

Implementation of the Virtual Data Embassy Solution

Summary Report of the Research Project on Public Cloud Usage for Government, Conducted by Estonian Ministry of Economic Affairs and Communications and Microsoft Corporation



This Summary Report is for informational purposes only. Microsoft and Estonia make no warranties, express or implied, and provide no legal opinions in this report. It is provided "as-is." Information and views expressed in this report, including online references, may change without notice. The Government of Estonia does not endorse Microsoft or its products or services.

Today, all countries are striving towards a more productive information society that could offer their citizens better digital services. Being "digital" and therefore dependent on information and communication technology (ICT), however creates unprecedented challenges that are growing increasingly complex.

In 1991, when Estonia most recently established its independence, one of its first challenges was to determine who its citizens were. Each person who could prove that his or her parents or grandparents were of Estonian nationality and spoke Estonian was automatically granted citizenship. A second core challenge was to organize land reform. Land and property had to be returned to their rightful owners. Both ancestral backgrounds and property history were established using surviving paper records and archives. However, since then states have become increasingly digitalized and no longer store such information on paper. Continuity, or in this case digital continuity, can only be ensured if such information is kept secure even if the state suffers a large-scale cyber-attack, natural disaster, or a conventional attack on a datacenter.

Estonia has been backing up important data outside of its borders for roughly ten years. Data is stored in Estonian embassies across the world. This approach makes sense. Backups are the most effective method of protecting data from being lost. However, in this process the government must nevertheless consider a number of issues, including: what should be backed up, who has the right to make this decision, and whether there should be a "crisis button" to delete all of the data. Recent events prove that modern methods of war could introduce new levels of fragility, where questions of control become ever more important.

All countries will face such questions when deciding how to protect their move from a paper to a digital world. What makes the Estonian situation even more complicated, is that data backups are not enough; services need to be continuously available as well. In Estonia, for example, 98% of banking transactions are done online, and most use either national identity cards or mobile IDs. It is of utmost importance that, even if a crisis develops in Estonia, digital authentication and authorization services remain operational.

Changes in the world are unrelenting, and solutions that have worked previously are no longer good enough. Geopolitical events in 2014 brought the question of continuity to the forefront of national conversations in Estonia. A brand new Government Cloud Policy stated that, to ensure service functionality and data continuity, capabilities need to be developed outside of our borders in addition to a national cloud in Estonia. Two new approaches were adopted:

- *physical embassy for data in a friendly foreign country;*
- *virtual embassy for data in a privately owned public cloud.*

In September 2014, we began considering the changes and requirements needed to implement our new policy initiative. As part of our efforts, we embarked on a research project with Microsoft to test whether a public-private cloud computing partnership model could contribute to our policy goals. This Project seeks to test the potential of public cloud solutions offered by a private sector company to extend the goals and objectives of our Government Cloud Policy.

As part of this research project, we have evaluated methods to ensure that the data and services of and for our citizens, e-residents, and institutions are kept safe, secure, and continuously available. Privacy, security, data protection, and data integrity are central to our government services. After the Snowden revelations, both governments and large corporations are facing a trust-deficit. For that reason, having an expert like or trust a particular type of technology is not good enough – every citizen needs to be able to trust it. For this reason, transparency and data protection need to be ensured.

As one of the most connected countries in the world, we strive continually to develop and improve our e-government services. The Virtual Data Embassy Solution proves that innovation cannot solely rely on new ways of using technology – we also need to revolutionize our way of thinking about the role of the state. In the traditional physical world, an embassy is a sovereign representation of a government in another country.

In fact, the building is respected as if it were the very territory of that country. As we move deeper into the information age and embrace a digital society, we face new questions and challenges. What is sovereignty in cyberspace? Can a country declare or be recognized by others as having data sovereignty in cyberspace? These are the hard policy questions that we are intent on addressing.

Estonia must be able to continue to function as a government, and as a people, even in the direst of scenarios, including the loss of our territory. Since we do not have paper backups of data, our demands for data protection, security, and privacy are unparalleled. Any breach could have catastrophic consequences. Our digital services need to not only be the best in the world, but also the most secure and resilient.

We believe that the Virtual Data Embassy Solution is at the cutting edge of national e-governance policies. However, it will only succeed in conjunction with a broader acceptance of national cloud and physical data embassy policies.

Taavi Kotka

Government CIO

Cloud computing enables and extends major opportunities for governments and their citizens. This is why Microsoft was particularly excited to partner with Estonia on the Virtual Data Embassy research project. A number of interesting technical and policy questions have been raised throughout its course. We support the thoughtful ideas that the Estonian government has brought to the table to further the important objectives of continuity of government and operations. These include the suggestion that governments should consider distributed cloud computing in a “public” cloud for storage of government data, with limited exceptions; that other governments should respect the integrity and sovereignty of another country’s data in the cloud as it would its physical territory; and that all governments should work together and with the private sector to evolve their interpretations of international laws and policies in cyberspace so that all nations can continue to benefit from the promise and power of technology.

We hope that more and more governments will emulate Estonia’s innovative approach to e-government and look to bring a greater number of services to citizens using cloud computing. However, governments must also recognize that attacks on the digital assets of a nation state have consequences, both technical and legal. As a result, to enable the scale and utility of the cloud, governments must be willing not only to recognize the inviolability of other governments’ digital assets, but also to work together to prevent attacks and to hold accountable those who commit them.

On many levels, technology is the easy part – online services today are robust enough to meet the volume and other needs of citizens’ digital interactions with governments. This research project has shown that e-government services can and will exist in the cloud, and that citizens will continue to benefit from the use of cloud computing. Certain laws or policies may need to be revised domestically or evolved internationally to ensure cloud computing can support particular government functions; however the core concept is viable. In time, we hope that the corpus of law will extend protections afforded to a physical embassy to the virtual world, recognizing virtual data embassies, as well as to all other government assets as they move from a physical world to the online environment. That will enable true e-government and represents an exciting prospect.

Matt Thomlinson

Vice President, Microsoft Trustworthy Computing

Table of Contents

1.	Executive Overview	5
2.	The Estonian Data Embassy Initiative	7
3.	Virtual Data Embassy Research Project	9
3.1.	Project Outline	9
3.2.	Policy and Legal Environment Research	12
3.2.1.	Domestic Policy and Legislative Environment Overview	12
3.2.2.	International Policy and Legal Environment Overview	14
3.2.3.	Policy and Legal Environment: Findings	15
3.3.	Technical Research	19
3.3.1.	Estonian Government ICT Architecture Overview	20
3.3.2.	Technical Project Overview	21
3.3.3.	Technical Architecture Findings	26
3.3.4.	Technical Performance Findings	40
4.	Conclusions and Recommendations	46
5.	Research Project Team	49

1. Executive Overview

According to IDC¹, 70% of chief information officers in 2016 will consider cloud-based delivery the preferred choice when implementing new services. Organizations are rapidly adopting cloud computing to gain speed, scale, and economic benefits. Governments are following suit. They are exploring cloud-based services to: create scalable, interactive citizen portals; collaborate more easily; deliver volumes of data to citizens in useful ways; and, maximize focus on mission-critical needs, while reducing ICT costs. They need, however, to be sure of the security, resilience and trustworthiness of the services they run “in the cloud.”

In 2013, the Estonian government began pursuing a *Data Embassy Initiative*, reflective of its innovative approach to e-government and of its need to ensure national digital continuity no matter what. Cloud computing, with its immense opportunities for resilience, security and continuity in light of physical or cyber emergencies, was a potential solution. In September 2014, the Ministry of Economic Affairs and Communications, the Ministry of Justice (Center of Registers and Information Systems), and the Office of the President of Estonia agreed with Microsoft to work on a research project to assess the feasibility of the virtual aspects of the *Data Embassy Initiative*.² In particular, the collaborative project tested how two separate government services – the official web site of the President of Estonia (www.president.ee) and the Riigi Teataja, or electronic State Gazette (www.riigiteataja.ee) – could be migrated and hosted on the Microsoft Azure™³ cloud computing platform.

In the following report, the project team summarizes its research, which took place over three months. It addresses the Estonian *Virtual Data Embassy Solution*, a key part of the *Data Embassy Initiative*. It also looks at the current Estonian government ICT architecture, for context, and describes the “data embassy” concept, the website migration process, and the verification testing that was conducted to ensure that the migration was successful and to assess the security and resilience of the cloud computing services.

Particular focus was given to the potential legal protections of a Virtual Data Embassy, as the success of the initiative fundamentally relies on the ability of citizens to trust the security and privacy of such embassies. The latter naturally draws in at least three actors: the Estonian government, the cloud service provider, and the country wherein the cloud provider is headquartered. The technical outcomes are also outlined, i.e. storage, network and compute architecture, with operational lessons, as well as security, identity, and data architecture findings set out. The report concludes with high level recommendations, which could be applicable to any government, as they consider cloud computing to achieve their national objectives.

Project’s three core findings are:

- i) The Virtual Data Embassy is consistent with Estonia’s existing domestic legal framework, with certain caveats;
- ii) Migrating and running the selected government services is technically feasible; and,
- iii) The government needs to be flexible to benefit from the latest technological advances and protections to ensure digital continuity.

¹ Worldwide and Regional Public Cloud ICT Services 2014-2018 Forecast, <http://www.idc.com/getdoc.jsp?containerId=prUS25219014>

² This Project, set up to assess the concept of Virtual Data Embassies, did not entail any commercial or financial commitment between Estonia and Microsoft.

