is difficult to integrate as well as to enforce in a public cloud. The Spanish system for identity management appears to be of particular interest in this regard. In fact, in Spain, the new identity card (DNIE) incorporates a device for the creation and verification of an electronic signature. The verification is performed against two formal systems for the validation of certificates: <i>Fabrica Nacional de Moneda y Timbre</i> of the Spanish Ministry of Industry, Tourism and Trade, and the Ministry of the Presidency. The system works for citizens in the private and public sectors, and the technical specification have now been made public in order to allow distinct developments. Moreover <i>Fabrica Nacional de Moneda y Timbre</i> serves as the PKI infrastructure for public administrations. The legal base of the common approach to the public and private sectors in the Spanish system for identity management is found in Spain's

		Law 59/2003 which is a transposition of the EU electronic signature directive 1999/93/CE. EMAIL: A complex access management system is difficult to integrate as well as to enforce in a public cloud, especially when the SaaS delivery model is considered.
Parameter: Encryption		
Requirement		
EHR, EAP and HR APPS: Encryption of data in transit and at rest is a security requirement. Having private key management has to be possible. EMAIL: Encryption of data in transit and at rest is a security requirement. Due to the nature of the application, email exchange outside the governmental sovereignty or domain will be unencrypted regarding the message's sender, receiver, and concerned routers. This should be accounted for during the specification of the requirement for encryption. Having private key management has to be possible.		
Community	Private	Public
EHR, EAP, EMAIL and HR APPS: In both private and community clouds, the distribution and revocation processes for encryption keys are easier to implement, as well as the mechanisms for key storage and protection.		EHR, EAP, EMAIL and HR APPS: It is difficult to integrate an external key management system into a public cloud. In the course of processing, it may be necessary to decrypt data sets and thereby expose their content in a public infrastructure potentially shared with thousands of other tenants; therefore strong isolation mechanisms are necessary. Traffic analysis can be implemented even when all the data are encrypted.

Parameter: Legal Compliance		
Requirement		
EHR, EAP, EMAIL and HR APPS: Both data location and legal jurisdiction need to be specified. The law in some countries imposes a requirement that data cannot leave the national territory. Full transparency against third-party providers is required. In order to guarantee accountability, court admissible logs are needed.		
Community	Private	Public
In community clouds, legal compliance can be easily achieved given: - the common legal requirements of the community's members, and - the control over the service supply chain.	Private clouds provide the highest confidence that legal compliance can be achieved.	In public clouds, achieving legal and regulatory compliance is difficult or even impossible due to: - uncertain data location and legal jurisdictions; - limited right to audit.
The requirements for court admissible logs can be satisfied by all cloud alternatives. Community and private clouds can better control log collection (level of detail) and retention.		

5.4. Selection of the solution and identification of threats and weaknesses

Legal restraints may require that personal health care data be retained within a specified physical locale. In this case, private and community clouds are the best solutions for the implementation of EHR services (as well as EAP, EMAIL and HR APPS). This is due to the capability of these cloud models to offer legal and regulatory compliance and control over the requirements for high levels of confidentiality and integrity. A community cloud solution that brings together different regional, national or international healthcare organizations is preferred over a community cloud that combines healthcare with other sectors.

Aside from the legal restraints, if the primary focus is on providing the healthcare benefits of an EHR system, large-scale public clouds have the best potential to meet the strict availability and resilience

requirements of an EHR system at a reasonable cost. If the legal and regulatory compliance challenges can be overcome, they will become preferable to private and community clouds unless the health authorities can afford significant investments in security specialists and over-provisioning of resources.


There are a number of threats that should be considered when using private and community governmental clouds for EHR services:

- lack of critical mass for infrastructure;
- politically motivated attacks;
- leakage of EHR in an irreversible failure covering a very long period of personal history;
- the extreme sensitivity of the reputations of governments and public bodies to the leakage of health records and other security incidents, including the use of a government-owned infrastructure to launch malicious attacks;
- loss of data integrity;
- data unavailability;
- poor definition of requirements and classification of assets;
- inadequate terms and conditions in contracts with business partner(s);
- lack of monitoring of contract execution;
- isolation failure (see report 2009) opening the door to information leakage (illegal monitoring) as well as operational problems;
- inadequate identity management and access control systems;
- lack of compliance with data protection regulations.

The above-mentioned list of threats obviously needs to be integrated with all the other technical threats included in [Annex IV](#).

In order to mitigate the threats mentioned, healthcare authorities should take into account the security measures and controls described in chapter 6.

6. Preparing a request for proposal



In this chapter we propose a set of security and resilience questions which can be used as guidance when preparing a request for a proposal. These should be seen either as security measures to be included in the contract specifications or as demands to be fulfilled by a third party. Moreover, the same set of questions should be used as a basis for the preparation of the mitigation plan in the risk treatment phase. The questions cover requests for service proposals from both public clouds, independent service vendors (ISVs), infrastructure service management providers, security services, and other service management providers having privileged access to infrastructures and platforms (eg, service-desk or helpdesk, capacity management, consultants, incident management, etc).

Moreover we advice to leverage the specific controls included in cloud oriented assurance framework such as ENISA Information Assurance Framework and CSA Controls Matrix when preparing the request for service proposal or the risks mitigation plan.

Finally we suggest evaluating the potential the Common Assurance Maturity Model (CAMM) (2) project, which is meant to be a framework for transparently rating and benchmarking the capability of a selected solution to deliver information assurance across the supply chain.

A. Preparedness

- Do you have information security management systems in place?

- How will you facilitate audit and necessary certifications?
- Can you support my data and service classification schema?
- What logging facilities do you provide, how is the integrity of logs ensured and access to logs controlled?
- What personnel security measures do you support?
- Show how you meet government requirements for the selection and vetting of personnel having access to data, infrastructure and management?
- How do you protect privileged access?
- How is my data isolated from other customers' data?
- How is court sanctioned access to my data controlled and guaranteed?
- What mechanisms do you support to manage data access rights by different roles or users?
- Do you support multi-level or multiple manager authorization?
- How do you prevent and detect privilege escalation and compartment jumping?
- Can you implement and guarantee separation of duties between different government entities?
- How is access to data controlled?
- What different levels of access are supported and how they are controlled across different user or operator categories?
- How are different types of authentication credentials supported?
- Is role-based access supported and how flexible is role management?
- What means of providing only specific data under subpoena or forensic investigation do you support?
- How will you ensure separation of interests between potentially competing client services?
- Is the end-user given appropriate guidance and tools to facilitate storage of information with different requirements in terms of sensitivity, availability, and compliance?
- How do you guarantee that the data classes associated with their owners? [Multi-level security policies?] Does the access management system properly translate the clearance of classification pairs in traditional ICT?
- Does the provider offer log segregation by user-customer?
- How will you provide transparency on outsourcing agreements you enter that have a material effect on a government's SLA with you?

- How will you facilitate consistency of policy with interoperating of services?
- How will the interface between the legacy systems and services and the cloud infrastructure be secured?
- What security and access control measures are supported by the legacy system?
- What breaches could result at the interface?
- How do you check the hardware sourcing process (especially for sensitive government operations)?

B. Service delivery

- What level of service availability can you guarantee?
- What is the availability level of the various components in the solution and how do these affect the availability of the service?
- What mechanisms exist to ensure data consistency?
- What measures do you take to ensure complete erasure of data?
- What defensive in-depth measures do you support to guard against unknown threats and vulnerabilities?
- What is your process for mitigating disruptions associated with rolling out configuration and software changes?
- What management process configuration and software changes do you follow?
- How do you ensure that your infrastructure and software is maintained free of known vulnerabilities?
- What is your policy and notification process for upgrades to platform software that require adaptation of client software applications?
- Are the categories of changes clearly defined?
- Ask for continuous notification of categories of change in order to perform risk analysis and assessment with the necessary frequency.

C. Response and recovery

- How quickly can service be restored after a disruption?
- How do you recover from a permanent CSP failure?
- Have you tested your BC/DR plan?
- Should an incident arise, what is the policy for incident notification and reporting?

D. Legal and regulatory compliance

- Can you guarantee compliance with requirements on the geographic location of data?
- If a court subpoena for data is received that conflicts with the local jurisdiction, what are the means of appeal?
- Regulatory due diligence – eg, do some simulations to verify compliance?
- Can you guarantee access to logs in order to demonstrate who had access to which data and when?
- How do you guarantee integrity and non repudiation of logs?
- Do the proposed terms of service express clearly who is responsible for which parts of the security policy in which cases?
- How the principle of accountability is applied and enforced?
- Suppose there is a commercial data protection clause in a law of an EU Member State concerning the protection of data held on citizens. Because of an incident with a foreign citizen there is an investigation. Will the other country's authorities be given access to the data?
- How do I monitor the fulfilment of the contract? What metrics are available to permit real-time monitoring of the SLA fulfilment, eg, jitter, load tolerance, delivery?

7. Conclusions and recommendations

Based on the analysis of the strengths, weaknesses, opportunities and threats relating to the security and resilience of the three cloud models – community, private and public – undertaken by ENISA with the support of the expert group, the following conclusions can be drawn:

- The cloud computing business model, on the one hand, has the potential to offer public administrations substantial benefits and improvements over current IT provisioning, including:
 - increased availability and reliability, as pooled resources simplify the handling and masking of hardware faults and of performance loss due to peaks in demand;
 - stronger security, as many CSPs have greater and better security expertise, management and controls than government agencies and enterprises;
 - better value, as the cloud model offers economies of scale and the provision of cloud services can easily be changed in response to the fluctuating IT needs of government agencies.
- On the other hand, it still shows weaknesses and exposures to significant threats that could undermine the full exploitation of all the benefits that such a model could offer. Weaknesses and threats are mainly linked to the lack of governance and control over IT operations and the potential lack of compliance with laws and regulations. National laws and regulations in the Member States of the European Union currently impose some restrictions on the movement of data outside national territory; moreover, a problem exists in the determination of the applicable body of law (governing laws) when data is being stored and processed outside the European Union or by a non-EU service provider. The main questions that each public organization, and more generally each EU central government, must address are:
 - whether current legal frameworks can be changed to facilitate the communication, treatment and storage of data outside national territory without exposing the security and privacy of citizens and national security and economy to unacceptable risks;
 - if so, whether moving citizens' data outside the national territory is a risk that may be undertaken;
 - whether the trade-off between the risks of losing control over data and the beneficial effects of geo-distribution is positive for them.

These considerations apply in general to all cloud deployment models (ie, public, private, community and hybrid), but the impact of the weaknesses and threats varies depending on the specific internal and external environment of public organizations in different Member States and depending on the deployment and delivery model considered. In legal and data governance terms, certainly public clouds represent the most risky solution when compared with community and private clouds, for the following reasons:

- Cloud owners (public cloud providers) and users (public bodies) have different missions and interests that can sometimes be in conflict.
 - Public clouds can be owned by non-EU companies.
 - Public cloud providers offer a lower degree of transparency about their security and resilience measures compared to any other cloud options or internal IT.
 - Cloud providers are not obliged to report on security and resilience incidents, and while it is possible for users to identify incidents that have an impact on service availability, identifying breaches of integrity, confidentiality and data protection and their impact is not an easy matter. As far as private and community clouds are concerned, we assume that the owner of a private cloud and the members of a community cloud will have a more transparent attitude toward the reporting of incidents to citizens or users and that, in the case of a cloud operated by a third party, the contract can include a clause obliging it to notify and report incidents.
- Internet connectivity is a fundamental building block of the cloud model; without connectivity it is obviously not possible to access cloud services. The quality and performance of communication services (capacity, latency, etc), are often not homogeneous in the European Union, and there are still areas (especially rural areas) in several Member States where the quality of the service is quite poor.
 - The lack of governance and control, as mentioned in point 2 of this section, appears to be an inherent weakness of the cloud model (especially with regard to public clouds and SaaS deployments), even though, as already stated several times by ENISA (eg, in the 2009 report), the situation can be improved by achieving transparency in the market and negotiating appropriate terms and conditions in contracts. It should be noted that the principle of transparency (ie, transparency for data subjects) is also mentioned in the European Commission draft Communication on 'A comprehensive approach on personal data protection in the European Union' (19) as well as in the recent speech of Commissioner Natalie Kroes on cloud computing and data protection (25).

By reaching an adequate level of transparency both in terms of the security and resilience requirements for public bodies and in terms of the security practices and controls applied by the cloud or service providers, it is possible to facilitate the matching of customer security requirements and the service security levels being offered. The recent NIST publication Proposed Security Assessment and Authorization for U.S. Government Cloud Computing (26) and the CAMM (Common Assurance Maturity Model) project in which ENISA is directly involved are both explicitly relevant here.

- For sensitive applications, private and community clouds appear to be the solutions that currently best fit the needs of public administrations since they offer the highest level of governance, control and visibility, even though, when planning a community or private cloud, special regard should be given to the scale of the infrastructure as most of the resilience and security benefits of the cloud model will not be realised if the necessary infrastructural critical mass is not reached.
- The public cloud option is already able to provide a very resilient service with an associated satisfactory level of data assurance and is the most cost effective. Moreover public cloud offers potentially the highest level of service availability, but due to the current regulatory complexity of intra-EU and extra-EU trans-border data transfer, its adoption should be limited to non-sensitive or non critical applications and in the context of a defined strategy for cloud adoption which should include a clear exit strategy. At the same time a number of emerging initiatives, including CSA Guidance, Control Matrix, and Consensus Assessment as well as the work of the Common Assurance Maturity Model (CAMM) (2) consortium are pushing the yardstick on providing the transparency and assurance that will allow using public cloud model in more sensitive applications.
- Regardless of the deployment model chosen, it is clear that a satisfactory level of service security and resilience can be achieved and maintained only if:
 - the service requirements are clearly identified;
 - an acceptable level of service is clearly defined;
 - the fulfilment of security and resilience parameters is continuously monitored;
 - there is coordination between monitoring, incident management and business continuity management processes.

7.1. Recommendations to governments and public bodies

1. Governments are recommended to adopt a staged approach, with the ability of backtracking each stage, because the complexity of the cloud environment introduces a number of

unknown variables that could be very difficult to manage. Public administrators (PAs) at any level should consider system interconnection and interdependencies (most of which may be unknown), especially when simultaneously moving multiple services to a cloud system(s). PAs should consider this caveat in the context of a dynamically changing environment and a currently incomplete understanding of vulnerability and attack mechanisms, and the complexity of related controls. PAs should not assume that the successful deployment of an application in a cloud environment is automatically a positive indication for proceeding with many other deployments; the security and resilience requirements of each application should be examined carefully and individually and compared to the available cloud architectures and security controls.

2. National governments should prepare a strategy on cloud computing that takes into account the implications for security and resilience that such service delivery models will have in the context of their national economies and services to citizens over the next 10 years. The early adopters in each Member State should be seen as possible test beds, but it will be essential to have, at least at a national level, a coherent and harmonized approach to cloud computing in order to avoid: 1) proliferation of incompatible platform and data formats (lack of service interoperability), 2) an inconsistent approach to security and resilience, including an inconsistent and inefficient approach to risk management, and 3) a lack of critical mass.
3. We recommend governments to study the role that cloud computing will play in the context of protecting critical information infrastructures. It is not unrealistic to assume that cloud computing, in all its possible implementations, will serve, in the near future, a significant portion of European Union citizens, small and medium-sized enterprises and public administrations, and therefore the critical infrastructures from which services are provided should be protected as such. In other words, a national strategy for cloud computing should aim to understand and address, among other issues, the effects of national and supra-national interoperability and interdependencies, and assess the impact of possible cascade failures, evaluate the opportunity to introduce an incident reporting scheme for cloud providers similar to one already adopted in the telecommunication sector (in particular we refer to the reporting mechanism introduced in articles 4 and 13 of the newly adopted Telecom Framework Directive²⁰) and to be prepared for crisis management in the event of large-scale incidents of this nature.

²⁰<http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0037:0069:EN:PDF>

4. We recommend national governments and European Union institutions to further investigate the concept of a European Governmental cloud as a supra national virtual space where a consistent and harmonized set of rules could be applied, both in terms of legislation and security policy and where interoperability and standardization could be fostered. Moreover such a European Union wide infrastructure could be used in the context of a pan European mutual aid and assistance plan for emergencies.

In more detail, if public bodies finally decide to move to cloud computing, they should:

- Taking into account national strategy, define their requirements (possibly using as a support those suggested in this report) in order to identify which cloud solution matches their needs. PAs should also consider human factors (eg, security and resilience awareness, resistance to new security policy models) and legal frameworks.
- As a matter of good governance, have in place an information security management process, which includes risk management, a policy for information security and resilience, asset management (physical and informational), etc, based on available good practices.
- PAs should focus on a comprehensive service catalogue and physical and information asset classification; per each service and asset, the appropriate security and resilience requirements should be specified. We reckon that most large institutions will not be able to complete such a task in a reasonable time frame, as we assume that in certain cases they don't have yet a complete picture of their assets. A viable alternative, in the context of the staged approach to cloud computing, would be to start with the definition of macro categories of assets and services (eg non sensitive and non critical, medium sensitive and medium critical etc.) and to elaborate a detailed asset classification according to the simple logic that the first service to be migrated to the cloud should be the first to be classified (prior migration).
- Define acceptable levels of service (a benchmark, eg, availability) for their requirements. They will use the benchmark(s) to measure the performance of their services.
- Identify the set of controls and their degree of specificity in order to reach a minimum acceptable level of data assurance and services resilience.
- Make sure that all the essential security, resilience and legal requirements are detailed in their service level requirements and specified in their service level agreements. These should be drawn up when planning a service migration. For instance, SLAs should include the right to audit, (or at least should include access to independent audit report), the means to recover data and application (ie avoid lock-in), and detail the level of monitoring and reporting, etc.

- Establish a metrics framework (including key goal and performance indicators) to continuously evaluate whether the following are met:
 - service level target(s);
 - the level of preparedness and preventative capability in case of incidents including both faults and malicious attacks;
 - the efficiency and effectiveness of the reaction and recovery phase following a disruptive event.
- Take into account relevant national and international regulations applying to third parties (eg, electronic digital signature directives, ISO third-party assurances) in order to ensure the trustworthiness of the communications between all the parties involved in the provision of the service (PAs, citizens, services provider-business parties, as well as systems). The authenticity of the identities of the parties and their authorization to perform an action, the point in time (ie, timestamp), and location should be assured.
- Apply, in the identity and access controls processes, the principles of need-to-know, least privilege and segregation of duties.
- Have tools, methodologies and governance structures to, for example, assure due diligence.
- Verify the financial stability and solvency of business partner(s), including specific relevant lines of business in order to avoid unexpected interruptions to services or lock in.
- Ensure that satisfactory telecommunication connections, critical dependencies (eg, electricity), processing power and storage capacity are guaranteed and maintained.
- Check the priority for the resumption of third-party communications and cloud services in the event of a disruption.
- Test the business continuity plan along the whole services supply chain.
- For highly critical applications, plan for cloud service unavailability. There should be a mechanism in place to make access to IT services possible even when the connection to the cloud(s) is not available.

Finally, cloud providers and independent service vendors should consider the recommendations included in this report as a possible source of information when aligning their business offers and values proposition with users' needs and requirements.

8. Glossary

AAA	Authentication, authorization and accounting
AD	Active directory
API	Application programming interface - specification of an interface published by a software supplier
ARP	Address Resolution Protocol (2)
Apps	In this report used as an abbreviation for Applications
Asset	The target of protection in a security analysis
Availability	The proportion of time for which a system can perform its function
BSDG	Bundesdatenschutzgesetz (German Data Protection Act)
BS	British Standard
CA	Certification authority
	Common Assurance Maturity Model
CC	Common Criteria
CEO	Chief executive officer
Classified information	Information that is labelled in a government or business classification system for confidentiality. A typical classification system consists of several levels: <i>unclassified</i> , <i>restricted</i> , <i>confidential</i> , <i>secret</i> or <i>top secret</i> . Classified information usually means 'restricted' or higher.
Confidentiality	Ensuring that information is accessible only to those authorized to have access (ISO 17799)
Co-residence	Sharing of hardware or software resources by cloud customers

CISO	Chief information security officer
CP	Cloud provider
CPU	Central processing unit
CRL	Certificate revocation list
CRM	Customer relationship management
CSO	Chief security officer
CSP	Cloud service provider
CTO	Chief technology officer
Data controller	The 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; where the purposes and means of processing are determined by national or Community laws or regulations, the controller or the specific criteria for his nomination may be designated by national or Community law.
Data processor	A natural or legal person, public authority, agency or any other body which processes personal data on behalf of the controller
Data subject	Identified or identifiable natural person (see EU Directive 95/46/EC) from whom data is collected and/or about whom that data is processed
DDoS	Distributed denial of service
De-provision	The process of enforcing the removal of a resource from use, or disallowing its use by a set of users
DNIE	Documento Nacional de Identidad Electronico
DPD	Data Protection Directive

Edge network	In this context, a network of computers which is able to process and store data for delivery close to its final destination
EEA	European Economic Area
EAP	Electronic Administrative Procedure
EDoS	Economic denial of service
Escrow	The storage of a resource by a third party which has access to that resource when certain well-defined conditions are satisfied
FIM	Federated identity management
Guest OS	An OS under the control of the cloud customer, running in a virtualised environment
Host OS	The operating system of the cloud provider which runs multiple guest Oss
HSM	Hardware security module
Https	Http connection using TLS or SSL
Hypervisor	Computer software or hardware platform virtualization software that allows multiple operating systems to run on a host computer concurrently
IaaS	Infrastructure as a Service (cloud architecture)
IDS	Intrusion detection system
Integrity	The property that data has not been maliciously or accidentally altered during storage or transmission
IP	Internet Protocol
IPS	Intrusion protection system
ISO	International Organization for Standardization
ISV	Independent Software Vendors

ITIL	The Information Technology Infrastructure Library (ITIL) is a set of concepts and practices for information technology services management (ITSM), information technology (IT) development and IT operations.
Jitter	The variation in the time between packets arriving, caused by network congestion, timing drift, or route changes
Latency	Time taken for a packet of data to get from one designated point to another
LAN	Local area network
LDAP	Lightweight Directory Access Protocol
Linkability	Linkability describes the extent to which a given data set allows one to establish identity between two or more pseudonyms.
LHA	Local Healthcare Authority
LMS	Leave Management System
MAC	Media Access Control (address of a network node in IP)
MAC (2)	Mandatory Access Control
MITM	Man in the middle (a form of attack)
MSS	Managed security services
NIS	Network and information security
NIST	<u>National Institute of Standards and Technology (USA)</u>
NRA	National regulatory authority (for telecommunications)
Non-repudiation	<u>The property whereby a party in a dispute cannot repudiate or refute the validity of a statement or contract</u>
OCSP	<u>Online Certificate Status Protocol</u>
OS	Operating system

OTP	One-time password (type of authentication token)
OVF	Open virtualisation format
PA	Public Authority
Perimeterisation	The control of access to an asset or group of assets
Port scan	Probing a network host to determine which ports are open and what services they offer
Protection profile	A document specifying security evaluation criteria to substantiate vendors' claims for a given family of information system products (a term used in Common Criteria)
Provision	The issuing of a resource
PV LAN	Private VLAN
QoS	Quality of service
RBAC	Role-based access control
Reliability	The extent to which any computer-related component consistently performs according to its specifications
Resilience	The ability of a system to provide and maintain an acceptable level of service in the face of faults (unintentional, intentional, or naturally caused)
ROI	Return on investment
ROSI	Return on security investment
RPO	Recovery point objective
RTO	Recovery time objective
RTSM	Real-time security monitoring
SaaS	Software as a Service (cloud architecture)

Security target	A document specifying criteria for the evaluation of security in order to substantiate a vendor's claims for the security properties of a product (a term used in Common Criteria)
Service engine	The system responsible for delivering cloud services
Side channel attack	Any attack based on information gained from the physical implementation of a system; eg, timing information, power consumption, electromagnetic leaks or even sound can provide an extra source of information which can be exploited to break the system.
SLA	Service level agreement
SSL	Secure Sockets Layer (used for encrypting traffic between web servers and browsers)
Subpoena	In this context, a legal authority to confiscate evidence
SWOT	Strengths, weaknesses, opportunities, threats
TFEU	Treaty on the Functioning of the European Union
Throughput	The amount of data transferred in one direction over a link divided by the time taken to transfer it, usually expressed in bits or bytes per second.
TLS	Transport Layer Security (used for encrypting traffic between web servers and browsers)
Tolerance	The ability of a <i>system</i> to respond gracefully to an unexpected hardware or software failure.
ToU	Terms of use
UPS	Uninterruptable power supply
VAS	Value-added services
VLAN	Virtual local area network

VM	Virtual machine
VPC	Virtual private cloud
VPN	Virtual private network
Vulnerability	Any circumstance or event that has the potential to adversely impact an asset through unauthorized access, destruction, disclosure, modification of data, and/or denial of service
WAN	Wide area network
XML	Extensible Mark-up Language

9. References

1. **European Commission.** [Online] 2010.
<http://europa.eu/rapid/pressReleasesAction.do?reference=IP/10/1602&format=HTML&aged=0&language=EN&guiLanguage=fr>.
2. **CAMM - Common Assurance Maturity Model.** [Online] <http://common-assurance.com/>.
3. **Official Journal of the European Union.** [Online] 2010. <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0037:0069:EN:PDF>.
4. **IDA.** [Online]
<https://www.ida.gov.sg/News%20and%20Events/20050907165443.aspx?getPagetype=21>.
5. **Etro, Federico.** [Online] 2009. <http://www.intertic.org/Policy%20Papers/RBE.pdf>.
6. **European Commission.** [Online] 2010.
http://europa.eu/press_room/pdf/complet_en_barroso___007_-_europe_2020_-_en_version.pdf.
7. **Ministerial declaration on eGovernment (Malmö, Sweden).** [Online] 2009.
<http://www.tecnimap.es/userfiles/ministerial-declaration-on-egovernment.pdf>.
8. **European Commission.** [Online] 2010. <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:52010DC0245%2801%29:EN:NOT>.
9. **ENISA.** [Online] 2007-2010. <http://www.enisa.europa.eu/act/res/technologies/tech/tech>.
10. **The Telegraph.** [Online] 2010. <http://www.telegraph.co.uk/technology/news/8186376/Cloud-computing-could-save-EU-economies-645bn-over-next-five-years.html>.
11. **O'Reilly.** [Online] 2010. <http://radar.oreilly.com/2010/06/randi-levin-on-cost-saving-thr.html>.
12. **Wikipedia.** [Online] http://en.wikipedia.org/wiki/Information_security.
13. **Office of Government Commerce.** ITIL - Service Operation. 2007.
14. **Institute of Electrical and Electronics Engineers (IEEE).** [Online]
http://www.ieee.org/education_careers/education/standards/standards_glossary.html.
15. **ENISA.** [Online] 2009. <http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment>.