³ <http://azure.microsoft.com/en-gb/>

Table 1: Project Overview

Key hypothesis tested: Operating government services from a public cloud as a Virtual Data Embassy offers resilience, security and scalability, as well as potential legal protections against compelled disclosure

	Legal aspects	Technical aspects	Operational aspects
Key findings of the research project	<p>There are no legal restrictions under existing Estonian domestic law to migrating government data to a Virtual Data Embassy in a public cloud, with limited exceptions. These include data relating to “critical” services, as defined by the Emergency Act (see below). Data in a Virtual Data Embassy could be protected under international law from compelled disclosure.</p>	<p>Technical aspects of application migration and running services in the cloud are relatively straightforward.</p> <p>However, textbook readiness is near impossible to achieve.</p>	<p>A multi-disciplinary approach is an effective way to assess the legal, technical, and risk management aspects of migrating to the latest technology (starting with public cloud).</p> <p>It can also be seen as an important step in gaining public trust in the cloud.</p>
Key challenges	<p>The <i>Virtual Data Embassy Solution</i> has not yet been tested under international law and changes/clarifications to Estonian domestic law might be required to support faster adoption of cloud computing.</p>	<p>Generally, developing a resilient, end-to-end approach to ensuring identity, data integrity, access control, and availability of government services is a key challenge. If digital identity systems remain on-premises and become unavailable, cloud services would fail. Furthermore, ownership and operation of the domains associated with the Domain Name System (DNS) is critical to ensuring digital continuity.</p>	<p>Public trust in cloud services is hard to gain and difficult to maintain.</p> <p>Challenges span management, operation, and public policy concerns. Operational and technical practices are yet to be defined, need to continue to evolve, and to be documented.</p>
Key recommendations	<p>Estonian security standard ISKE should be updated, building upon international standards, as appropriate, to address cloud computing.</p> <p>Risk assessments should be conducted to establish acceptable risk tolerance and associated requirements.</p> <p>For the migration of critical services, countries should consider developing digital continuity legislation and a strategy to increase assurances of diplomatic and other international law protections.</p>	<p>Cloud computing should be utilized to increase security and resilience of government infrastructure.</p> <p>An overarching cloud strategy and government action plan facilitating cloud migration should both be developed to enable technical and operational agility and increase cost-effectiveness.</p>	<p>The trusted relationship between the government, private sector company, host country, and the country where the provider is headquartered should be deepened for success of the Initiative.</p> <p>Operational procedures should be prepared and tested in advance rather than in a crisis.</p> <p>Governments should keep data governance and security models up to date across the data lifecycle and required of government entities.</p>

2. The Estonian Data Embassy Initiative

Estonia is highly dependent on information technology. Estonian citizens are able to perform nearly every public and private sector transaction in digital form, and a vigorously implemented “paperless” policy means that some essential registries, e.g. the land registry, only exist digitally and only have evidentiary value in digital form. Moreover, its innovative approach to e-identity for non-residents signals the beginning of Estonia’s transformation into a “country without borders.” As a result, Estonia needs to reassure not only its citizens but also its e-residents of the viability and durability of the state itself and of their status within it, even in the face of cyber-attacks, natural disasters and other national or internal emergencies. Such trust in ICT is not easily won, however, and is even more difficult to maintain.

This requires more than just the preservation of critical data sets and ICT services on Estonian physical territory. A solution needs to be developed for situations, admittedly improbable, during which the Estonian state might need to operate some services outside its current borders. This is the Estonian government’s concept of “digital continuity” in the context of the development of e-government. In 2013, the *Data Embassy Initiative* emerged as a possible answer, with a data embassy being defined as a physical or virtual data center in an allied foreign country that stores data of critical government information systems and mirrors of critical service applications.

Three Core Elements

In essence, the Initiative consists of additional security measures that would allow Estonia to ensure continuity in government and operations, including: digital and data continuity (backup); data integrity (non-repudiation); and core government services in the event of a physical or cyber emergency. To achieve these goals, Estonia plans a three-part solution consisting of: i) maintenance of data backups and live services within Estonia’s borders (Government Operated Cloud); ii) backups at physical Estonian embassy locations or dedicated data centers in allied countries chosen by the government (Physical Data Embassy); and iii) backups of non-sensitive data in private companies’ public cloud (Virtual Data Embassy). All three parts of the Initiative should be seen as complementary to one another.

Within the government operated cloud, Estonia plans to have additional data centers and backups for e-government services located within its physical borders. However, the concept of digital continuity, requires that the official version of services, including government services like the State Gazette, is available and can be used and updated in real time and in all circumstances.



Figure 1: Three elements of the Data Embassy Initiative

The Physical Data Embassy element provides additional measures of security, with a server resource that is completely under Estonian government control, but located outside of Estonia’s physical borders. The Physical Data Embassy includes two approaches, which are currently being explored. The first would utilize government cloud solutions that have been developed by Estonia’s closest allies. For instance, Estonia

could selectively sign bilateral agreements to procure existing cloud computing in dedicated data centers of an allied country. The second option, further elaborated upon in the international policy section below, seeks to use existing Estonian embassies to house backups for registries, taking advantage of the embassies' established diplomatic status. This approach would extend Estonian jurisdiction to the e-government services in question and ensure that they are afforded the same protections, including immunity, as a physical embassy, consulate, or ambassadorial residence. Indeed, transforming server rooms in physical embassies into data embassies would allow Estonia to create a network that would ensure its digital continuity, even in the face of determined efforts to damage it or take it offline completely.

The third element of the initiative, the Virtual Data Embassy Solution – seeks to further augment digital continuity by using appropriate commercial cloud computing products and services⁴ as an additional security measure for Estonian e-government services. This aspect forms the core of this report, which focuses on the feasibility of such an approach, explores solutions to challenges encountered, and identifies a number of benefits that the Solution would offer. Notably, while Virtual Data Embassies might offer a higher guarantee of availability, certain data or services (e.g. state secret data) may not, at present, be hosted in a privately-owned cloud service due to data protection, privacy, and data integrity concerns. Nevertheless, while the use of public clouds does not eliminate all risk, their capacity to deal with the most widespread cyber-attacks currently exceeds that of many organizations. Moreover, their location outside of the physical borders of Estonia and within the global cloud environment makes the public cloud well suited to meeting the digital continuity goals of the Virtual Data Embassy Solution.

CLOUD BENEFITS FOR GOVERNMENTS

- **Citizen services.** Ability to drive innovation with data services in the cloud that citizens can reuse.
- **Infrastructure.** Reduction in data centers and public sector ICT can drive hardware efficiencies.
- **Flexibility.** Allows the meeting of real-time needs, or offloading of onsite data to the public cloud as needed to improve operational efficiencies.
- **Collaboration.** Enables more effective communicating and collaboration.
- **Continuity of operations.** With centralized data storage, management, and backups, data recovery can be faster and easier.
- **Creative IT.** Since cloud services can be centrally managed, IT workers are freed from a “keep-the-lights-on” approach, providing more time to foster creative problem-solving.

Table 2: Cloud computing benefits for government

⁴ **Private cloud** is cloud infrastructure operated solely for a single organization, whether managed internally or by a third-party, and hosted either internally or externally. In a **public cloud** the services are rendered over a network that is open for public use. The **hybrid cloud** is a composition of two or more clouds that remain distinct entities but are bound together, offering the benefits of multiple deployment models. **On-premises software** is installed and run on computers on the premises of the person or organization using the software, rather than at a remote facility.

3. Virtual Data Embassy Research Project

3.1. Project Outline

This research project was initiated in September 2014 to assess the feasibility of the Estonian *Virtual Data Embassy Solution*. The entities involved included Microsoft, the Estonian Center of Registers and Information Systems (RIK), the Estonian Ministry of Justice, the Office of the Estonian President, and the Estonian Ministry of Economic Affairs and Communications. The core objective was to demonstrate that selected non-restricted⁵ e-government services could be run from outside Estonian physical borders in the event of a cyber-attack, natural disaster, or other circumstances. Moreover, the project team wanted to show that the impacted services could be restored quickly, fully, and in an orderly manner.

The research project was also intended to help to establish that public cloud platforms could be trusted to host and run the selected e-government services. While the Estonian general public has a high level of trust in what is currently being provided, the perception of cloud computing in the country has suffered over the past two years. The project sought, therefore, to demonstrate that both government and citizens could rely on cloud computing without compromising their privacy, security, or control.

What emerged early on was the important role that the public cloud could potentially play not only in ensuring digital continuity, but also in verifying security and privacy. For example, independent third party auditors⁶ could help to build trust by verifying elements of cloud security, including encryption, regular data back-ups, data access limitations, and visibility into the availability of and changes to the service of a specific cloud platform. In addition, they could verify privacy safeguards are in place, for example such that ensure that data is not used for advertising purposes, and that EU Model Contracts⁷ for transfer of personal data to third countries are complied with.

#	Workstream (Government Service)	Description
1	President of Estonia website www.president.ee	Migration and Hosting of Official Website of the President of Estonia onto the Microsoft Azure™ platform
2	“State Gazette” website www.riigiteataja.ee	Migration and Hosting of the “State Gazette” Riigi Teataja Web Site onto the Microsoft Azure™ platform

Table 3: Workstreams of the research project

⁵ Restricted data, as per the Public Information Act (further reference in the domestic policy section below), includes information that is intended for internal use of a government institution or other public body and sensitive personal data, the disclosure of which violates private life or business secrets.

⁶ Microsoft does not allow customers to individually audit its cloud services or its data centers. Instead, independent third party audits are conducted and shared with customers. Together with certifications they demonstrate how Microsoft obtains and meets its security and compliance objectives. In addition, Microsoft has developed an extensible compliance framework that enables it to design and build services using a single set of controls to speed up and simplify compliance across a diverse set of regulations and rapidly adapt to changes in the regulatory landscape. More information on specific compliance programs is available here: <http://azure.microsoft.com/en-us/support/trust-center/faq/>

⁷ http://ec.europa.eu/justice/data-protection/document/international-transfers/transfer/index_en.htm

The two services selected for the research project, as per Table 3 above, only involved non-restricted public data and services. Data classification, e.g. organizing data into categories in terms of information security, could in the future help guide further implementation of similar initiatives. For example, a future project could test the feasibility of using public cloud services to securely host sensitive data.