16. **Article 29 Data Protection Working Party.** [Online] 2010.
http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2010/wp169_en.pdf.
17. **Council of Europe .** [Online] 2010.
http://www.coe.int/t/dghl/cooperation/economiccrime/cybercrime/Documents/Reports-Presentations/2079_reps_IF10_yvespoullet1b.pdf.
18. **Paolo, Balboni.** [Online] 2010. http://common-assurance.com/wp-content/uploads/P_Balboni_Security-and-Privacy.
19. **European Commission .** [Online] 2010.
http://ec.europa.eu/justice/news/consulting_public/0006/com_2010_609_en.pdf.
20. **NIST.** [Online] <http://csrc.nist.gov/groups/SNS/cloud-computing/>.
21. **European Commission.** [Online] 2010. <http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-report-final.pdf>.
22. **ISO.** [Online] 2005.
23. **HIMSS.** [Online] <http://www.himss.org/content/files/EHRAAttributes.pdf>.
24. **Wikipedia.** [Online] http://en.wikipedia.org/wiki/Human_resource_management_system.
25. **Kroes, Neelie.** [Online] 2010.
<http://europa.eu/rapid/pressReleasesAction.do?reference=SPEECH/10/686&format=HTML&aged=0&language=EN&guiLanguage=en>.
26. **NIST.** [Online] 2010. <http://www.cio.gov/pages-nonnews.cfm/page/Federal-Risk-and-Authorization-Management-Program-FedRAMP>.
27. **ENISA.** <http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment>.
[Online] 2009. <http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment>.
28. **European Commission.** [Online] 2010.
<http://europa.eu/rapid/pressReleasesAction.do?reference=IP/10/1602&format=HTML&aged=0&language=EN&guiLanguage=fr>.

Annex I – Legal analysis

Note of caution: the analysis included in this legal annex is done with limited resources, in the context of a report focused on information security and resilience, and it is not meant to be exhaustive, not definitive. It aims to highlight the most relevant legal compliance issues and to provide suggestions for addressing them.

Methodology

The methodology we use for the legal analysis is as follows.

STEP 1: We first aim at answering these six fundamental questions:

- Which services (**‘Services’**) identified in the scenario is the PA/GOV considering migrating to the cloud?
- Are there specific laws or regulations that apply to the Services and what are the relevant duties or obligations imposed upon the PA/GOV (eg, data retention, data protection, interoperability, medical file management, disclosure to authorities, etc)?
- What is the nature of the data or information that would be processed with these Services?
- What are the specific legal provisions that apply to the types of data or information that will be processed and what are the relevant duties or obligations imposed upon the PA/GOV (eg, data protection, intellectual property, confidentiality, security, etc)?
- What is the data or information flow (internal²¹ and external²²) during the operation of these Services?
- Who are the subjects (natural and/or legal persons) involved in the operation of the Services and what are their roles (responsibilities, duties, obligations, and liabilities)?

The questions above will be answered with respect to each specific scenario we are dealing with.

We will mainly take into consideration EU law. Where the relevant laws or regulations have not been harmonized at the European level, we will point out the issues and provide a few examples of the applicable laws of Member States.

STEP 2: Once the questions above have been answered, we will be able to:

²¹ Within the PA/GOV

²² From one PA/GOV to another PA/GOV and/or from the PA/GOV to citizens

- make a list of applicable laws and regulations and the relevant duties and obligations of the PA/GOV; and
- identify the legal issues and the related legal risks.

STEP 3: We will then analyse the impact on legal issues and related legal risks will have for a PA/GOV (and more generally for all the stakeholders) migrating the Services to the cloud. More specifically, we will identify the pros and cons and the benefits and risks of migrating the Services to the cloud. Concerning the risks, we will suggest how to deal with them. In order to improve the migration of the Services to the cloud, we will conclude by issuing recommendations on solutions and/or workarounds for the competent Authorities.

Main regulatory concerns

See [paragraph 3.3.](#)

Regional e-health cloud – scenario No 1

E-health services represent a sector in which the governments of Member State face significant challenges. In fact, on the one hand, high quality and high performances are to be delivered while, on the other hand, there is increasing pressure to cut public expenditure. In other words, there is a need to put into practice the recurrent mantra: ‘Do more with less’. In this context, the introduction of new technologies, new business models and services (eg, unique patient identifier, EHR, EHF, online scheduling of reservations for health examinations, online provision to patients of the related examination records) can provide solutions to these challenges.

E-health is a fast-growing market for providers of related services. Through its e-Health Action Plan – issued in 2004 – the European Commission has boosted its growth.²³ Moreover, due to the growth in

²³ European Commission, *e-Health – making healthcare better for European citizens: An action plan for a European e-Health Area*, Communication from the Commission to the Council, the European Parliament, the European Economic and Social Committee and the Committee of the Regions, COM(2004) 356 final, Brussels, 30 Apr 2004. For more on the Commission’s e-health strategy, see *ICT for Health and i2010: Transforming the European healthcare landscape: Towards a strategy for ICT for Health*, Office for Official Publications of the European Communities, Luxembourg, 2006. The ultimate goal is to enable access to the patient’s electronic health record and emergency data from any place in Europe. See also Article 29 WP (WP 131/2007) Working Document on the processing of personal data relating to health in electronic health records (EHR); COM (2008) 414 Proposal for a Directive of the European Parliament and of The Council on the

market opportunities and demand, e-health has been selected as part of the Lead Market Initiative.²⁴ In 2008 the EU Commission issued a recommendation on the cross-border interoperability of electronic health record systems.²⁵ It foresees ‘the adoption of a comprehensive legal framework for interoperable electronic health record systems. Such a legal framework should recognize and address the sensitive nature of personal data concerning health and provide for specific and suitable safeguards so as to protect the fundamental rights to protection of the personal data of the individual concerned.’ Furthermore, it encourages Member States ‘to implement the interoperability of electronic health record systems as an integral part of regional and national e-health strategies.’ The recommendation invites Member States ‘to report, on a yearly basis, to the Commission on the measures they have taken in relation to the implementation of the cross-border interoperability of electronic health record systems.’

The legal framework applicable to the provision of health services in Europe is quite complex. Sources of primary and secondary care levels interact to impose duties and obligation on all the players: public administrations, local healthcare authorities, hospitals, private practices, doctors, administrative staff, etc. The legal framework is not consistent throughout the European Union Member States.²⁶ The reason has to be traced back to the fact that healthcare is a domain which largely remains under the competence of the Member States. Before the enactment of the Treaty of Lisbon (1 December 2009),²⁷ the European Union (and previously the European Community) only had a supporting, coordinating and complementary role in this domain (parallel complementary competence).²⁸ It could use ‘soft law’ instruments (eg, recommendations) in order to coordinate and promote specific actions in this field, yet measures for harmonization were explicitly excluded. The Treaty of Lisbon should open a new phase in EU harmonization in this sector, by clarifying and expanding EU competences in the sphere of public health. Article 4.2(k) of the (Consolidated Version of the) Treaty on the Functioning of the

application of patients' rights in cross-border healthcare; COM(2008) 415 A Community framework on the application of patients' rights in cross-border healthcare.

²⁴ See: <http://ec.europa.eu/information_society/activities/health/policy/lmi_ehealth/index_en.htm>.

²⁵ See: <http://ec.europa.eu/information_society/newsroom/cf/itemlongdetail.cfm?item_id=4224>.

²⁶ See European Commission (2009) Study on the Legal Framework for Interoperable eHealth in Europe, pp 11 et seq; available at: <http://ec.europa.eu/information_society/activities/health/docs/studies/legal-fw-interop/ehealth-legal-fmwk-final-report.pdf>.

²⁷ <http://www.consilium.europa.eu/showPage.aspx?id=1296&lang=en>

²⁸ See, eg, Schutze, *Co-operative federalism constitutionalised: the emergence of complementary competences in the EC legal order*, *European Law Review* 2006, p 179.

European Union²⁹ classifies ‘common safety concerns in public health matters, for the aspects defined in this Treaty’ as shared competences between the Member States and the Union. Furthermore, the EU has a parallel complementary competence for the ‘protection and improvement of human health’ on the basis of article 6. Significant direct and indirect interventions by the EU in this field are as follows:

- The European regulatory framework for medical devices³⁰;
- The European Directive on the transparency of measures related to the pricing and reimbursement of medical products³¹.
- The European Directive on electronic commerce³², which contributes to the functioning of the internal market by ensuring the free movement of information society services, including e-health services, between Member States.
- EU legislation on the free movement of professionals, including health professionals³³.
- The European regulatory framework for personal data protection³⁴ and for the protection of privacy in electronic communications³⁵.
- In June 2008 the Commission finally published a proposal for a directive on patients’ rights in cross-border healthcare³⁶. It is notable that article 16 of the proposed directive specifically

²⁹ Available at: <<http://eur-lex.europa.eu/JOHtml.do?uri=OJ:C:2010:083:SOM:EN:HTML>>

³⁰ Available at: <http://ec.europa.eu/consumers/sectors/medical-devices/regulatory-framework/index_en.htm>

³¹ Directive 89/105/EEC; available at: <http://ec.europa.eu/enterprise/sectors/healthcare/competitiveness/pricing-reimbursement/transparency/index_en.htm>.

³² Directive 2000/31/EC sets forth the information requirements imposed on service providers in the information society, rules on commercial communications, rules regarding contracts concluded by electronic means and the liability of intermediary service providers; available at: <http://ec.europa.eu/internal_market/e-commerce/directive_en.htm>.

³³ Directive 2005/36/EC aims at ensuring that Member States enact uniform, transparent, and non-discriminatory rules recognizing professional qualifications and experience to allow professionals to work temporarily or permanently through the European Union; available at: <http://ec.europa.eu/internal_market/qualifications/future_en.htm>.

³⁴ Directive 95/46/EC Directive on Privacy and Data Protection, http://ec.europa.eu/justice/policies/privacy/index_en.htm

³⁵ Directive 2002/58/EC Directive on Privacy and Electronic Communications, http://eur-lex.europa.eu/smartapi/cgi/sga_doc?smartapi!celexapi!prod!CELEXnumdoc&numdoc=32002L0058&model=quichett&lq=en

relates to e-health and provides for the adoption of 'specific measures necessary for achieving the interoperability of information and communication technology systems in the healthcare field, applicable whenever Member States decide to introduce them. Those measures shall reflect developments in health technologies and medical science and respect the fundamental right to the protection of personal data in accordance with the applicable law. They shall specify in particular the necessary standards and terminologies for the interoperability of relevant information and communication technology systems to ensure safe, high-quality and efficient provision of cross-border health services.' It is also worth mentioning that the proposed article 14, requests the Commission to adopt 'measures enabling a pharmacist or other health professional to verify the authenticity of the prescription and whether the prescription was issued in another Member State by an authorized person through developing a Community prescription template, and supporting interoperability of ePrescriptions'.

The scenario

In order to be able to undertake a thorough analysis of the issues and, given a lack of harmonization in the e-health sector throughout Member States, a local scenario has been drafted. In this respect, consideration has been given to the fact that a number of Italian LHAs are contemplating entering into a joint agreement with a national telecommunication company for the creation of their own cloud. The LHAs plan to migrate, to the cloud, services such as EHR, EHF, the online scheduling of reservations for health examinations, the online provision to patients of related examination records and other, less critical, services, eg, back-end services, human resources, payroll, e-learning.

Data and service availability, data authenticity, integrity, trustworthy information security, resilience of services, protection of personal data and legal compliance (especially regarding data protection legislation) are the main concerns to be investigated.

Data types and flow of data between subjects involved

Both the EHR and the EHF contain several items of information on an individual's health that relate to current and past clinical events (eg, medical findings, hospitalization records, and emergency care) that are aimed at documenting that individual's medical history. The personal data are inter-linked using different computerized tools, which in any case allow the data to be easily retrieved and browsed by the various health care professionals and/or bodies providing medical care to that individual over time.

³⁶ European Commission, *Proposal for a Directive of the European Parliament and the Council on the application of patients' rights in cross-border healthcare*, COM(2008)414 final,
http://ec.europa.eu/health/ph_overview/co_operation/healthcare/docs/COM_en.pdf

More specifically, the EHF is a file set up at a healthcare body that acts as the sole data controller (eg, a hospital or a nursing home) where several health care professionals are employed. Conversely, the EHR is a file set up by pooling the data from different data controllers, which as a rule – though this is not always the case – operate within the same geographical area (eg, a healthcare unit and a private laboratory operating in the same region and/or area). For instance, health files may also make up the set of health care information held by the individual data controllers that participate in an EHR initiative at regional level. Given the above, the EHR and the EHF have to be seen as very interrelated.

‘Online scheduling of reservations for health examinations’ means that it is possible for patients to make appointments by interacting with the LHA online reservation system.

‘Online access to examination records’ means that the patient is enabled to access an ‘examination record’ online; here, examination record means the written record drawn up by a physician on the patient's clinical status following clinical examination and/or test results. In some cases it also allows the downloading of the ‘findings’, ie, the results of the clinical examination and/or test performed on the patient such as an X-ray, an ultrasound scan recording and/or blood tests, along with the written examination record drawn up by the physician.

All the services that the LHAs plan to migrate to the cloud entail the processing not only of personal data but also of sensitive data (a special category of data concerning health, cf article 8 of the Directive 95/46/EC) by different data controllers and data processors – except for back-end services, payroll, and e-learning, for example, where personal data is processed but the processing of sensitive data is less likely. Internal³⁷ and external³⁸ data and information flow during the operation of these Services. A significant number of subjects with different data protection roles are involved and multiple communications and transfers of data occur among them.

Legal issues

Generally speaking, the discipline concerning EHR, EHF, online scheduling of reservations for health examinations and provision to the patients of the related examination records online is set forth in the healthcare laws of Member States, in legislation on patients’ rights,³⁹ and in data protection

³⁷ Within the PA/GOV.

³⁸ From one LHA to another LHA and/or from the LHA to the citizens.

³⁹ See also the European Charter of Patients’ Rights (2002); available at: <www.patienttalk.info/european_charter.pdf>.

regulations, where rules are provided on how healthcare providers must keep and share health records, their contents, archiving, and access right for patients, etc.⁴⁰

The present analysis will focus on data protection and, more generally, on legal compliance. The point of reference will be European and Italian regulations on the matter. We expect that the general line of reasoning and conclusions will be applicable to most Member States. Our aim here is to point out the relevant issues related to the scenario at stake and to offer a method of analysis that can be used for assessments in other Member States.

Data protection law is by far the most relevant legislation. Specifically, the Italian Data Protection Authority issued guidelines on the provision of EHR, EHF,⁴¹ and online examination records.⁴² For an analysis of issues related to 'governmental sovereignty and control over information and data', 'government procurement', 'CSP professional negligence' and 'subcontracting cloud services and change of control of CSP' we refer to the 'main regulatory concerns' as these are general issues (as opposed to issues that are specific to the present scenario). Other issues such as 'data protection and data security', 'interoperability / transfer back / vendor lock-in' will be touched upon, to the extent they are specific to the present scenario and for the rest we will refer to the relevant section in the 'main regulatory concerns'.

Data protection and data security

⁴⁰ 'In many Member States, the responsibility for important areas of the healthcare system has been delegated to the regional level. The decentralization has been implemented, however, in diverse forms and to various extents. [...] Italian regional governments, through their departments of health, are responsible for pursuing the leading national objectives set by the National Health Plan at the regional level. Regional health departments are required to guarantee the benefit package to be delivered to the population through a network of local health units and public and private accredited hospitals. They are responsible for legislative and administrative functions, for planning healthcare activities, for organizing supply in relation to population needs and for monitoring the quality, appropriateness and efficiency of the services provided. The regional level has legislative functions and executive functions as well as technical support and evaluation functions.' European Commission (2009) Study on the Legal Framework for Interoperable eHealth in Europe, p 21; to see how decentralization has been implemented in various Member States, see also pp 21-24; see also pp 18 and 75 et seq; available at: http://ec.europa.eu/information_society/activities/health/docs/studies/legal-fw-interop/ehealth-legal-fmwk-final-report.pdf.

⁴¹ Italian Data Protection Authority (2009) Guidelines on the Electronic Health Record and the Health File (published in Italy's Official Gazette No 178 dated 3 August 2009); available at: <http://www.garanteprivacy.it/garante/doc.jsp?ID=1672821>.

⁴² Italian Data Protection Authority (2009) Guidelines on Online Examination Records (document adopted on 25 June 2009 and submitted for public consultation in accordance with the Notice published in Italy's Official Gazette No 162 dated 15 July 2009); available at: <http://www.garanteprivacy.it/garante/doc.jsp?ID=1634292>.

Here we just deal with the very specific issues related to the present scenario, referring to the introduction for an overview of general data protection and data security issues.

Transfer

Article 29 Working Party, in the (WP 131/2007) Working Document on the processing of personal data relating to health in electronic health records (EHR), has stressed that '[c]onsidering the elevated risk to the personal data in an EHR system in an environment without adequate protection, (...) any processing – especially the storage – of EHR data should take place within jurisdictions applying the EU Data Protection Directive or an adequate data protection legal framework.'⁴³ Moreover, unless the data subject or patient has given his explicit consent (which must be in writing in a number of Member States if it concerns sensitive data) personal data in EHR 'should be transferred to countries outside the European Union/European Economic Area only in anonymized or at least pseudonymized form.'⁴⁴

Data security

An appropriate high level of data security for the complete performance of the system must be in place (article 17 of the Directive 95/46/EC).

Identity management, access control, and data integrity

More precisely, in order to make a system acceptable from a data protection perspective, access by unauthorized persons must be virtually impossible and prevented. Concurrently, availability of the system for authorized professionals must be virtually unlimited where there is a genuine need to know. This is what the Article 29 Working Party explicitly recommends in its WP 131/2007 document.⁴⁵ Furthermore, '[t]he legal framework for setting up an EHR system would have to foresee the requirement for implementing a series of measures of a technical and organisational nature appropriate for avoiding loss or unauthorized alteration, processing and access of data in the EHR system. Integrity of the system must be guaranteed by making use of the knowledge and instruments

⁴³ Article 29 Working Party, in the (WP 131/2007) Working Document on the processing of personal data relating to health in electronic health records (EHR), p 19; available at:

<http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2007/wp131_en.pdf>.

⁴⁴ Id

⁴⁵ Article 29 Working Party, in the (WP 131/2007) Working Document on the processing of personal data relating to health in electronic health records (EHR), p19; available at:

<http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2007/wp131_en.pdf>.

representing the present state-of-the-art in computer science and information technology.”⁴⁶

Moreover, The Working Party pointed out that encryption should not only be used for the transfer but also for the storage of data in EHR systems.⁴⁷

In conclusion, the legal framework concerning security measures should especially foresee the necessity of:

- the development of a reliable and effective system of electronic identification and authentication as well as constantly up-dated registers for checking on the accurate authorization of persons having or requesting access to the EHR system;
- comprehensive logging and documentation of all processing steps which have taken place within the system, especially access requests for reading or for writing, combined with regular internal checks and follow up on correct authorization;
- effective back-up and recovery mechanisms in order to secure the content of the system;
- preventing unauthorized access to or alteration of EHR data at the time of transfer or of back-up storage, eg, by using cryptographic algorithms;
- clear and documented instructions to all authorized personnel on how to properly use EHR systems and how to avoid security risks and breaches;
- a clear distinction of functions and competences concerning the categories of persons in charge of the system or at least involved in the system with a view to liability for shortcomings;
- regular internal and external data protection auditing⁴⁸.

All these very stringent security requirements have to be coherently respected in the cloud environment also, irrespective of whether the CSP is to be considered a controller or a processor. In fact, the obligation to have in place or to guarantee high data security rests on the data controller who will actually transfer it to the data processor by means of a contract or letter of appointment (article 17(2) (3) and (4)).

⁴⁶ *Id*, p 19

⁴⁷ *Id*

⁴⁸ *Id*, pp 19-20

Below are some examples of specific provisions concerning security measures as set forth in either the Italian Data Protection Authority (2009) *Guidelines on the Electronic Health Record and the Health File* (G_EHR) or the Italian Data Protection Authority (2009) *Guidelines on Online Examination Records* (G_OER), which provide for rigorous ID management, access control and data integrity.

Duty/Obligation	Source
Given the sensitiveness of the personal data processed via an EHR/EHF, specific technical arrangements should be made in order to ensure the appropriate security level (section 31 of the DP Code) - without prejudice to the minimum measures data controllers are required to take in any case pursuant to the Code (section 33 et seq.). If data storage/filing systems are used, suitable arrangements should be made to protect the data against unauthorised access and theft and/or loss, in whole or in part, of the storage media and/or fixed/portable processing devices; to that end, encryption technologies might be applied to file systems and/or databases, or other protection measures might be implemented to prevent the data from being intelligible to unauthorised entities. The following measures should also be taken: • suitable authentication and authorisation systems should be applied to the persons in charge for the processing as a function of the respective access/processing requirements (eg, as for browsing, changing and adding records); • procedures should be in place to regularly check quality and consistency of authentication credentials and authorisation profiles applying to the persons in charge for the processing; • criteria should be laid down to encrypt and/or keep separate the data suitable for disclosing health and sex life from any other personal data; • accesses and operations should be logged; • audit logging should be in place to control database accesses and detect abnormalities. As for EHRs, secure communication protocols should be deployed by implementing encryption standards for electronic data communications between the various data controllers.	Part II Section 10 (G_EHR)

The highly sensitive nature of the personal data that are processed in connection with the online access to examination records requires specific technical arrangements to be made to ensure the appropriate security level as per section 31 of the Code; this is without prejudice to the minimum measures that every data controller is required to take in pursuance of the Code (see section 33 et seq.) with particular regard to those set forth in Rule 24 of the Technical Specifications applying to minimum security measures (Annex B to the Code) – whereby the transfer of data suitable for disclosing an individual's genetic identity is only permitted in encrypted format.

Section 6
(G_OER)

Online consultation of examination records via web-based services on the Internet

Where the service to be provided consists in enabling a data subject to access the website of the health care body that has performed the relevant examination in order to download and/or view the respective record(s), specific precautions should be implemented such as the following:

- Secure communication protocols based on encryption standards for electronic data transfers, including digital certification of the systems delivering network-based services (https SSL - Secure Socket Layer - protocols);
- Suitable arrangements to prevent acquisition of the information contained in the electronic file if the latter is stored in local and/or centralised caching systems after being consulted online;
- Suitable authentication systems based either on standard credentials or, preferably, on strong authentication procedures;
- Short-term (maximum 45-day) availability of the online examination record;
- Possibility for the user to prevent online viewing of the relevant examination records and/or delete such records, in whole or in part, from the online access system.

Emailing of the examination record(s)

If the data controller plans to send a copy of the examination record(s) to the data subject's email address based on a specific request by the latter, the following precautions will have to be implemented as regards digital records:

- The examination record(s) will have to be sent as an attachment to the email message rather than as text embedded in the body part of the

message;

- The file containing the examination record(s) will have to be protected so as to prevent unlawful and/or unwanted acquisition of the information by entities other than the relevant addressee(s). To that end, the file may be password-protected, or else an encryption key may be applied and notified to data subjects via different communication channels (see Rule 24 of the Technical Specifications – Annex B to the Code). This requirement may fail to be met if the data subject expressly requests to do so, after being duly informed, since the sending of examination records to the email address specified by the data subject does not give rise to a transfer of medical data between two data controllers and consists actually in the communication of data between the health care provider and the data subject at the latter's request;
- The email addresses will have to be validated by means of an ad-hoc online checking procedure to prevent sending electronic documents – albeit protected by encryption – to addressees other than the specific user that has requested them.