The first government service, a “monument” website,⁸ is the official website of the President of Estonia (www.president.ee). The second is a key government registry, the Riigi Teataja or electronic State Gazette (www.riigiteataja.ee). The State Gazette is a critical, public government web service, hosting Estonia’s official body of laws. No paper copies of the laws are kept and legal acts are only in effect if they have been published online on the State Gazette. Without the State Gazette, the Estonian legislative system would not be able to function, so testing the potential for migrating it to the cloud was of utmost importance.

While both services only include public, non-restricted information, protecting them from cyber-attacks was nevertheless essential to ensure continuous access to legal information and to maintain Estonia’s international reputation. Selecting these government services allowed the participants to evaluate the technologies used⁹ in a realistic context, to examine the existing software architecture, and to identify areas where such services could be enhanced through the use of public cloud technologies.

#	Hypotheses to be tested
Hyp 1	Services migrated to the public cloud and able to run successfully.
Hyp 2	Minimal architectural changes required to migrate the services.
Hyp 3	A number of areas for improvement to be identified, given that existing services were originally developed to run on-premises.
Hyp 4	Minimal time and effort required to modernize the government services after migration.
Hyp 5	Cloud built-in capabilities, such as auto-scaling, Distributed Denial of Service (DDoS) protection, etc., leveraged to enhance the performance, stability, security of, and public trust in the services.
Hyp 6	Design of the operational procedures, such as failover and fail-back processes, are key to the services running in the cloud.
Hyp 7	Estonia’s ISKE security standard should be updated to address public cloud.
Hyp 8	Legal basis exists for asserting international law protections against compelled disclosure of data stored in a Virtual Data Embassy.
Hyp 9	Under Estonian domestic law, non-restricted data can be migrated to the public cloud.

Table 4: Hypotheses for the research project

⁸ Monument websites in this context are websites with symbolic status, which contain only public location. However, defacement or other attacks on those websites would be seen as damaging to the reputation of the country.

⁹ Several technologies and products were used during the research project. Cloud technologies: Microsoft Azure™. Website of the President of Estonia: Operating system; FreeBSD, application stack; PHP and MariaDB. Electronic State Gazette: Operating system; CentOS, application stack; Apache, Java and PostgreSQL.

The research project specifically explored the technical feasibility of migrating and running the selected services in the public cloud, and sought to understand how the transition from on-premises to the public cloud could be achieved optimally from both a technical and operational standpoint. In particular, it sought to demonstrate that the services could be run in the cloud including in a pre-agreed, limited, controlled failover scenario under non-peak operation time. Another key goal was to use the migration to record findings regarding areas in which the existing government services could be optimized and improved for their operation in the public cloud.

Sections 3.2 and 3.3 below contain further detail on the design of the research project and the tested hypotheses (see Table 4). In Section 3.2, the Estonian and international legal environments, which provide the overarching framework to the *Data Embassy Initiative*, are outlined and ways in which the current government ICT architecture can be adapted to using cloud services are discussed, along with where the existing website and web services can be enhanced using cloud technologies. Section 3.3 considers the technical dimensions of the project, describing the approaches and methodology used for migrating the selected e-government services onto the public cloud, and how their effective operation and performance was assessed. For the project's key findings, conclusions and recommendations, please see section 4.

3.2. Policy and Legal Environment Research

The Project team found that there are no legal restrictions under existing Estonian domestic law to migrating government services to a Virtual Data Embassy in a public cloud, with limited exceptions. These include data relating to “critical” services, as defined by the Emergency Act. It was also established that the Estonian data in a Virtual Data Embassy could be protected under international law from compelled disclosure. To enable migration of government services with restricted data, however, changes to laws must be addressed, for example to clarify digital continuity. Furthermore the migration of restricted data raises other legal issues that have not been specifically addressed in domestic law or international practice.

In particular, the *Virtual Data Embassy Solution* is unique because it seeks to be recognized by the international community as having legal protections associated with diplomatic and consular missions while applying these protections to novel technologies and circumstances. It applies existing international laws, many of which are based in treaties that were ratified nearly fifty years ago, in new ways, and its success depends on acceptance by governments around the world. This means that the research project had to examine not only the feasibility of the technical implementation, but also its legal feasibility, in particular with regard to the availability of diplomatic, consular and sovereign immunities for data stored on the Virtual Data Embassy. The following sections address this question, taking as starting points the legal and policy landscapes in Estonia, and internationally.¹⁰

3.2.1. Domestic Policy and Legislative Environment Overview

Estonia is a pioneer in e-government and ICT adoption and use. This has been true both in terms of technology and legislative adoption. Over the last decade, Estonia has continuously adopted an inventive and forward-thinking legislative environment to support its commitment to a digital society. The basic policy documents governing e-government in Estonia, as identified by the Project team, are the [Principles of the Estonian Information Policy](#), approved in 1998, and reviewed in 2005 and 2010. In 2007, [the Estonian Information Society Strategy 2013](#) entered into force, for the first time clearly setting out the government objective of developing information society in the country as a strategic choice. In late 2012, the government built on the document with the [Estonian Digital Society Strategy 2020](#), committing the country to a dramatic leap forward when it comes to e-government, with the introduction of concepts such as virtual data embassy, digital identity, and virtual residence for non-citizens.

In addition to these e-government policies, successive Estonian governments have adopted and implemented various security policies and frameworks. The first such was the [2005 Information Security Policy](#), which was implemented as Estonia conducted its first e-elections. It was updated in 2009. Also worth mentioning is the 2008 [System of Security Measures for Information Systems](#), which establishes security measures for information systems processing data in state and local government databases and for related information assets. The 2013 [Security measures for vital service information systems and for the related information assets](#) establishes that a provider of a vital service must ensure the constant application of security measures and institutes reporting requirements. Concurrently, the [National Cybersecurity Strategy](#) was adopted in 2007 with a review published in the summer of 2014. The latter outlines the government cybersecurity objectives, including for the *Virtual Data Embassy Solution*.

¹⁰Microsoft offers no opinion on the state of Estonian domestic law for the purposes of this Report.

In 2008, Estonia also adopted [ISKE \(Three-Level IT Baseline Security System\)](#), a system of security measures for information systems that contain and process non-state secret data on state and local government databases. ISKE establishes procedures for the specification of security measures, creating a three-level baseline security system for high-, medium-, and low-risk systems, and describes organizational, physical and ICT security measures to protect data. ISKE is based on [IT-Grundschutz](#), a German standard that complies with the ISO 27000 family of standards, while offering granular technical information to support implementation. While ISKE largely maps to IT-Grundschutz, the Information Systems Authority, Estonia's regulatory authority that manages ISKE, has added some Estonia-specific content. In particular, ISKE contains additional content that is relevant to Estonia's national identification cards and X-road.

ESTONIAN POLICY FRAMEWORK: CORE DOCUMENTS

- **Information Society Services Act (2004):** The Act, [amended](#) in 2014, implements the European Union (EU) Directive [2000/31/EC](#). It establishes requirements pertaining to information society service providers, as well as the organization of supervision and liability in the case of violation.
- **Public Information Act (2001):** The Act, last amended in 2015, regulates access to information, other than classified, created and maintained by government and local government institutions and other public bodies. It defines restricted information and sets grounds for granting access to restricted information. Since 2008, the Act regulates also establishment and administration of public databases, and supervision over the administration of such databases. Adoption of security standards, including ISKE, is delegated by this Act.
- **Electronic Communications Act (2004):** The Electronic Communications Act, [amended](#) in 2014, implements the [EU Regulatory Framework for Electronic Communications](#). Its purpose is to create the necessary conditions to promote the development of electronic communications networks and communications services, while ensuring the protection of the interests of the users of such services.
- **Digital Signatures Act (2000):** The Digital Signatures Act provides for the use of digital signatures and digital ink, and the conditions of certification and oversight procedures for time-stamping services. The Act was [amended](#) in 2007, and again in 2014.
- **Personal Data Protection Act (1996):** The Personal Data Protection Act was amended in 2003, to fully comply with the EU Data Protection Directive [95/46/EC](#), and again in January 2008. The [2008 version](#) of the Act introduced several changes. Firstly, classifying data into two categories: (1) 'personal data' and (2) 'sensitive personal data', the latter being the sub-class under special protection. Secondly, all processed personal data to be protected and registered by chief processors with the Data Protection Inspectorate.

Table 5: Estonian policy framework

Beyond Estonia's ICT development and information security policies, Estonian laws guiding government actions during emergencies are also relevant in the context of the *Virtual Data Embassy Solution*. This is important for the concept of digital continuity, in particular to establish when the usage of the Virtual Data Embassy would be activated to ensure continuity of government services during different types of emergencies. Estonian legislation recognizes three levels of emergency: a state of war, a state of emergency, and an emergency situation. First, during a state of war, the [War-Time National Defense Act \(1994\)](#) applies; in addition, the [Emergency Act \(2009\)](#) applies to the extent that it does not conflict with the former. Second, during a state of emergency, the [State of Emergency Act \(1996\)](#) applies. Third, during an emergency situation, only the Emergency Act (2009) applies.

3.2.2. International Policy and Legal Environment Overview

This section considers certain international laws that might apply to Estonian government data and critical services hosted in data centers located outside Estonia's physical borders and owned and operated by a third-party provider.¹¹ When considering government data hosted with a third party, including a cloud service provider, numerous areas of international law are potentially relevant. These include sovereignty, data protection, data custodianship, diplomatic protection, consular protection, and sovereign immunity. This research project recognizes that while many international laws pre-date the internet and additional work between governments is needed to address previously unanticipated circumstances, existing laws can be applied to cover new technologies and circumstances.