The following measures will have to be implemented in all cases with a view to processing data to provide such online services to users

1. Suitable authentication and authorisation systems will have to be deployed in respect of the persons in charge for the processing as a function of their roles and access/processing requirements – eg, by considering whether they may browse, modify and/or supplement the information; biometrics-based strong authentication will have to be implemented if the processed data are suitable for disclosing an individual's genetic identity;

2. The data suitable for disclosing health and sex life will have to be kept physically and logically separated from any other personal data that is processed for administrative and/or accounting purposes.

Furthermore, the data controller should envisage ad-hoc procedures to immediately disable the online consultation and/or terminate the emailing of examination records related to a data subject that has notified the theft and/or loss of his/her own authentication credentials, or else any other circumstances that may endanger the confidentiality of the respective personal data.

In any case, all the security measures required to comply with the ban on dissemination of medical data set forth in the Code should be implemented (see

sections 22(8) and 26(5) of the Code).

Data subject or patients' rights

As mentioned in the introduction, the controller has an obligation to guarantee the data subject's right of access to data as described in article 12 of the Directive 95/46/EC. However, when health-related personal data is being processed, specific restrictions on a patient or data subject's right of access, on the one side, have to be combined with specific rights of access to health records, which are established in national provisions on patients' rights, on the other side. In this specific case, it is not only that the Directive 95/46/EC has been implemented in an inconsistent way but also that the rights of patients have been defined and implemented differently under various national laws.

Interoperability / Transfer back / 'Vendor lock-in'

Interoperability of electronic health record systems is a necessary condition that has been laid down in both the 2008 EU Commission Recommendation on cross-border interoperability of electronic health record systems⁴⁹ and the 2008 EU Commission proposal for a directive on patients' rights in cross-border healthcare.⁵⁰

Transfer back and vendor lock in have to be taken into consideration in the e-health sector as they represent serious threats to the continuity of the service. In fact any (temporary) unavailability and/or inefficiency of the services will lead to significant liability for the e-health providers (ie, LHAs in this specific case).

Final considerations

The main issues related to the migration to the cloud of services such as EHR, EHF, the online scheduling of reservations for health examinations, and the online provisioning to patients of related examination records are to be identified in:

- (i) patient data transfer;
- (ii) patient data security; and
- (iii) interoperability.

⁴⁹ See: <http://ec.europa.eu/information_society/newsroom/cf/itemlongdetail.cfm?item_id=4224>.

⁵⁰ European Commission, *Proposal for a Directive of the European Parliament and the Council on the application of patients' rights in cross-border healthcare*, COM(2008)414 final;
http://ec.europa.eu/health/ph_overview/co_operation/healthcare/docs/COM_en.pdf

Migrating services such as EHR, EHF, the online scheduling of reservations for health examinations, and the online provisioning to patients of related examination records to the cloud has the potential to offer security and advantages in interoperability. In fact, it is most likely that solid CSPs have more competent, qualified personnel, and greater financial capabilities than any LHA, enabling these CSPs to assure the highest security standards and to foster the global interoperability of these services. However, it can be extremely difficult for CSPs to be able to offer services that fulfil all the stringent and often non-harmonized regulatory requirements mentioned in our analysis. LHAs should look for suitable offers from CSPs, and obtain the necessary regulatory guarantees by carefully negotiating the relevant contracts. Alternatively, LHAs concerned with inadequate regulatory compliance could start familiarizing themselves with cloud technologies by migrating less critical services, i.e., back-end services, payroll, e-learning, and HR (although one should bear in mind that sensitive data will most likely be processed when HR services are being provided).

Moving forward, regulations should be changed in order to introduce greater clarity and consistency into the EU regulatory framework for the protection of personal data, including patient data, whilst creating workable regulatory conditions for CSPs, and thus reap the full benefits of cloud computing as applied to e-health services. In this respect:

- A strong effort should be made to harmonize data protection laws in EU Member States.
- Data protection roles (ie, of data controller and data processor) in the cloud computing environment have to be once and for all clarified.
- Self-regulatory initiatives, such as codes of conduct or codes of practice like the ‘binding corporate rules’ for international data transfers, should be streamlined so that they fully guarantee patients’ rights in global cross-border healthcare.
- Clear homogeneous and appropriate security measures should be set forth throughout the European Union, so that such measures may be embedded – to the maximum extent possible – in cloud services that aim for the so-called ‘privacy by design’ (19)⁵¹. On the specific concept of “appropriate security measures” we recommend Art.29 Working Party to provide further clarifications.

⁵¹The principle of ‘privacy by design’ means that privacy and data protection are embedded throughout the entire life-cycle of technologies, from the early design stage to their deployment, use and ultimate disposal. This principle features *inter alia* in the Commission Communication on A Digital Agenda for Europe – COM(2010) 245..

- The rights of patients and data subject should also be harmonized throughout the European Union so that patients can expect and enjoy consistent protection. We believe that in this way a CSP will be in a position to prepare offers for Healthcare services more easily in line with the necessary legal and regulatory requirements
- CSPs should clearly explain how LHAs (and, more generally, their clients) can migrate to another CSP (avoiding risk of 'vendor lock-in' for LHAs) and thus also assure continuity of service from the LHAs during the transfer back and the migration of information and data. CSPs should also be encouraged to offer service interoperability.

Governmental IT-services in municipal public administrations – scenario No 2

Governments are increasingly using IT to conduct their administrative tasks. Because of its promising advantages, public administrations are considering the migration of their services into cloud computing infrastructures. The following analysis will focus on the specific legal issues arising when offering cloud-based services, and will only address those that are particular to governmentally-run clouds. What rules and regulations have to be observed? What duties and obligations have to be fulfilled? How high is the risk of liability? The answers to these questions depend vastly on the nature of services that are being offered in the cloud.

The Scenario

The following analysis will focus on a scenario in which provinces or regions set up and offer a cloud as a service to municipalities (private or community cloud). For the purposes of the legal analysis, four parties (subjects) can be isolated: the region or province that is the cloud service provider, the third party managing the cloud, the municipalities using it as tenants and the citizens.

The region or province owns the cloud infrastructure, the management of which it has outsourced to a private entity. This third party will be procured by the government, and will be entrusted with the set-up and operation of the cloud infrastructure. The services that are being offered to municipalities are either based on a PaaS or a SaaS delivery model. In particular, these services are:

- Electronic request management: allowing citizens to request electronically from home a subsidy, aid, license, news of the status of their requests, etc.
- Municipal management platform/back office: accounting, human resources, etc. This could include specifically services for invoices, citizen databases (such as criminal records) and various automated reports.
- Online payment platform: for paying taxes, fines, etc.

Data types and flow of data between subjects involved

Because personal data is involved, the respective national rules on data protection need to be observed.⁵² In this scenario, the nature of the data can range from seemingly unimportant personal details to highly sensitive information, such as criminal history and records relating to the suspension

⁵² As set forth by the Data Protection Directive 95/46/EC

of voting and licenses. Data sensitivity will have to be taken into account for the services that are being run in the cloud, and also for the cloud infrastructure itself.

Due to the nature of a so-called private cloud and community cloud, meaning that the infrastructure operates under the control of the region, province or by a community of these entities, under normal circumstances there will be no unsolicited data flow. Data is handled solely by the region or province that is operating the cloud, the municipality and by the citizen (via his electronic client).

Legal issues

Applicable law: Data Protection Directive

Directive 95/46/EC does not differentiate between data processing by private or by public entities, which means the rules and regulations set forth in Directive 95/46/EC will have to be observed by regions or provinces and municipalities just as must be observed by private entities. Great attention has to be given to the fact that sensitive data may be processed in the cloud. This triggers the rules in article 8 of Directive 95/46/EC, which contains specific provisions on sensitive data. If the specific service processes data that reveals racial and ethnic origin, political opinions, religious and philosophical beliefs, trade-union membership and data concerning health or sex life, in order to operate it will have to fulfil one of the exceptions enumerated in article 8(2) as enacted in national law.

Even though it is not applicable to this scenario, it should be noted that pursuant to article 13(1), Directive 95/46/EC, Member States may restrict the applicability of Directive 95/46/EC for matters of national and public security, or the prosecution and prevention of crime. Thus, depending on the local law in a Member State, certain types of data handled by the municipalities may not be subject to regulation under Directive 95/46/EC at all. In this scenario, no such data is being processed.

The Data Protection Directive assigns various duties and obligations depending on an entity's role in handling personal data. It distinguishes between data controller and data processor (cf article 2(d) (e), Directive 95/46/EC). Because the municipalities determine, as cloud tenants, the specific purposes and means of data processing, they are data controllers. As sole provider of infrastructure services, working on behalf of its customer, the region or province is the data processor. The third-party contractor itself does not handle data. However, he will have to be part of the data security concept (see *supra*).

The data controller is responsible for compliance with data protection legislation as implemented in national law. He may be held responsible if the processor does not comply with the rules (article 17(2) and (3), Directive 95/46/EC). In order to mitigate this risk, the controller has to request specific assurances from the processor. More specifically, rules and guidelines about data handling have to be put in place. This can be arranged in the form of contracts or agreements or a legislative act between

the region or province and the municipalities, as required by article 17(3), Directive 95/46/EC. These agreements should take into account the sensitivity of the data depending on the specific service that is going to be offered. Guidelines on terms and conditions between these parties will be made below (under legal recommendations). It should also be noted that the controller and processor will be responsible for the selection and supervision of a reliable third-party contractor who manages the cloud. This has to be part of the security concept for the governmental cloud service.

Pursuant to article 20 and depending on national law, checking may be necessary prior to processing. This depends on the type of service and types of data being processed.

Applicable law: government procurement

Because a private third party will be contracted to set up and maintain the cloud infrastructure, the extensive EU regulations on public procurement will have to be observed.⁵³ In this regard, there will be no significant difference to procurement in other areas of governance, so that provinces, regions and municipalities will be able to apply their existing knowledge and experience with the applicable laws and regulations.

Applicable law: contracts

In this scenario, the municipality enters into a contract with the region or province which in turn employs a private third-party to manage the cloud. The legal requirements on the regions or provinces will have to be reflected in the contractual chain. This applies especially to the requirements on data controlling and processing discussed above. Specific recommendations addressing contractual terms that can reflect and assure compliance with these requirements will be included below.

Applicable law: civil and criminal procedural law

Cloud service operators need to be aware of the fact that data being stored in the cloud may be requested as evidence in civil and criminal proceedings. With the existence of multi-party relationships in this scenario, there are multiple subjects for subpoenas or discovery requests that could possibly transfer evidence. This raises the issue of the appropriate party to whom such requests should be addressed. Similarly, care will have to be taken that the principle of governmental sovereignty is preserved, meaning that governments remain in control of their data and may surrender it only when required so by law. Placing data under the control of private entities may pose a risk; by law, private entities may be obliged to surrender evidence under their control in certain circumstances (16). Such a case could be one in which a company that is based in another country runs the cloud infrastructure

⁵³ http://ec.europa.eu/internal_market/publicprocurement/legislation_en.htm

on which governmental data is hosted. Through its extra-national affiliations, the entity would be within reach of courts of foreign country. For example, such issues could arise as a part of the “pre-trial discovery” process that is known to Anglo-American jurisdictions.⁵⁴

The Article 29 working party has produced a document⁵⁵ that addressed the practical issues raised by, and handling of, such request.

Final considerations

For governments and PAs in general, one of the main legal issues is sovereignty and control over the data that is being handled. A governmental body that is by law entitled to handle the data retains responsibility for its proper handling and should ensure that its obligations to protect the data extend by contract to its third party providers. Where cloud infrastructure hosting extends beyond the local legal jurisdiction, the public body must consider the implications and related safeguards offered by their provider(s). If governmental data is being handled by private parties in foreign jurisdictions, this creates the risk that foreign courts subpoena the private entity and thus reach into the government’s data. A government body therefore should ensure that its outsourcing providers impose adequate security measures, and that procedures and mechanisms are in place so that only relevant data would ever be surrendered in response to legitimate demands by the judicial authorities.⁵⁶

The legal compliance by all parties involved in a governmental cloud has to be implemented through their contractual relationships. In practice, they either negotiate clauses that assure compliance or choose to contract only with partners whose standard terms and conditions include the required assurances. All stages of the contractual chain between the municipalities, province or region and the external service provider have to be carefully negotiated and/or evaluated. This is especially relevant where data protection law is applicable, because EU and national data protection laws require IT-security measures. This may transform into, for example, service level agreements and provisions on technical and organizational measures for IT-security. In the following table, recommendations regarding data security will be made that flow directly or indirectly from Directive 95/46/EC. Additional recommendations and guidelines on cloud service contracts can be found in a previous ENISA study on

⁵⁴ For details, see Geercken/Holden/Rath/Surguy/Stretton, *Computer und Recht International* 2010, pp 65.