This is a complex area, particularly as a sovereign owns the data. Although Estonia would have well-founded arguments that its data is protected from compelled disclosure, the untested state of the law makes it impossible to be certain whether a claim of international legal protection would be respected, either by the host state or by third-countries. Whether the protection afforded is that of an embassy or consular facility and whether Estonia seeks protection as an extension of the sovereign state itself also pose challenges. This underscores the need for an international legal framework, based on principles of transparency, due process, and respect for human rights and privacy, which sets out clear processes for governmental access to data in the cloud.

Estonia could pursue several options to protect its data. The [Vienna Convention on Diplomatic Relations](#) (VCDR) and/or the [Vienna Convention on Consular Relations](#) (VCCR) could be applied. For example, Article 24 of the VCDR provides that "[t]he archives and documents of the mission shall be inviolable at any time and wherever they may be." This could apply to modern storage methods, including those that store government data outside of embassy premises. It follows that Estonia might have a solid basis to argue that data associated with its diplomatic mission should retain this protection even when held by a third-party cloud provider. Indeed, the U.S. State Department's position is that documents can retain their Article 24 protection when in the hands of third-parties acting as an agent or contractor to the state.¹²

Similarly, Article 33 of the VCCR provides that archives and documents "*of the consular post*" are inviolable. The VCCR's definition of consular function includes a notable catch-all for "*any other functions entrusted to a consular post by the sending state,*" provided the host state does not object. Estonia might also seek protection of its data under the customary international law principle of sovereign immunity. The doctrine has been interpreted to extend to all non-commercial property of a state situated abroad, so the government could assert that its data is considered "*non-commercial property*" situated abroad in a data center located outside of Estonia. As to all of these issues, the Virtual Data Embassy presents novel circumstances, so there is no clear precedent for how these various protections would apply.

Clearly, foundations exist that support the extension of a sovereign's right to inviolability of its data to the internet and cloud storage. Governments around the world need to begin to come together in support of an interpretation of both treaties and customary international law that recognizes sovereign data protection rights and obligations.

¹¹ Any laws or conventions cited in this document may or may not be relevant to circumstances outside those of this research project, and should not be interpreted beyond that scope.

¹² See Eileen Denza, *Diplomatic Law* 198-88 (3d ed. 2008)

3.2.3. Policy and Legal Environment: Findings

Although Estonia has a well-developed ICT policy landscape, it has not yet adopted an overarching cloud computing policy. Nevertheless, recalibrating Estonia's existing information assurance frameworks might be prudent, as cloud computing dramatically changes the way data moves across platforms, devices, services, as well as borders. The primary drivers for this are the digital continuity concept, which is central to the *Data Embassy Initiative*, and the need to operate some services outside Estonia under other circumstances, e.g. for e-residents. In the following sub-sections, we discuss the conclusions of the research project across three policy areas that are pivotal to ensuring trust in cloud computing: security; data protection; and, digital continuity.

3.2.3.1. Information Security Provisions

When analyzing security provisions, the research project focused on the ISKE standard, highlighted above. ISKE is based on the German IT-Grundschutz and sets the standard that helps the government estimate their availability, integrity and confidentiality needs. As with the IT-Grundschutz, ISKE is applicable to cloud computing but does not specifically address it. Moreover, ISKE's application to cloud services might be challenging because the output of the methodology it outlines is a collection of specific security measures. Requiring cloud service providers to implement these would eliminate the opportunity to take advantage of the cloud service provider's most up-to-date resources and expertise. As a result, an update of ISKE's auditing provisions might be necessary.

In the interim Estonia should continue to leverage ISKE to evaluate all the components of its information domain that are external to the cloud system, e.g. the network, client systems and software. This would allow the public cloud system to be treated as an independent component. This approach is also supported in IT-Grundschutz for components that cannot be adequately modelled using the catalog or that were not foreseen in the standard's scope.¹³ This would mean that a supplemental security analysis for the cloud service component would be needed, requiring verification that the cloud service provider successfully meets specific security objectives. This could for example be aligned to [ISO/IEC 27001/2](#). Specifying the security requirements at the level of the security objectives would allow cloud service providers flexibility in implementing controls that match today's evolving threat environment.

3.2.3.2. Data Protection

As highlighted above, two legislative documents govern data protection in Estonia: the [Personal Data Protection Act](#) and the [Public Information Act](#). Together they dictate how to divide Estonian public sector information into different categories. For example, information with no restrictions placed upon it (Article 28 of Public Information Act) could be stored in the cloud. However, this is not true for all public sector information, as per Chapter 5 of the Act. For the latter, justified interest in doing would have to be demonstrated. For personal data and sensitive personal data further restrictions on cloud usage may be introduced. Hosting the data in the Estonian Government operated cloud on servers located within embassies abroad would ensure the security of this type of data and achieve compliance with the Data Protection Act.

¹³ Discussed in BSI-100-2 Section 4.6.2, pp. 72.

Another important point to consider is that existing law stipulates that the cloud provider, for example the provider of the Virtual Data Embassy, becomes an authorized processor of data and is thus required to comply with the instructions of the chief processor. According to Public Information Act (Article 43), the chief processor of a database is the state or local government agency, other legal person in public law or person in private law performing public duties that organizes the introduction of the database and the administration of services and data. They may also, based on a procurement contract or a contract under public law, authorize a person in private law to perform the tasks of processing of data and housing of the database. This means, as there is not a direct mandate for the outsourcing of database hosting to the private sector, an amendment to Public Information Act might be advisable.

3.2.3.3. Digital Continuity

The research project examined whether there is a need for new legislation to be introduced to enable the implementation of the *Virtual Data Embassy Solution*, e.g. a Digital Continuity Act, or whether an amendment of existing rules, e.g. the Emergency Act (2009), would suffice. In particular, the team examined how to define the conditions for the activation of the Estonian digital continuity mechanisms. The challenge of any such legal framework would be to ensure the *de jure* “complete preservation and persistence” of Estonia and its *de facto* functioning to a certain extent in the cloud. This is particularly important given that the data would be located outside Estonian territory.

The Estonian analysis determined that the Emergency Act already sets out the guidelines for behavior in case of an emergency (Article 2), i.e. the government should form a permanent crisis committee. The latter could act as the authority that has the right to activate the mechanisms for digital continuity. However, even with the appropriate authority established, further and more in-depth analysis is warranted on emergency procedures to be put in place. The Emergency Act has clear processes to analyze the likelihood of emergencies and the tools and capabilities needed for an effective response. The Estonian government could engage in an effort to understand the scope and extent of its ICT needs currently being met through the vital services set forth in Article §34, and emerging emergency needs and capabilities, which may not yet be so vital as to require “*the continuous operation of a vital service.... [as] the consistent functioning of the organizer of the vital service and the ability to restore the consistent functioning after an interruption.*” Amendment of the Act to enable this emerging needs capability could benefit Estonia in a number of areas beyond scenarios involving ICT, including medicine, public health, and transportation.

The Emergency Act (Article 7) already lays out the need for an Emergency Response Plan, but the text of the Act is silent as to the scenarios for which a response plan shall be prepared. Some governments have specifically incorporated ICT into their national response plans, recognizing that technology will play an important role in a major emergency. This might be appropriate here, since as many government services in Estonia only exist in digital form it is particularly important that they are restored immediately to minimize any disruption. The Emergency Response Plan, alongside other methods for ensuring continuity, creates qualitatively better opportunities for *de jure* preservation of the Republic of Estonia and also creates certain opportunities for the state to partially function *de facto*.

Ensuring digital continuity requires more than the preservation of critical data sets and ICT solutions on Estonian territory. The need may arise for operating some services outside of Estonia’s borders and the ultimate challenge for digital continuity is to develop a solution where the Estonian state would endure despite a volatile security situation. Simply defining the guidelines for behavior in line with the Emergency Act might be insufficient and the implementation of the Digital Continuity Act could be necessary. If a volatile international security situation were to arise, the Digital Continuity Act would ensure the *de jure* “complete preservation and persistence” of Estonia and its *de facto* if limited functioning in the

Government Cloud. It would also ensure the functioning of the Government Cloud in parallel with the state's regular information systems.

Digital continuity is particularly important for managing data back-up and service functionality under different national security situations. A legislative framework focusing on digital continuity should, for example, define under what circumstances critical data and services need to be operational and backed up from outside of Estonia, and which services need to be operational from outside the country within 24 hours in case there is serious threat to Estonian security. The activation mechanisms to define those conditions should be established in a manner similar to the War-Time National Defense Act. In addition, a procedure needs to also be determined to regulate the moment a secondary site becomes the primary site, and when and if the database is switched back to Estonian physical territory.

The need to act will depend on the level of threat. For instance, a substantial database stored within Estonian borders during peacetime due to its restricted content, might need to be migrated to the Virtual Data Embassy in an emergency because armed conflict would outweigh the risks of international legal uncertainty regarding the status of the Virtual Data Embassy. Three national security environments could initially be used to define data storage and service functionality:

- 1) **Full Control:** Under the full control mode of operations, the Estonian government operates from within its territory in-country and the core ICT operational staff has no constraints with regard to their physical location or logical ability to access computer services.
- 2) **Fragile Control:** Under the fragile control mode of operations, the Estonian government operates from within its territory but the core technical and policy staff may have constraints with regard to their physical location or logical ability to access computer services. For example, this could be due to a significant cyber-attack or a volatile security situation.
- 3) **No Control:** Under the no control mode of operations, the Estonian government operates outside of its territory and it is expected that the core technical and policy staff may have multiple constraints on their physical location or logical ability to access computer services. In this scenario, the Data Embassy must be fully able to support a "failover to cloud" procedure that leaves the properly elected officials in control of the entity and computing resources that represent Estonia's government and society.

In addition to the above, data access and levels of restriction must be considered in the context of digital continuity. For example, a government could choose that certain sensitive information, such as state secret data, not be stored in a privately owned public cloud due to the risk profile of the data. However, if digital continuity is under threat, for instance due to a change in the level of control, such circumstances may necessitate operating government services from the cloud, including those services with sensitive data.

Estonia already operates a central catalog for all of its national information systems ([Administration system for the state information system RIHA](#)) which should be updated to include information about how different security situations affect digital continuity. Moreover, information about the frequency and speed of data backups and the appropriate data center types should be included in the classification. This applies for data as well as services. Such classifications would allow the government to determine which registries and services can utilize privately owned public cloud services.

3.2.3.4. Evolving International Law

The research project's analysis showed that there is a need for broader discussion and agreement on how existing international law enables governments to protect their data when stored in third-party data centers. This new set of policy considerations is an important subset of the broader policy and legal issues

arising from government surveillance. This research should encourage governments to respect sensible limitations on their ability to access user data, and to work together to develop a robust, principled and transparent international framework that resolves conflicts.

In conclusion, the research project found that it is possible to host Estonian government data and services in servers owned and maintained by a private sector cloud provider. However, to enhance adoption, Estonia should consider limited amendments to existing domestic legislation, particularly for government critical services using restricted data. The level of data involved and its interaction with Estonian data protection rules will be essential. So too will be an update of the ISKE regulation to include guidelines for hosting data in the cloud. When it comes to the international environment, it was found that while legal protection exists, countries may have differing interpretations of relevant rules under domestic law, and that such varying interpretations could impact the implementation of international legal principles in any specific case. A further effort to develop an accepted understanding of the international legal framework in this space should be undertaken.

RECOMMENDATIONS

1. Information security standards, ISKE in this case, should be updated, building upon international standards, as appropriate, to address cloud computing.
2. Risk assessments should be conducted to establish acceptable risk tolerance and associated requirements.
3. For the migration of restricted data, countries should consider developing digital continuity legislation and a strategy to increase assurances of diplomatic and other international law protections.

3.3. Technical Research

The research project sought to understand how cloud based application packaging¹⁴ is able to help overcome the environmental dependencies that would have previously prevented the on-premises applications from being moved to either a different physical or online location.

Public cloud platforms provide the first real steps towards abstraction of the physical computing environment, which includes servers, networking equipment, and storage systems, through the use of different cloud application packaging types. Cloud-based application packaging can take the form of specialized installers, containers, virtual disks, or entire virtual machines. Virtual machine and virtual disks can be used as a type of cloud application packaging to virtualize physical host and hardware. They may also be used to capture system specific settings including device drivers, network interfaces and Internet Protocol (IP) addresses, routing paths, DNS settings for both hosts and subdomains, cryptographic keys, user and machine credentials as well as many other application specific settings. In the context of this project it was also proposed that the use of cloud application packaging would enable consolidation and portability of applications, as well as make the operations, maintenance, and application development simpler tasks by mitigating some of the more disruptive elements in application lifecycle management.

Furthermore, it was assumed a major benefit of a *Virtual Data Embassy Solution* would be a consistent (seamless) online environment based on the latest versions of compatible hardware and software. The Solution is expected to eliminate the variability of stand-alone hardware and software environments, which have been developed (and upgraded) at different intervals over time. For example, when each application is free to dictate all layers of the software and hardware stack, there is a real risk that it can be extremely complex to move that application to a new physical, e.g. physical embassy, or cloud environment, e.g. Virtual Data Embassy. The research project looked at how, when using a public cloud platform, significant layers of the application stack could be consolidated, standardized, and scaled in ways that make it possible to move and operate all applications in a more standardized and repeatable fashion.

Finally, it was important to show how the use of a public cloud platform could help optimize for different sets of skill and operational models required by administrators of e-government services, building a critical mass of knowledge overall. To this end, the research project team sought to show how a basic IaaS¹⁵ platform could support two different applications using common operations up to the level of the specific application requirements. While it was assumed that the use of an IaaS platform would not completely eliminate the need for subject matter experts, it was found to reduce the amount of support needed. Moreover, software automation allows resources, previously allocated to support on premise applications, to be redeployed in other areas.

¹⁴ *Application packaging* is a process of binding the relevant files and components to build a customized application for a customer.

¹⁵ Software as a Service (SaaS) is a software delivery business model in which a provider or third party hosts an application and makes it available to customers on a subscription basis. SaaS customers use the software running on the provider's infrastructure on a pay-as-you-go basis. Infrastructure as a Service (IaaS) is similar to traditional hosting, where a business will use the hosted environment as a logical extension of the on-premises datacenter. The servers (physical and virtual) are rented on an as-needed basis, and the IT professionals who manage the infrastructure have full control of the software configuration. Platform as a Service (PaaS) offers hosted application servers that have near-infinite scalability resulting from the large resource pools they rely on. PaaS also offers necessary supporting services like storage, security, integration infrastructure and development tools for a complete platform.

3.3.1. Estonian Government ICT Architecture Overview

Before elaborating on the technical specifics of the research project further, it is important to understand the evolution of the Estonian ICT environment. Over the last 20 years, Estonia achieved tremendous success in combining various technologies with effective, efficient local and central governance. Starting from the ambitious “Tiger Leap” program in 1997, which kick-started Estonia’s e-transformation, and continuing with the development of inter-organizational data exchange layer “X-Road”, Estonia has invested smartly in developing a fully functional e-infrastructure. The existing e-government services allow Estonian citizens to electronically manage a substantial portion of their interactions with the government.

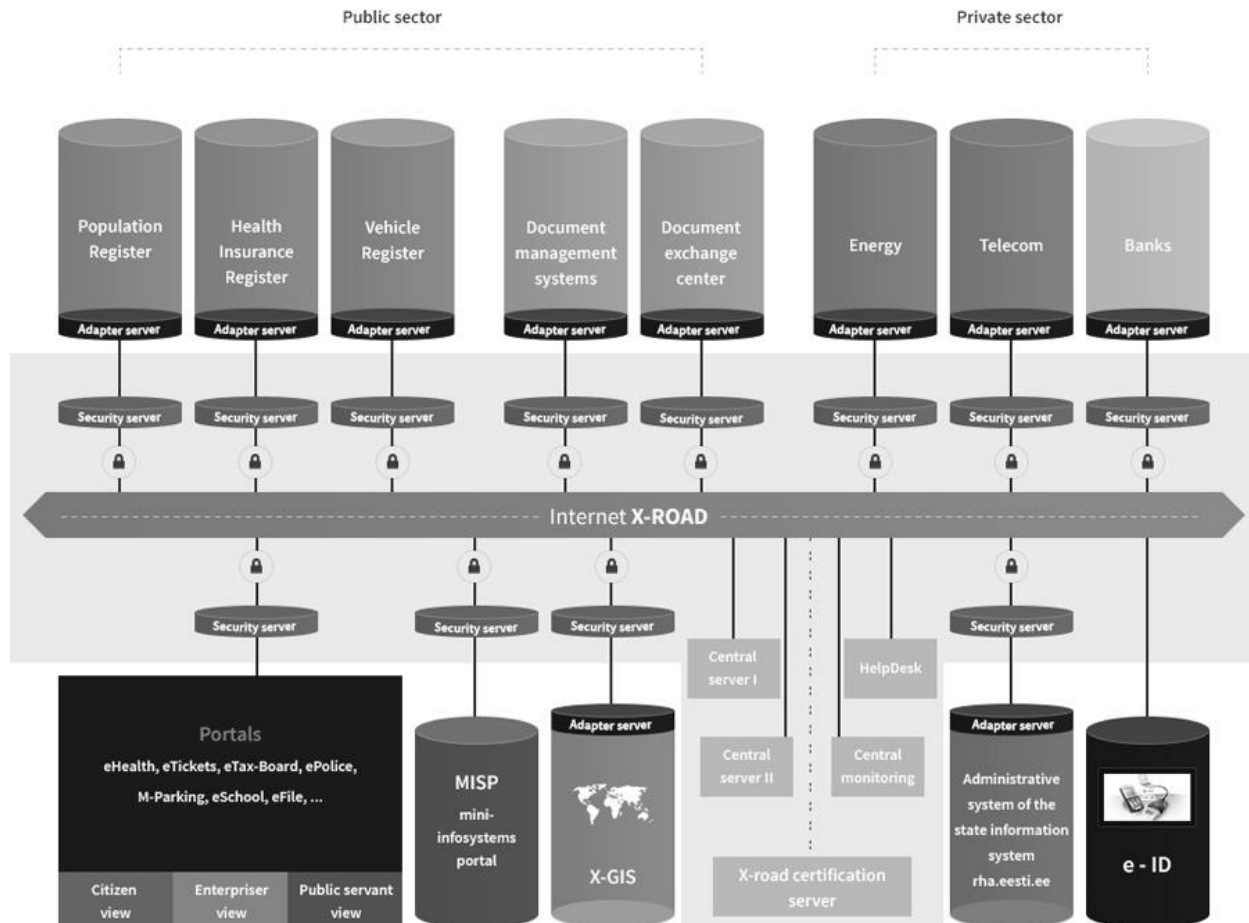


Figure 2: Architecture of the X-road

Currently, the Estonian e-government system is built around a three-tier architecture defined by: Services, which use presentation and application programming interfaces; Data Transport (X-Road); and, Databases. It utilizes a well-designed public key infrastructure, which allows the government to securely encrypt, transport, and store information from a variety of government data registries. The core platform is built around a data transport system, which provides a distributed connectivity platform to prevent any single point of failure and which supports multiple data protocols, including SOAP, XML RPC, LDAP, etc. The system relies on a primary certificate authority for the Data Transport (X-Road), while communication between organizations happens at a peer to peer level. Each major e-government service has a data registry supported by one or more databases that are connected to X-Road via a standard gateway called Security Server and custom service adapters. The Database tier is a vendor agnostic model and supports multiple commercial and open source databases including Oracle, MySQL, and PostgreSQL.