⁵⁵ http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2009/wp158_en.pdf

⁵⁶ This includes checking whether the evidence is rightfully requested (by subpoena or during discovery). See Article 29 Group Working Document 1/2009 on pre-trial discovery for cross-border civil

litigation, adopted on 11 February 2009, WP 158; available at:
http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2009/wp158_en.pdf

cloud computing⁵⁷ (15). Note that national law or governmental policy may require additional, specific assurances for data that is being managed by a third (private) party.

Cloud specific recommendations regarding compliance with the data protection directive

Duty/Obligation	Source	Cloud specific
Spontaneous, unannounced security auditability	Articles 16, 17(1)	No
Transparency to data subjects regarding parties involved	Articles 10 – 12	No
Transparency to data subjects regarding all steps in data processing ('data flow')	Articles 10 – 12	No
Data breach and security incident notification	§ 42a BDSG (Germany)	No
Risk-adapted security policies	Articles 17(1)	No

The column 'Cloud specific' indicates whether this requirement is specific to cloud computing or applies to IT outsourcing generally.

⁵⁷http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1662374.

Governmentally provided cloud infrastructure – scenario No 3

Governments may not only choose cloud computing as a means to fulfil their own IT-related functions but may also choose to operate cloud infrastructures as a service for their citizens – the subject of this scenario. This latter service is aimed mostly at small and medium enterprises which may rent cloud infrastructure from the government on a cloud delivery model. These firms can then use it themselves to run and offer Software as a Service (SaaS). This means that private business will be conducted on governmental cloud ‘premises’, which raises several legal issues that are specific to this type of scenario. However, all of the traditional open issues surrounding cloud computing still apply.⁵⁸ The following analysis will therefore focus on the legal issues arising specifically when offering cloud-based services, and will only address some of the most dominant issues that are particular to governmentally-run clouds.

Data types and flow of data between subjects involved

The cloud model of ‘Infrastructure as a Service’ leaves particularly open the types of data that are involved and how data is being processed and transmitted. As a result the type of personal data being handled cannot be pre-determined. In this cloud model, there is no factual limit to the data flow between the parties involved, and to and from third parties. Control over the data is held by cloud tenants and their customers, and there is a potential for personal data, sensitive data, confidential information (eg, know-how), and intellectual property to circulate in a governmental-run cloud.

Legal issues

The present scenario is similar to a quasi-typical B2B cloud environment. Therefore, the main data issues related to protection already pointed out in the ENISA (2009) *Cloud Computing Risk Assessment*⁵⁹ and the ones highlighted above in the relevant section of “Main regulatory concerns” – especially concerning ‘Confidentiality and intellectual property’, ‘Data Controller-Data Processor’, ‘Appropriate technical and organizational data security measures’, ‘Data transfer to countries outside the EEA’, ‘Data subject’s right of access to data’, Interoperability / Transfer back / “Vendor lock-in” provisions; ‘CSP professional negligence’ – apply *mutatis mutandis* here.⁶⁰

⁵⁸ ENISA (2009) *Cloud Computing Risk Assessment*, pp. 97ss. Available at:
<http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport>.

⁵⁹ *Id.*

⁶⁰ For an analysis of such issues in the B2B cloud environment, see also Council of Europe Discussion Paper (2010) *Cloud computing and its implications on data protection*. Available at:

Applicable law: Treaty on the Functioning of the European Union

When a government located within the EU decides to support cloud computing, it needs to consider the rules set forth in Art. 107 – 109 TFEU.⁶¹ These provisions proscribe governmental aids that distort or threaten to distort competition by favouring certain undertakings. A governmentally-run cloud infrastructure could violate these rules in three regards: Firstly, it may distort competition by favouring cloud computing over other, traditional means of IT outsourcing. Secondly, competition may be affected if infrastructure is only rented to national customers and finally, it may be anti-competitive to other, private cloud providers.

Applicable law: E-Commerce Directive

Because they may be considered information society services,⁶² government clouds are subject to the rules contained in the E-Commerce Directive.⁶³ In the extent they act as hosting providers, Art. 14 of the E-Commerce directive shield them from liability for the information stored at the request of a recipient of the service. In this scenario, the recipient is the cloud tenant who in turn sets up SaaS offerings. Art. 14 (1) (a) of the Directive provides that the government cannot be held liable to the extent they do not have actual knowledge of illegal activity or information and, as regards claims for damages, are not aware of facts or circumstances from which the illegal activity or information is apparent. Art. 14 (1) (b) requires the provider to remove or disable access to the information upon obtaining such knowledge or awareness. Furthermore, Art. 15 clarifies that there is no general obligation to monitor the information that is transmitted or stored. However, hosted content may nonetheless pose a resilience risk on the government. Recital 45 of 2000/31/EG states that limitations of liability do not affect the possibility of injunctions. In practical terms, this means there is a risk that an entire cloud service may be taken down by means of a court injunction.

Applicable law: contracts

<http://www.coe.int/t/dghl/cooperation/economiccrime/cybercrime/Documents/Reports-Presentations/2079_reps_IF10_yvespouillet1b.pdf>; Balboni, P. (2010) *Security and Privacy in Cloud Computing: The European Regulatory Approach*, Executive Action Report, No.335, The Conference Board, October 2010. Available at: <http://common-assurance.com/wp-content/uploads/P_Balboni_Security-and-Privacy-in-Cloud-Computing.-The-European-Regulatory-Approach.pdf>.

⁶¹ Treaty on the functioning of the European Union, Notice No. 2010/C 83/01, pp. 91

⁶² As defined by Article 1(2) of Directive 98/34/EC, as amended by Directive 98/48/EC.

⁶³ See Article 2(a) of 2000/31/EG.

The relationships between the government and other parties involved are two-fold. On one side is the procurement contract to the private outsourcer running and managing the cloud. On the other side is a contractual chain from the government over the software provider to the consumers. Contractual issues in the contract between government and outsourcer revolve around the assurance of data protection compliance and government procurement (see above). This contractual chain however can fulfil a very specific role: They are the only means by which the government can control the fulfilment of certain duties and responsibilities for what is run in the cloud, and to shield itself from potential liability for it.

Final considerations

Limitations on types of data and data flow

As the cloud owner, the government has duties, responsibilities and obligations imposed by law and by contract. Within the contractual chain that reaches down from the government over the SaaS vendors to their customers, specific provisions may ensure that the government complies with all relevant laws. In practice, this may be by way of standard terms and conditions that are made mandatory and non-negotiable for all contractual partners.

Treaty on the Functioning of the European Union

Generally speaking, the rules contained in articles 107 – 109 TFEU apply only if governmental aid selectively reduces economic or other burdens.⁶⁴ In this scenario, that would be the case if the governmental cloud were not to operate at below-market prices. If it operates at market prices, its influence on the market is no different to that of a private cloud infrastructure offered within the EU. To exclude a violation of article 107 TFEU, the governmental cloud service should not be priced lower than the service offered by its commercial competitors in the EU. Non-economic incentives over private competitors, such as strict adherence to data protection rules and regulations, cannot be deemed as threatening market distortion, because legal compliance can be expected from any company.

Prices that are lower than competitors may nonetheless be compatible with the internal market under article 107(3)(c) TFEU. This provision permits aid in order to facilitate the development of certain economic activities or areas. However, the aid must not adversely affect trading conditions to an extent that is contrary to the common interest. Pursuant to article 108 TFEU, it is up to the European

⁶⁴ ECJ, Order of 18 February 1960, *De gezamenlijke Steenkolenmijnen in Limburg / ECSC High Authority* (30/59, ECR 1961 p. 48) (FR1961/00091 NL1961/00093 DE1961/00099 IT1961/00089 EN1961/00048 ES1961-1963/00049), available at: http://eur-lex.europa.eu/smartapi/cgi/sga_doc?smartapi!celexplus!prod!CELEXnumdoc&lg=en&numdoc=6195900030

Commission to decide whether governments negatively affect trade by providing aid to SaaS vendors in the form of non-competitive pricing.

E-Commerce Directive

There is no general obligation to monitor the information that is transmitted or stored in a cloud hosting environment. However, article 14 may offer providers acting as hosts a safe harbour from liability for illegal content where, among other conditions, they remove or disable access to the information upon obtaining knowledge or awareness that it is illegal or infringing material. The government should ensure that a notice-and-takedown procedure and system is in place in case illegal or infringing material needs to be removed from the cloud. Precautions need to be taken by those service providers who wish to avail themselves of the E-Commerce Directive safe harbour, by designing cloud services in a way, for example, that particular cloud tenants can be taken down separately in compliance with a court order.

Annex II – Scenarios

Healthcare scenario – scenario No 1

It is Europe, year 2011, and local healthcare authorities need to implement a new service for the citizen: electronic health records.

An electronic health record (EHR) is an electronic record of patient health information generated by one or more encounters in any care-delivery setting. Included in this information are patient demographics, progress notes, problems, medications, vital signs, past medical history, immunizations, laboratory data and radiology reports.

- The EHR is a secure, real-time, point-of-care, patient centric information resource for clinicians.
- The EHR aids decision-making by clinicians by providing them with access to patient health record information where and when they need it and by incorporating evidence-based decision support.
- The EHR automates and streamlines the clinician's workflow, closing loops in communication and response that result in delays or gaps in care.
- The EHR also supports the collection of data for uses other than direct clinical care, such as billing, quality management, outcomes reporting, resource planning, and public health disease surveillance and reporting.

In order to fulfil European recommendations⁶⁵ and national requirements and to exploit the full value of e-health services, interoperability between different local and national electronic health records has to be guaranteed. For instance the Commission issued an e-health action plan in 2004 and, in July 2008, a recommendation on cross-border interoperability of EHR systems (<http://www.semantichhealth.org/PUBLIC/20080516/P01-03->

⁶⁵ European Commission e-Health – making healthcare better for European citizens: An action plan for a European e-Health Area Communication from the Commission to the Council, the European Parliament, the European Economic and Social Committee and the Committee of the Regions, COM(2004) 356 final, Brussels, 30 Apr 2004. For more on the EC's e-health strategy, see ICT for Health and i2010: Transforming the European healthcare landscape: Towards a strategy for ICT for Health, Office for Official Publications of the European Communities, Luxembourg, 2006. The ultimate goal is to enable access to the patient's electronic health record and emergency data from any place in Europe.

[Semantic%20WS%20Gerard%20Comyn.pdf](#)) so that doctors can gain access to vital information on patients from other Member States whom they happen to be treating.

Cloud scenario

Given the strong focus on interoperability and the potential positive impact on business efficiency by cloud models, a number of local healthcare authorities (LHAs) are considering entering into joint agreement with a national telecommunication company for the creation of their own cloud.

The LHAs plan to migrate, to the cloud, services such as electronic health records⁶⁶ electronic health files (EHF⁶⁷), the online scheduling of reservations for health examinations and other, less critical, services, eg, back-end services, human resources, payroll, and e-learning.

Marco Rossi, CEO of an LHA, is one of the main sponsors of the cloud approach, but he is conscious that representatives of other LHAs in his region are reluctant to move services onto the cloud and therefore he needs to marshal strong arguments, particularly with regard to data and service availability, data authenticity, integrity, trustworthiness in information security, resilience of services, personal data protection and legal compliance (especially with regard to data protection legislation).

He knows that at the next and decisive meeting with the CEOs of all other LHAs who are potentially interested in the 'regional eHealth hybrid cloud', he will have to address the following questions:

- What is the real added-value of the regional e-health hybrid cloud in terms of resilience and reliability?
- Can the regional cloud offer at least the same level of data assurance and security that the LHAs currently have? Can these requirements be captured in agreements on infrastructure service levels between the LHAs and the cloud-infrastructure provider? Which services or information will potentially be more at risk and which of these could be even more secure than they are at present?
- Which type of access can be developed so that the security levels required can be applied?

⁶⁷ "...the Health File is a file set up at a health care body that acts as the sole data controller (eg, a hospital or a nursing home) where several health care professionals are employed. Conversely, the electronic health record is a file set up by pooling the data from different data controllers, which as a rule - though this is not always the case - operate within the same geographical area (eg, a health care unit and a private laboratory operating in the same region or area). For instance, health files may also make up the set of health care information held by the individual data controllers that participate in an EHR initiative at regional level...."

- How should the LHAs deal with audit, legal and regulatory compliance? What type of audit and workflow processes need to be implemented?
- Which deployment model (private, public, hybrid, community) best suits local healthcare authorities? Should public services (eg, storage of health records) reside in the same cloud service as back-end services (eg, payroll, HR, etc)?
- Which service model (IaaS, PaaS or SaaS) best suits the needs of the LHAs? Which of these provides the best match (if any) between service models and services (eg, online collection of medical files, online appointment scheduling and other less critical services, eg, back-end, HR, payroll, e-learning)? Taking the different service requirements into account, what types of SLAs need to be applied?
- How are the design, deployment and administration of the infrastructure to be managed across the various LHAs?
- Are there interoperability problems between the cloud and legacy systems that produce medical data in certain hospitals and how can these be overcome? What are the minimum requirements for the interoperability of e-health records?
- What is the real added-value of the regional e-health hybrid cloud in terms of a reduction in the cost of IT acquisition and maintenance?
- How can the LHAs ensure effective security controls in the resulting end-to-end solution? What forms of audits, SLAs, financial penalties or incentives, etc, will work best to deliver adequate assurance?
- What is going to change (if anything) in terms of the provision of the relevant services to patients and the usability of such services by professionals (eg, hospital doctors, GPs, LHA staff)?

Electronic health record (EHR)

An *electronic health record* (EHR) is a repository of information regarding the health status of a subject of care in computer processable form, stored and transmitted securely, and accessible by multiple authorised users. It has a standardised or commonly agreed logical information model which is independent of EHR systems. Its primary purpose is the support of continuing, efficient and quality

integrated healthcare and it contains information which is retrospective, concurrent, and prospective⁶⁸.

- The EHR is a secure, real-time, point-of-care, patient centric information resource *for clinicians*.
- The EHR aids decision-making by clinicians by providing them with access to patient health record information where and when they need it and by incorporating evidence-based decision support.
- The EHR automates and streamlines the clinician's workflow, closing loops in communication and response that result in delays or gaps in care.

Attributes and essential requirements⁶⁹

The EHR system must:

- Provide secure, reliable, real-time access to patient health record information where and when it is needed to support care.
- Guarantee patient health information confidentiality and security.
- Be available and reliable 24/7.
- Be responsive enough to integrate with the clinician workflow.
- Be accessible where needed—at inpatient and ambulatory care sites, with remote access.
- Capture and manage episodic and longitudinal electronic health record information.
 - Check information captured or imported for reasonableness and provides time stamps, information source, and amend audit trail.
 - Comply with approved industry standards for message vocabulary and content.
 - Accept information from external systems and automated data capture devices such as patient monitors, laboratory analysis equipment, and bar-code scanners.
 - Ideally accept and integrate health record information from outside the immediate organization, including medication dispensing information from community pharmacies.
 - Provide tools for unique patient identification and information integration across systems and settings without a common patient identifier.

⁶⁸ ISO/TR 20514:2005(E)

⁶⁹ The attributes of service and its essential requirements are mainly based on:
<http://www.himss.org/content/files/EHRAttributes.pdf>

- Permit efficient data entry of all orders and documentation by authorized clinicians. This includes prescription writing and refill management. Ideally supports various means of clinician entry (eg, keyboard, voice, pointer device, or handwriting recognition). Ideally, documentation includes clinical reasoning and rationale.
- Support electronic signature where permitted by law.
- Accept patient self-reported health information.
- Ideally differentiate between patient historical data (applicable across visits and across the continuum of care, eg, allergies) versus episodic data (applicable for one visit, eg, breathing sounds from last respiratory assessment) and support copying data forward as appropriate to support continuity of care, accuracy of ordering, and efficiency of clinical documentation.

Electronic health files

A *health file* presents the same essential attributes as electronic health records, the only difference being that the file is set up at a healthcare body that acts as the sole data controller (eg, a hospital or a nursing home) where several healthcare professionals are employed. Conversely, an *electronic health record* is a file set up by pooling the data from different data controllers, which as a rule - though this is not always the case - operate within the same geographical area (eg, a healthcare unit and a private laboratory operating in the same region or area). For instance, health files may also make up the set of healthcare information held by individual data controllers who participate in an EHR initiative at regional level. (<http://www.garanteprivacy.it/garante/doc.jsp?ID=1672821>)

Regional electronic patient record brokerage and exchange system

A *regional electronic patient record brokerage and exchange system* is a regional point of reference for all the healthcare information related to a patient living in that region.

The information (electronic patient records) is either stored directly in a regional repository shared with every national and international interested party (hospitals and clinics, general practitioners, etc) or is kept at local level (local healthcare authorities, hospitals, etc) and referenced by a link from the regional electronic patient records brokerage and exchange system.

The same attributes and essential requirements identified for the electronic healthcare record apply to this service.

Local and regional community cloud – scenario No 2

Cloud use-case involving local authorities in Spain

The provincial government of Jaén in the south of Spain wants to improve the participation of citizens in the Information Society and to promote e-government services. It is doing so via a project called Jaén Digital Province which has four strategic objectives:

1. to provide a digital communication infrastructure throughout the province;
2. to improve access to and participation in the Knowledge and Information Society for citizens;
3. to put the services and resources of town councils online;
4. to put the services and resources of the provincial government online.

The provincial government has recognised the value of putting in place common organizational and technological solutions for local administrations.

Here is what the provincial government has done so far to reach its four objectives.

Digital infrastructure

The digital infrastructure is based on the provincial communications network known as Heraclea, which provides broadband access to all the municipalities. The network has a total of 113 connections, of which 77 are town councils and the remaining 36 offices of the provincial government. So far, the provincial government has connected 97 municipalities to the provincial government via GIGADSL, which provides the gateways to online municipal services.

It also provides access to free software repositories via two administrative networks: the Andalusian regional government's network called NEREA and the national government's SARA network. The Andalusia regional government, of which Jaén is a province, has developed its own Linux distribution network called Guadalinux (<http://www.guadalinux.org/>). Guadalinux has its own office tools and offers citizens free access to its operating system software and tools.

Digital citizenship

The Digital Citizenship programme promotes the equal participation of the citizens of the Province of Jaén in the Knowledge and Information society. In 2001, the provincial government was among the pioneers in the introduction of centres with public internet access. Thirty-four of the province's 97 municipalities were equipped with telecentres. Now all municipalities have telecentres. Jaén now has a network of 161 telecentres, of which 62 centres are in villages. The telecentres provide citizens with free access to the Internet and, especially, programmes and activities focused on e-learning, e-commerce and other e-services.

Digital town councils

Jaén and other Andalusian provincial governments have joined forces to develop a common platform called the Digital Town Council ICT Model, which enables online administration and progressive

implementation of digital citizenship. The model has three layers: (1) an online services portal; (2) a municipal website; and (3) a town council back office, which manages census, land management, registry, water, taxation, accounts and payroll. As of April 2010, 23 municipalities had used the platform to establish their own online services portals. All municipalities are expected to follow suit.

Digital provincial government

The provincial government has made e-government a reality by putting its services online. The paper version of the *Provincial Bulletin* disappeared in December 2003; the provincial government now publishes only an electronic version. It has promoted free and open access software not only on its own website but also on those of the municipalities.

The provincial government has put its strategic plan online, together with indicators to measure the success of its implementation (known as the balanced scorecard). It has produced a *Guide to the Services of the Provincial Government* which has also been put online. With a view to achieving paperless administration, the provincial government has put its electronic register and communications online, as well as its plans for spending management, subsidies and tax management.

Background information

Jaén Digital Province is a project for enhancing technological cooperation and participation in the Information Society. It links the provincial government and the 97 town councils of the province. It is based on the following principles:

- Interoperability and free software as a basis for the components of the model.
- Working online through the sharing of communication networks (both within the province and at the Andalusian and national levels).
- Definition of common models of management and information systems for the enhancement of online government. The ICT model of the Digital Town Hall is part of the local Andalusian software repository. Not only is it to be implemented in the municipalities of Jaén, but it is also recommended for and available to other Andalusian and Spanish municipalities.

The project is based on an initiative of the Spanish Ministry of Industry, Tourism and Trade (MITYC), which provides the infrastructure, platform and a set of applications for all Spanish councils (of which there are 8,300). The governments have yet to move to the cloud, however, as each council needs to download, install and configure each application and thus needs to have its own infrastructure.

Among the services provided by the Ministry of Industry to the local councils are these:

- 1 LocalWeb, an application for generating websites,
- 2 SIGEM, an application for managing the administrative procedures of a file or record,

- 3 LocalGIS, an application for cartography management,
- 4 Registro, an application for the maintenance of the municipal population census,
- 5 Catastro, an application for the registration of citizens and company properties,
- 6 E-Easy, an application supporting the creation of enterprises and billing for local bodies.

ICT requirements of the Jaén model [fictitious, two years into the future]

The Jaén provincial government initiated its Jaén Digital Province plan following consultations with all of its 97 municipalities. Collectively, they identified their service and resource requirements, including a distributed network, multiple servers, adequate storage, a system allowing multiple and various types of applications and services. Specifically, they envisaged needs for:

- *An on-demand self-service* whereby any municipality could have as much server time and network storage as needed, automatically, without requiring human interaction with each service's provider.
- *Access to the network access via different devices*, including work stations, telecentres, mobile phones, laptops and PDAs.
- *Pooling of computing resources* to serve many different users, some municipal and provincial government officials and some members of the public.
- *Rapid elasticity*, ie, the network can respond rapidly and automatically to changes in demand from particular municipalities or the provincial government.
- *Measured resource usage*, so the system could meter and report the differing levels of usage (eg, storage, processing, bandwidth and active user accounts) by the municipalities, provincial government and even citizens.
- *Migrating existing services*, whereby they could continue to use some tailor-made or specialised services in the new system.
- *Positive cost-benefit*, where all stakeholders could benefit from costs lower than the diverse computing resources they have been using in essentially stand-alone environments. Furthermore, they wanted to pay only for actual usage, and not for resources they might not actually use.
- *Rapid deployment* – they wanted to be able to quickly provide new services without having to procure, certify and validate additional hardware and software.
- *High availability and reliability* – they wanted a system that was always available and always reliable or, at least, as near to always as possible. If there was a failure in a central server, they wanted another server to take over instantly. If there was a failure in the

network (eg, due to a power failure or a break in the communications line), they wanted instant back-up.

- *Ease of use* – they wanted a system with a short learning curve for those developing and using e-government services.

Based on their requirements, there was agreement that they should move to the cloud, but in a way that would allow them to take some legacy systems with them. There was some debate about which service model would be most appropriate.

Service models

Software as a Service (SaaS) – Many citizen-consumers need to use popular applications (e-mail, word processing, spreadsheets) as well as specialised applications (for obtaining parking permits, utilities, local council registries, library access, etc) as well as back office applications (payroll and taxation) which could be cloud-based and accessed via a web browser with different user devices (work stations, laptops, mobile phones, and PDAs). The citizen-consumer (and back office worker) would need only his or her device. The cloud would provide the software, applications, storage, and back-up.

Platform as a Service (PaaS) – The provincial government and municipalities use popular, commercial applications, but also need to develop their own specialised applications using programming languages and tools supported by the cloud provider. The cloud provider would manage the network, servers, operating systems and storage.

Infrastructure as a Service (IaaS) – The government and municipalities considered whether they needed control over the actual infrastructure (servers, operating systems, storage, applications, etc) but decided this was not necessary and that it would be more economical to let a cloud provider deal with these matters.

After considering deployment models – a private cloud, community cloud, public cloud or a hybrid – they decided to go for a community cloud.

Resilience

While the economic and other benefits of moving to the cloud were clear, service resilience, information security and legal compliance were concerns. The cloud provider's record in that regard was better than their own. Even so, their concerns regarding resilience included the following:

- *Data protection* (integrity, privacy and authenticity) – some of their services used personal data; hence, they needed assurance that the cloud provider would comply with Spain's data protection laws.
- *Availability, reliability, quality of service* – they needed services that would be always available and reliable (ie, that they did what was expected every time).

- *Back-ups and continuity* – if there were a crash of the main hosting server, another had to be able to take over ‘instantly’. Various externalities, such as power disruptions, physical and cyber attacks, needed to be taken into account.
- *Access control* – some services could be accessed by anyone, others (eg, social services) were controlled and limited to selected individuals. Access control needed to include measures for authenticating the users and to provide audit logging and monitoring. In Spain, many electronic public administration services can use the Spanish electronic identity card, which helps support the business case for Identity as a Service.
- *Audits and certification* – perhaps the toughest requirement set by the municipalities was that service provided by the cloud had to be subject to audit and, in some cases or services, had to be certified as being in compliance with security standards (ISMS, ISO 27001).

Electronic administrative procedures

An application was developed to manage the administrative procedures of a file or record. It allows citizens to request electronically (from their homes) subsidies, aid and licenses, or to make payments, receive news of the status of their requests, as well as information on the lack of any documents with instructions on how to attach them, and to receive, finally, a notification of the outcome of their efforts.