The collective e-government system allows Estonian citizens to utilize a citizen-facing portal to manage a substantial portion of their interactions with the government via electronic means. Estonia has invested considerable resources in extending the capabilities of the current web platform to support digital signatures via multiple web browsers, so that interactions with the aforementioned systems can be authorized and digitally signed. However, while the current e-government framework represents a solid foundation, the government came to the realization that it needs to evolve to ensure it can utilize global cloud computing resources in the implementation of the *Data Embassy Initiative*. To this end, the principles outlined below have been developed to ensure Estonia can keep in step with the fast evolving technology developments, not just for this project, but for the future:

- **Single digital identity:** Each citizen should at birth be assigned a unique identifier that is associated with that person's rights and obligation within the government framework. This identifier should be linked to a secure digital identity that the citizen uses in execution of their rights and obligations.
- **Citizen control of data:** Each citizen should have control over his or her data and should be given information about how it is being used and when it is being accessed. Exceptions should be allowed for criminal investigations.
- **Collect once, use many:** Government agencies should only collect information from the citizen once and be ready and able to share information with other government agencies, when required.
- **Central catalog of ICT systems:** All government information systems should be cataloged and registered centrally. Information should be made public to the extent possible.
- **X-Road usage mandatory:** X-Road should represent the main and only channel of communication among the different agencies.
- **No legacy systems.** The government should be vigilant to ensure that any government ICT systems approaching the end of their lifecycle, e.g. 13 years, are phased out as soon as possible.

The following sections explore how the *Virtual Data Embassy Solution* could be implemented within the context of the current Estonian government ICT architecture. The implementation process and the steps taken are elaborated, followed by a section that compares and contrasts the results of the testing that was conducted, before ending with findings that were considered significant in the process.

3.3.2. Technical Project Overview

At the onset of the research project a number of steps needed to be taken, from which a number of implications for the project as a whole, were derived. While not part of this document, the process of selecting from the different cloud options available and selecting the government services to be migrated was critical to the success of the research. To this end comprehensive risk assessments were conducted, which allowed the selection of the most appropriate services, as well as highlighting a number of opportunities for improvement overall.

Building on this, consideration was given to how to best migrate the services to the cloud. The challenges encountered in this process, as well as solutions used, are presented in the section below. The next section talks about how the different starting architectures across the two project workstreams were worked around in order to ensure that in the cloud operational efficiencies were achieved, e.g. common file storage, common backup procedures, and common load balancing technologies. Lastly, an overview of the testing conducted, once migration had been completed, is presented.

3.3.2.1. Migration to the Public Cloud

The migration of existing e-government services to a public cloud platform, and the feasibility of doing so, represented a major part of the research project. If it had emerged that the migration process was too complex, costly, time consuming or that it required significant architectural changes, doubts would have been cast on the overall viability of the *Virtual Data Embassy Solution*. An assessment was therefore made at the beginning of the project to understand it would be feasible and whether any significant changes would be required.

Two main approaches were considered for migrating the selected government services, although it has to be pointed out that these cannot be seen as binary either/or options. It is expected that in other similar situations a combination of these two would be used, as one is a better fit for newer operating system migration and the other for more complex applications:

- **Virtualize the application and perform an “in-place” base operating system upgrade:** This approach consists of three distinct steps: 1) any physical servers are virtualized and any existing virtual servers cloned; 2) the operating system is upgraded to the latest version in-place; and, 3) all is uploaded onto the cloud platform. Typically this approach is beneficial if there is a need to maintain a direct clone of the entire operating system environment on-premises. However this approach requires significantly more time and bandwidth since the images to be uploaded can be large, e.g. over 50GB.
- **Deploy directly onto the cloud platform:** This approach sees the base operating system provisioned directly by the cloud platform before any application components, such as databases or web servers, are deployed onto the operating system and the application content and data is synchronized. This approach has the benefit of requiring only the core application and its data to be uploaded to the cloud which frequently represents a much smaller data footprint.

For the migration to be possible, it had to be established whether the Microsoft cloud platform supported the operating systems currently used by the government services and their applications. As highlighted in the research project description, the two government services use FreeBSD and CentOS, which are not Microsoft products. Importantly, the cloud platform selected supports a number of non-Microsoft products, including UNIX and Linux operating systems, as well as FreeBSD and CentOS. However, it typically only supports the latest three versions of an operating system. This means that an operating system upgrade was required prior to migration, as the versions used by the selected applications were not supported. Furthermore, on-premises virtual machines can also be transferred onto the Microsoft’s cloud platform directly, as it supports the open industry-standard virtual hard disk (VHD) format.¹⁶ This is used by a number of on-premises hypervisors.¹⁷

Given the base operating system upgrades required by both application workstreams, the second option, “deploy directly in the cloud”, was selected as being quicker and less risky. One reason for this choice was that the software installation media for the application components was readily available, making it straightforward and fast to re-install them on the public cloud operating system. The second reason was that this option meant the content and data could be synchronized directly with the on-premises master, again with speed and ease. Finally, and most critically, the versions of FreeBSD and CentOS used originally

¹⁶ About the VHD format: <http://msdn.microsoft.com/en-gb/library/windows/desktop/dd323654%28v=vs.85%29.aspx>

¹⁷ A hypervisor or virtual machine monitor is a piece of computer software, firmware or hardware that creates and runs virtual machines.

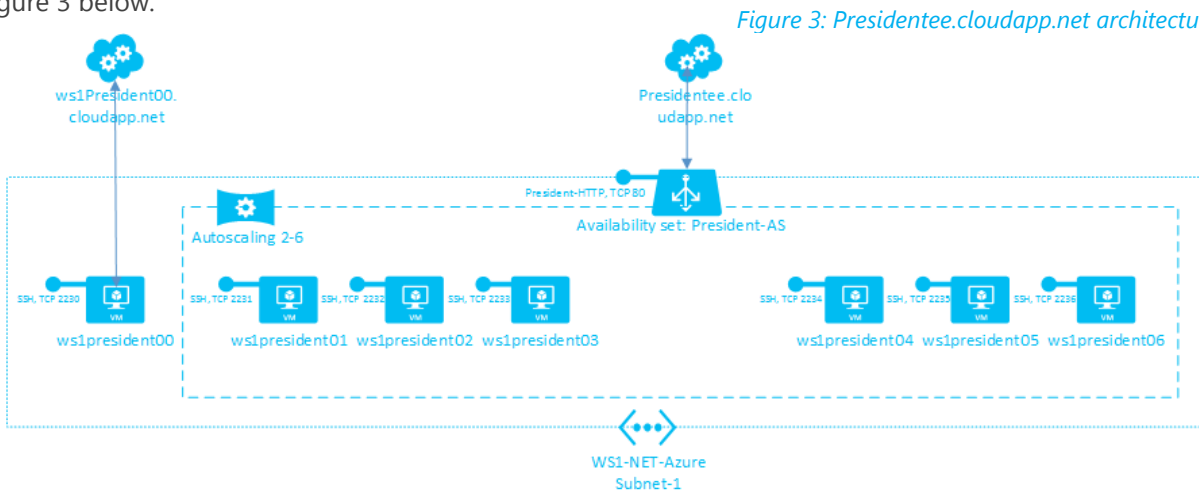
did not natively support an in-place operating systems upgrade. As a result, the in-place upgrade, core to the first option, would have to be performed manually, introducing a higher level of risk.

3.3.2.2. Target Architecture for Public Cloud

While the two project workstreams had different starting architectures, the research project aimed to deploy the applications in a similar deployment architecture to help achieve operational efficiencies, such as common file storage, backup procedures, and load balancing technologies. To this end, target deployment architectures were utilized as set out in sections 3.2.2.2 and 3.2.2.3.

Workstream #1 – Presidentee.cloudapp.net

Presidentee.cloudapp.net was a Microsoft cloud application version of www.president.ee site. The cloud service “presidentee.cloudapp.net” represented six virtual machine instances behind a load balancer configured in an availability set. The cloud service had been configured to use auto-scaling with minimum of two instances to a maximum of six instances. It is important to point out that two instances represent a minimum required to be able to achieve the level of availability needed. The overall design is depicted in Figure 3 below.



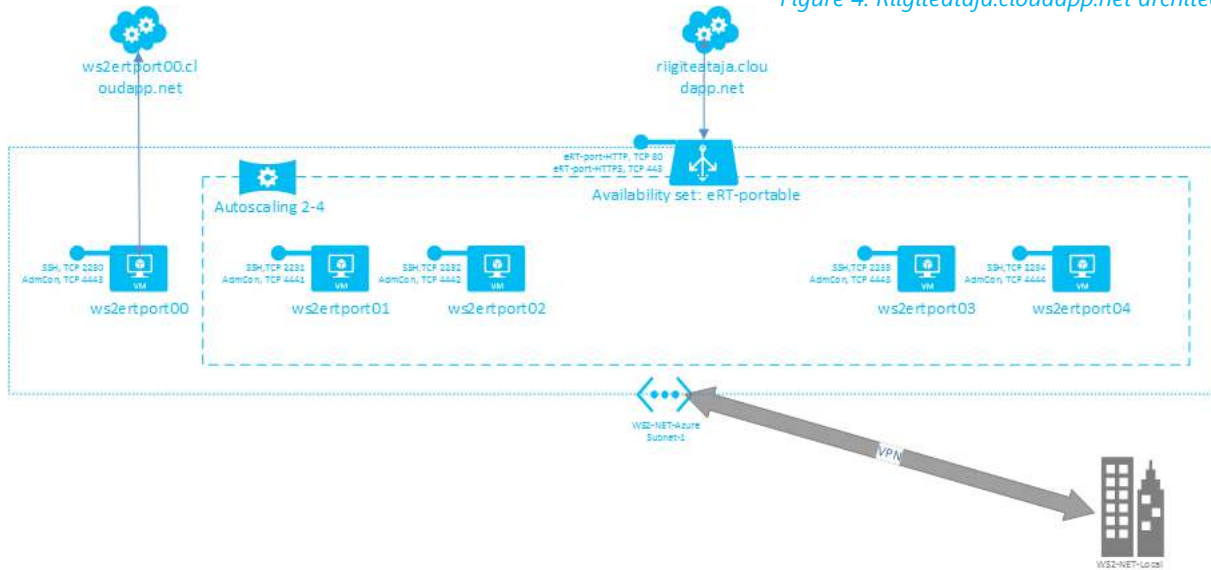
The cloud service had Hyper Text Transfer Protocol (HTTP) protocol on Transmission Control Protocol (TCP) port 80 published on the internet. The load balancer ensures that the network traffic is equally divided between all running virtual machine instances. Moreover, when a new virtual machine instance starts or stops, it is automatically included in the load balancer pool. As mentioned above, the cloud service was serviced by up to six virtual machines kept on geographically replicated storage and configured as an availability set. For the purposes of the research project, this meant in practice that data was replicated between the Azure™ Amsterdam and Dublin datacenters and that the virtual machines continuously ran on different update domains to prevent any cloud service downtime, such as that which might occur because of the cloud platform’s fabric maintenance.

All virtual machines used were sized Standard A1 (1 core, 1.75 GB memory), had a fixed IP reservation and a single disk sized 50 GB. The IP address was reserved in the WS1-NET-Azure Subnet 1 IP address segment. As a result, the cloud platform fabric always assigned the same IP address to the virtual machine. Finally, the content was synchronized from on-premises publishing servers using the custom rsync protocol over the Secure Shell (SSH) to ws1president00 server and, in the second stage, from ws1president00 to ws1president[01-06].

Workstream #2 – Riigiteataja.cloudapp.net

Riigiteataja.cloudapp.net was the portable version of www.riigiteataja.ee site running on the Microsoft cloud platform. The cloud service “riigiteataja.cloudapp.net” represented four virtual machine instances behind a load balancer configured in an availability set. The cloud service had been configured to use auto-scaling with minimum of two instances to a maximum of four instances. It is important to point out that two instances represent a minimum required to be able to achieve the level of availability needed. The overall design is depicted in Figure 4 below:

Figure 4: Riigiteataja.cloudapp.net architecture



The cloud service had both Hyper Text Transfer Protocol and secure (HTTP, HTTPS) protocol on Transmission Control Protocol (TCP) port 80 and 443 published on the Internet. The load balancer ensures that network traffic is equally divided between all running virtual machine instances. When a new virtual machine instance starts or stops, it is automatically included in the load balancer pool. As mentioned above, the cloud service was serviced by up to four virtual machines stored on geographically replicated storage and configured as an availability set. For the purposes of the research project, this has meant that in practice data was replicated between the Azure™ Amsterdam and Dublin datacenters and that the virtual machines continuously ran on different update domains to prevent any cloud service downtime, which might occur because of the cloud platform’s fabric maintenance.

All virtual machines used were sized Standard A5 (2 cores, 14 GB memory), had a fixed IP reservation and two disk drives (OS disk sized 30 GB, data disk sized 200 GB). The IP address was reserved in the WS2-NET-Azure Subnet 1 IP address segment. As a result, the cloud platform fabric always assigned the same IP address to the virtual machine. For synchronization purposes, a site to site Virtual Private Network (VPN) was also created between the WS2-NET-Azure virtual network and the network of the Ministry of Justice. Moreover, the cloud service riigiteataja.cloudapp.net was also assigned a reserved public IP address to make external DNS handling easier. Finally, the content was synchronized from on-premises publishing servers using custom rsync protocol over SSH to ws2ertport00 server and in the second stage from ws2ertport00 to ws2ertport[01-04].

3.3.2.3. Performance Testing Methodology

Once the migration of the two applications was completed, testing activities began across two dimensions: performance and demand. Two types of performance tests were used: load and stress testing. Load testing was used to understand the behavior of the system under a range of normal load conditions, and to compare the transactional response time with the on-premises solution. Stress testing, on the other hand, was used to determine the solution’s robustness under peak load and to prove it would automatically scale-out elastically under sustained peak load situations.

Load testing involves simulated client and end user activity that could take place, if a large number of human end users were attempting to access the services at one time. The patterns and usage scenarios are designed to ensure that the correct activities are available to users. For example, if one million users were to go to the president’s website due to an external event, this would be considered a load test scenario. The users are not trying to do unusual but their sheer numbers could impact performance. Conversely, the stress testing scenarios tend to simulate client and end user behavior designed to break or cause problems for the site, e.g. where multiple users try and overload it by playing videos to attempt to consume all the resources available, thus preventing the site from functioning. The load and stress testing were also repeated under two demand scenarios. The first was a normal demand scenario, which consisted of replicating the existing normal usage conditions whilst catering to organic expanded demand, as might occur if additional users were to utilize the e-government service. The second was a malicious demand scenario, which might occur if Estonian e-government services were under a cyber-attack intended to render the services incapacitated or unavailable. Section 3.3.4 covers the results of the testing.

Normal Demand Usage Scenario

The normal demand usage scenario was to demonstrate the cloud platform’s ability to dynamically scale up the number of application servers, storage systems, and route traffic appropriate to the end user demand. For example, if a text based and media content update to the President.ee website were to occur, load testing should show that the cloud platform has supported increased demand for the media files and web server content. In an ideal situation, the cloud platform would dynamically scale up compute, storage, and network resources using the cloud platform load balancer to deliver the content to users via the closest Virtual Data Embassy, irrespective of where the data center is based.

Under normal load testing, it was expected that the cloud platform could demonstrate automatic scaling, eliminating the need for procurement, setup, or redeployment of applications by the administrator. The normal load testing scenario was also designed to show the reliability of the cloud platform under normal failure events, such as a (non-malicious) application crash due to a software bug in either the application or underlying guest operating system. If such an instance were to occur, the cloud platform should automatically use a replacement application instance with no staff involvement.

Malicious Demand Usage Scenario

The malicious load and stress testing was to demonstrate that the cloud platform is resilient to malicious attempts to consume compute, storage, or network resources, which would prevent a normal end user from accessing the application or monument websites within a reasonable response time. Under malicious load and stress testing, the cloud platform was expected to implement parameterized automatic scaling, which would eliminate the need for operator involvement or the procurement, setup, or redeployment of applications by the administrator.

The simulated malicious load testing scenario was designed to demonstrate the reliability of the cloud platform against malicious attacks that would attempt to exploit known (e.g. Heartbleed SSL) or unknown (e.g. zero day attacks) software bugs in either the application or underlying guest operating system. In a malicious failure scenario under load, the cloud platform should automatically use a replacement application instance with no operations staff involvement and begin to report malicious usage to key staff people for examination of possible mitigation techniques beyond simple load balancing.

Table 6: Normal and malicious demand scenario testing

3.3.3. Technical Architecture Findings

The research project demonstrated that the standard Microsoft public cloud platform is able to host and run existing government applications fully across the two workstreams, confirming the first hypothesis (see Table 4 previously). The research project also confirmed that the performance, stability and security of the services could be enhanced using built-in capabilities of the Microsoft public cloud. The second hypothesis was also confirmed early on: while only minimal architectural changes were required, the team has amassed a number of valuable detailed findings, which have been broken into six major architectural and operational areas in the sections that follow. It also became clear that no matter what, textbook readiness is impossible to achieve.

- Security, Identity, & Data Findings;
- Operational Findings;
- Application Findings;
- Compute Findings;
- Storage Findings;
- Network Findings.

While the specific, detailed findings are elaborated below, one of the most critical findings was around the level of dependency that the government and users have on the correct operation of core internet systems, such as DNS. Today, the DNS and the associated public key cryptographic systems are essential to ensuring that users get to the right location, when using a web browser or another Internet application. DNS is the core system that translates the hostname component of a Uniform Resource Identified (URI), such as <http://www.president.ee/et/>, to a corresponding IP address, in this case “194.204.33.69”. In the process, a number of critical intermediate servers and systems are vulnerable to attack and therefore the assertion of ownership over these DNS servers and name authorities is pivotal. If the common means of accessing a government service (via DNS and URI) misdirects the user to a rogue site, it is irrelevant how secure the application or service running in the cloud is.

It also emerged that in order to fully address the security implications, it had to be better understood how specific government domains, e.g. “.ee” domain, could be further protected by standards developed by the [Internet Engineering Task Force](#) (IETF) and by organizations, such as the [Internet Assigned Number Authority](#) (IANA), and the [Internet Corporation for Assigned Names and Numbers](#) (ICANN). As it stands, currently nearly all the main root servers reside in the United States. While this has some benefits for the Virtual Data Embassy, there is still a need to address other government agencies to determine proper means of asserting ownership over top level domains like “.ee”.

Another major finding was the need to have a fully functioning e-Identity system running in the cloud both prior and during any failover¹⁸ to cloud scenarios. For example, for the successful implementation of the switchover from running the President.ee website on-premises to in the public cloud, a number of operational, policy and technical procedures need to be defined and followed to ensure an authentic and orderly transition (see domestic policy and legislative findings above). These procedures are critical to

¹⁸ Failover is switching to a redundant or standby computer server, system, hardware component or network upon the failure or abnormal termination of the previously active application, server, system, hardware component, or network.

avoiding human error or malicious behavior by rogue insiders. Given the important role e-Identity system plays in Estonia ICT infrastructure, this would necessarily involve authorization from both government and technological perspective through the authentication and signing of key documents and procedural elements. The e-Identity system must be able to run from outside the Estonian borders; i.e. in the public cloud, for that to be possible under all circumstances.

It is also worth highlighting the application development process as a unique area for additional study. During the research project, Microsoft was able to apply general principles that are utilized in the Security Development Lifecycle (SDL)¹⁹ for this application design. In an ideal SDL scenario threat modeling, security and performance are all part of the design process from the onset and become instrumental to making certain architectural trade-offs. This was not possible in this case due to the constraints of the research project but the overarching principles were used to flag areas where increased security, robustness, or performance improvements could be made to the existing applications.