Gov cloud as business incubator - scenario No 3

The Minister of Communications and Technology and the Minister of Industry and Development were having an informal discussion after a meeting with the Prime Minister.

The topic under discussion was, as it often had been for the past year, cloud computing.

The Minister of Industry and Development said to his colleague: 'I had a meeting last week with my Japanese counterpart and he was explaining me about their J-SaaS project.

'The Japanese Ministry of Economy, Trade and Industry developed a computing infrastructure about a year ago. The system is called J-SaaS. J-SaaS is a computing platform that works as an incubation bed for SaaS providers and users.

'SME-class ISVs (Independent Software Vendors) can bring in their application packages to offer them as SaaS. SME users can use the service for both production or for experimental use at a low cost. Japanese ISVs used to sell software packages to SMEs who have little IT competence and CAPEX allowance. If ISVs' software products can be offered as SaaS, market barriers can be lowered. ISVs, however, could not offer SaaS because they don't have the infrastructure to provide SaaS. Thus the government prepared the infrastructure so that the vendors and users can make use of it and promote IT-based management, on the user side, and business in a new deployment model on the vendor side.'

'I was not aware of it,' replied the Minister of Communications and Technology. 'It sounds like a very nice idea that could be adopted in our country and in the whole of the EU (?) as well, especially if we consider that 99% of the businesses in the EU are small, medium and micro enterprises. We might also consider extending the services offered by this Gov cloud to include IaaS and PaaS as well.'


'Yes, I agree with you.' said the Minister of Industry and Development, 'but before presenting this idea to the public we should have concrete and solid answers to those very typical questions about cloud computing:

1. Is it secure enough?
2. Are ISVs trusted enough (trust management between consumers and ISVs)?
3. Where are the data going to be stored?
4. Will the Gov cloud be resilient enough? Can the Gov cloud offer a better SLA than the one currently offered to SMEs?
5. Is the resource concentration going to increase incentives for attackers?

6. Is our current legal framework adequate to face possible cloud challenges?
7. Which deployment model (private, public, hybrid, or community) best suits the purpose?
8. Is it actually going to be cost effective for the SMEs?
9. Can the Gov cloud adopt a model of grants and subsidies for SMEs to promote the migration process?
10. Will it distort the market?
11. What about critical infrastructure?
12. Liability – can the government support it?
13. Will it need to be supra-national to be resilient?
14. Can we define our expectation concerning risk clearly?

J-SaaS model

Structure of J-SaaS



Having as terms of reference the structure of the Japanese J-SaaS, it appears clear that the government can offer a variety of difference services. The government strategic view will drive business decisions, which will take into account, among others factors, the following possible variables:

Type of potential customer

- Micro enterprise
- Small and medium-sized enterprises (SMEs)
- Companies doing business in highly regulated business sectors
- Large enterprises (moving specific services).

All the above mentioned categories should be considered both as final users of the platform and as companies leveraging the governmental cloud to offer their IT services to other companies.

Needs of potential customers and governments

- Support for national businesses
- Integration with government capable services and data sets
- Reliable and cost effective IT infrastructure on demand potentially able to cover the whole spectrum of IT services.
- Reliable and cost effective IT platform on demand that can be integrated with or that is compatible with the internal service platform
- A test-bed with low availability requirements
- Sector specific services
- Value-added security services
- Trusted brand
- Legal compliance
- Clear responsibility and assignment of liability.

Annex III - Reservoir architecture description

Virtualised cloud architecture for community clouds

In this section, we describe an abstract virtualised cloud architecture that provides Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). The architecture introduces a virtualised infrastructure layer on top of the physical infrastructure. This abstraction layer is designed to manage a federation of heterogeneous physical infrastructures. Every site is partitioned by a virtualization layer into virtual machines (VMs) that are fully isolated runtime modules that abstract away the physical characteristics of the resource and enable sharing of resources. This architecture is based on three distinct layers:

- **Service/platform manager (SM):** is responsible for the instantiation of the service application by requesting the creation and configuration of a VM for each service component, in agreement with the service definition, thus ensuring compliance with the service level agreement (SLA). In the case of a platform, it deploys code on the appropriate platform, eg, on a java service container.
- **Virtual infrastructure manager (VIM):** is responsible for the placement of VMs onto host machines (HMs). It receives requests from the SM to create and re-size VMs and decides the best placement for these VMs to optimize a site utility function given a set of constraints (set by the SM). The VIM not only manages provisioning of the VMs, but also of the virtual networks (VNs) and virtual storage (VS) that is required. The VIM has full control of the HM. A policy engine (PE) is a VIM component that is responsible for the placement and migration of VMs onto a HM. The VM represents a virtualized resource hosting a certain type of VM. VIMs issue generic commands to manage the lifecycle of VMs, and VMs are responsible for translating these commands into commands specific to the virtualization platform abstracted by each VM.
- **Physical infrastructure manager (PIM):** This layer manages the physical machine, networking and storage equipment. It manages the addition and removal of resources to and from the pool of shareable resources.

As is shown in the figure, each layer has a management component to manage the services and platforms, the virtual infrastructure, and the physical infrastructure.




FIGURE 2 END-TO-END THREATS TO RESILIENCE

Threats to end-to-end resilience of a virtualised cloud architecture

Different types of users access the cloud via a network connection: the service provider who deploys and manages his multi-tier application on the cloud, the virtual infrastructure administrator who manages the infrastructure and the end-user who accesses applications running on the cloud infrastructure. These different types of cloud users are sources of threats. They need to be identified and countermeasures need to be put into place.

End-to-end cloud resilience can be threatened at any layer in a virtualised cloud architecture. The table below classifies some of the main threats for each layer as described in Figure 2 and the components that are threatened:

Threat	Layer	Threatened Component
Billing service failure or reduced availability	Service or platform	SM

Reduced availability of or failure in creating VM	Virtual infrastructure	VIM, VM, VN, VS
Reduced availability of or failure in shutting down VM	Virtual infrastructure	VIM
Reduced availability of or failure of the migration function	Virtual infrastructure	VIM
Reduced availability of the monitoring function	Service or platform, virtual infrastructure	
Network breaks	Physical infrastructure	Network
Compromise of network management	Physical infrastructure	Network
Application interference	Service or platform, virtual infrastructure	SM, VIM, VM, VN, VS
System overload, inability to scale	Virtual infrastructure	VIM
Compromise of hypervisor or OS	Virtual infrastructure	VM
Compromise of management interface	Virtual or physical infrastructure	VIM, PIM
Compromise of identity management system or provider	Service or platform, virtual or physical infrastructure	Service, VIM, PIM
DDoS or DOS attack on another healthcare authority effecting your systems	Service or platform, virtual or physical infrastructure	SM, VIM, PIM, service, platform
Compromise or failure of the accounting and billing system	Service or platform	SM

EHR service unavailability, or other service unavailability	Service	Service
Loss or compromise of EHR information	Service	Service
Data inconsistency	Service	Service
Disasters	All	All
Lost encryption keys	All	All
Bankruptcy of cloud provider or partner	All	All

Threats to end-to-end resilience in community clouds




FIGURE 3 END-TO-END THREATS TO RESILIENCE - COMMUNITY CLOUD

In the case of community clouds based on a federation of data centres, the various data centres are connected by physical networks and also by virtual networks. End-to-end resilience then faces different kinds of threats than it does in the case of a single cloud. As well as the threats to the different layers presented in the previous section, additional threats related to the structure of the underlying cloud federation need to be taken into account.

Figure 3 shows a community cloud composed of two sites A and B. The two sites are connected by networks and virtual networks. The two sites can provision resources from each other's sites. The figure shows the three types of clouds users (end-user, service/platform administrator and virtual/physical infrastructure administrator). Threats can occur at any location: at the cloud user site, on the network between the cloud user site and at the primary cloud site, on the network between the two cloud sites, and at the second cloud site.

The table below classifies the main threats based on the location of the threat in Figure 3 and the components that are threatened:

Threat	Location	Threatened Components
Billing service failure or reduced availability	Cloud site	SM
Reduced availability of or failure in creating VM	Cloud site	VIM, VM, VN, VS
Reduced availability of or failure in shutting down VM	Cloud site	VIM
Reduced availability of or failure of the migration function	Cloud site	VIM
Reduced availability of the monitoring function	Cloud site	
Network breaks	User to cloud network, community network	Network
Compromise of network management	User to cloud network, community network	Network
Application interference	Cloud site	SM, VIM, VM, VN, VS
System overload, inability to scale	Cloud site	VIM
Compromise of hypervisor or OS	Cloud site	VM
Compromise of management interface	Cloud site	VIM, PIM
Compromise of identity management system or provider	Cloud site	Service, VIM, PIM
DDoS or DOS attack on another healthcare authority effecting your systems	Cloud site	SM, VIM, PIM, service, platform
Modifying network traffic	User to cloud network, community network	Network, virtual network

Making an informed decision

Compromise or failure of the accounting and billing system	Cloud site	SM
EHR service unavailability, other service unavailability	Cloud site	Service
Loss or compromise of EHR information	Cloud site	Service
Data inconsistency	Cloud site	Service
Lost encryption keys	Cloud site	All

Annex IV – List of threats

ID	Threat Description
Threats applicable to all the scenarios	
1.	Network breaks (temporary loss of routing components under assumption that this can be recovered)
2.	Traffic Loss
3.	Network management (ie, network congestion / mis-connection / non-optimal use, etc)
4.	Application interference (access to code – which is then subverted)
5.	CP administrator error
6.	Malicious insider (ineffective screening, ineffective logs, etc)
7.	Exhausted CP resources (losing performance)
8.	System overload, inability to scale
9.	Compromise of hypervisor or OS
10.	Social engineering attacks (impersonation – eg, security on demand)
11.	Data leakage on up/download, intra-cloud (sniffing, spoofing, man-in-the-middle attacks, side channel attacks, etc)
12.	Management interface compromise (manipulation, availability of infrastructure)
13.	Identity management system or provider (is point of failure for the system)
14.	DDoS
15.	DOS on another healthcare authority effecting yours
16.	Undertaking malicious probes or scans
17.	Modifying network traffic

ID	Threat Description
18.	Privileges escalation
19.	Theft of computer equipment
20.	Compromise or failure of the accounting and billing system
21.	Replay
22.	Guest-hopping (compartmentalisation)
23.	EHR service unavailability
24.	Other service unavailability
25.	Loss or compromise of EHR information
26.	Duplication of data
27.	Data inconsistency (ineffective data update)
28.	Loss or compromise of operational logs
29.	Loss or compromise of security logs (manipulation of forensic investigation)
30.	Interception of data during migration and periodic update of data in the cloud
31.	Compromise of data during migration and periodic update of data in the cloud
32.	Disasters (natural)
33.	Unauthorized access to premises (including physical access to machines and other facilities)
34.	Backups stolen
35.	Lost encryption keys
36.	Unsecure deletion of data (when requested by the patient)
37.	Loss of governance, control, specifications (SLA may be incomplete in covering all KPIs)
38.	Bankruptcy of cloud provider or cloud partner

ID	Threat Description
39.	Acquisition of cloud provider or cloud partner (increases likelihood of strategic shift and may put non-binding agreements at risk)
40.	Liability with respect to regulations (eg, notification of data security breaches to customers)
41.	Legal and data protection issues
42.	Conflicting privacy and security requirements
43.	Conflicts between regional or local guidelines – necessity to apply
Threats specific to the community cloud model and more in particular to the federated approach proposed in Reservoir project.	
44.	Billing service failure or reduced availability: the IaaS user requests his bill based on pay per usage. Any reduction in availability of the billing service will affect IaaS resilience.
45.	Reduced availability or failure of creating VEE (Virtual execution environments): setting up the VEE involves downloading the system images, instantiating them with the configuration parameters, starting them up, and creating the virtual network between the VEE. Any reduced availability in initiating VEE will affect IaaS resilience.
46.	Reduced availability or failure in shutting down VEE: deletes a VEE following an un-deployed request. Any reduced availability in shutting down VEE will impact IaaS resilience.
47.	Reduced availability or failure of the migration function: migration allows to relocate a VEE from one host to another without any service interruption. Any failure in the migration functionality will impact overall IaaS resilience.
48.	Reduced availability of the monitoring function: monitoring of VEE allows assessing the state of the infrastructure, and taking corrective actions if necessary, such as restarting, scaling or migrating VEE. Any reduced availability in the infrastructure monitoring will impact overall IaaS resilience.
49.	Billing service failure or reduced availability: the IaaS user requests his bill based on pay per usage. Any reduction in availability of the billing service will affect IaaS resilience.