Finally, the Project team identified numerous areas for improvement across the organizational components of the work, e.g. defined procedures for running services in cloud, clear division of responsibilities between the vendor and the government. This confirmed the third hypothesis being tested. It is important to note that the architecture of the information systems was poorly documented, a concern that needs to be addressed to ensure digital continuity for Estonia. This is especially important when it comes to disaster recovery plans, as secure and well-established operational procedures for disaster recovery, failover authorization, and migration activities depend on formal documentation, testing and planning.

RECOMMENDATIONS

1. Cloud computing should be utilized to increase security and resilience of government infrastructure.
2. An overarching cloud strategy and government action plan facilitating cloud migration should both be developed to enable technical and operational agility and increase cost-effectiveness.

3.3.3.1. Security, Identity and Data Architecture Findings

The findings in this section have been arranged to put some of the most significant and urgent issues first. While this should not be taken as a reason to ignore or deprioritize the deeper technical findings, it is critical to address some of the broader issues which will impact not only these two government services but also all others.

Foremost on this prioritized list are the areas of security, identity, and data architecture, which are all deeply linked. For example, consider a given data item, e.g. a credit card number, birth certificate or medical record, which needs to be kept very secure. To protect this information, it is important to understand who is attempting to access it, and what that individual's authority or access rights are. The ability to ensure that all logs, records, and access can be associated with the correct individual, not by someone fraudulently impersonating them, is pivotal. An effective security and e-identity architecture is therefore critical to a well-functioning data governance system.

¹⁹ <http://www.microsoft.com/security/sdl/default.aspx>

Data governance and security rest on well-defined data classification, data handling and access policy, data lifetime management, data auditing, and access auditing tools. While Estonia does have in place a policy around handling secret and sensitive data, the research project showed that there is a need for a broader data classification policy that extends across the spectrum of high, medium, and low “business impact” (HBI, MBI, LBI), and that adapts to the country’s digital continuity needs, which may change over time. Such an approach would help the government identify which applications and services could be migrated to the cloud in their current form and which would need to be modified or not moved at all. For example, an application might need to be refactored so that one or two data items in a database require specific encryption and handling, while all other fields are able to be stored and processed without additional encryption.

While the President.ee site is an open site for citizens and general internet users, the administration of the State Gazette requires the usage of a Smart ID-card. The migration of services such as the State Gazette should enable government employees to publish new legislative acts even when the service is running in the cloud. However it currently depends on services that are presently only based on-premises, such as:

- Time stamping service. Every new published legislative act and its previous iterations need to be stamped digitally with timestamp.
- Digital signing service. Every new published legislative act and its previous iterations need to be digitally signed.
- Using Document Exchange Center (DEC) via X-road to submit the legislative acts. The State Gazette application checks after a determined interval if new legislative act have been submitted.

Without time stamping and digital signatures, legislative acts can’t be published in the State Gazette. As a result, these services, including authentication with ID-card, downloading legislative acts from the DEC, digital stamping, time stamping, must also be moved to the cloud, if the Estonian government is to ensure digital continuity.

Other findings by the project team related to security, identity and data architecture included:

1. Data architecture and data security policies are essential for data integrity

Data integrity is of utmost importance for e-government systems, as it is central to security of and trust in the government services. To ensure data integrity the following elements are needed:

- A well-defined and change-controlled data architecture, which contains a data model describing the data entities, their attributes and other metadata, as well as the relationships between entities;
- Change-controlled datasets, which record any change within the individual data repositories and are labelled with a data architecture version;
- A data security policy, which defines: the rules to be enforced with regards to any change in data entity state; who can perform certain operations; and, when these changes can be made.

These elements allow for data integrity validation routines, which can then be applied to each new dataset and help confirm the integrity of the data. If the data change is deemed invalid, the last known uncorrupted state could be restored and an exception routine triggered. This approach ensures that data integrity is monitored and maintained at all times against a defined data architecture baseline and security policy.

2. A holistic data governance and security approach is required to facilitate migration to the cloud

A data governance approach describes the methodology for management and control of data and data related assets within the e-government systems. It ensures that data is categorized in terms of business impact, for example across High, Medium, and Low Business Impact (HBI, MBI, and LBI). Such a categorization is key to determining how the details of the data governance approach and related security policies are applied. It ensures consistency and allows data entities to be used across multiple applications, which is particularly important for sensitive personal data, such as birth date, name, etc.

The data governance approach can also be used as part of the cloud migration strategy. In particular, the business impact classifications are a key parameter in determining which systems, components, and data are appropriate for migration to the cloud, and in defining the transition roadmap. One possible approach is to have a data governance model that allows for LBI and MBI data and systems to run in the cloud under full control situations, however in a fragile or no control situation, even HBI data and systems could be run in the cloud.

3. Designing systems for 'separation of duties' and 'least privilege' is vital to maintaining security

Effective system security balances allowing user access to the data and systems with minimizing the risk of misappropriation or corruption of data. A generally accepted security principle for minimizing this security risk is described by the National Institute of Standards and Technology in its [Special publication 800-14 Generally Accepted Principles and Practices for Securing Information Technology Systems](#), which focuses on *separation of duties* and *least privilege*.

Separation of duties refers to dividing roles and responsibilities so that a single individual cannot subvert a critical process, e.g. in banking no single individual is given authority to issue checks under normal circumstances, rather, one person initiates a request for a payment and another authorizes it. *Least privilege* refers to the security objective of granting users only such access as they need to perform their official duties. Data entry clerks, for example, may not have any need to run analysis reports for their database. This can be achieved by designing a multi-layered security model, where access is restricted to the minimum privileges by function and component for each role. In addition, this facilitates management of the individual user to role mapping, which further reduces the security risk of exploits, which could arise from users changing roles or leaving the government team.

4. Role Based Access Control (RBAC) and claims-based security increases overall security

Access to Estonian e-government services today is primarily through individual accounts, which rely on shared secrets and traditional password based security for both user authentication and secure data transfer. As a result, if an account with high-level privileges is compromised, for example an operating system root user account or data replication account, potentially the entire system might be at risk.

Using the RBAC approach with claims-based security would dramatically decrease the impact of an individual user's credentials being captured or compromised. Utilizing security groups where individual users are members, rather than individual accounts, which may be shared between users of the system, is much more secure. Among other things, it allows access to be controlled in detail at the user level, which has the additional benefit of allowing role permissions to be applied at the security group level and then automatically filtered down to the user accounts which are members of that security group. It also creates the ability to revoke access for an individual user without affecting other users.

5. Overall system control increases if operating system root credentials are restricted

Retaining control over the overall system is essential to security. The latter is optimized if the operating system root credentials are not shared with all administrators by default, but are instead restricted to a small sub-set of people performing administrator activities. Through the RBAC approach described above, the collection of administrator-level privileges can be decomposed and grouped into its

constituent roles, and appropriate levels of more restricted permissions applied accordingly. For example, the “application administrators” security group could enjoy a lower level of privilege than the operating system root account group to enable better control and avoid escalation of privilege attacks.

6. Isolation of user roles and accounts between environments prevents accidental changes

Applications should have well established operating environments, such as development, test, staging, and production. For each of these environments a set of user roles should be created, which only have permission to the target environment. User accounts can and should be mapped to the target operating environments on an “as needed” basis to prevent unauthorized or accidental access and changes. Typically, there are user roles that apply to all environments, e.g. development, test, staging and production, or to a sub-set of environments. Additionally, it would be advisable to have a subset of user accounts authorized for the test user role on the staging, than on the development and test environments. This approach also improves the ability of a system to be audited and have dynamic security role assignment. For example, a developer may be granted temporary access to the {Production Env: Developer Role} to troubleshoot an issue, then have that privilege revoked.

As an additional example, a “PortalDeveloper” role would manifest as a “PortalDeveloperDev” and a “PortalDeveloperTest” user groups, which would each be given separate privileges to access the relevant portal components and systems for the development and test environments. If user “John Smith” is mapped to the PortalDeveloper role and requires access to both the development and test environments, then his user account would be added to both the PortalDeveloperDev and PortalDeveloperTest user groups. A good way to manage this is to maintain a matrix, which maps individuals to roles and then roles to the relevant environments.

7. Isolating service accounts between application instances reduces security risk

Currently within the Estonian e-government systems, multiple instances of the same application are using the same service accounts, i.e. accounts specifically created to isolate a service or application. This presents an increased security risk because the fact they share credential means that if one application instance was to be compromised, all systems of the target application could also be compromised. To reduce the risk, per application instance service accounts should be created, so that there is a one-to-one relationship between a physical service account and an application instance. This will ensure that if a service account is compromised, access is sand-boxed to only one application instance rather than allowing the attacker to break-out across all application instances.

3.3.3.2. Operations Architecture Findings

During the research project, it became clear that many of the activities required to move the selected application and e-government services to the public cloud required extensive manual operations in order to be successful. This included steps as important as opening a text editor and manually keying in a new IP address for a DNS entry to redirect users to the cloud based implementation of an e-government service. If operations such as this are not flawlessly executed in a timely fashion, all other efforts may be for naught. Furthermore, it was discovered just how important the documentation and implementation of standard operational procedures are to repeatability under duress or other stressful situations (e.g. transition to fragile control). Most of the procedures required for the delivery of highly available solutions above a certain level, e.g. 99.9% availability, have proved to be significantly dependent on human activity and skill.

Therefore to “failover to cloud” successfully it is essential that operational procedures are designed and tested to ensure the ability to execute a defined procedure correctly and efficiently. Well defined and

tested operational procedures can often mean the difference between failure and success. Once information systems have been developed, tested and transitioned into production, e.g. “live,” operational procedures become essential to successful service delivery.

Key detailed findings related to service operations included:

1. Documented disaster recovery and cloud fail-over procedures ease pressure on teams

The event which triggers the disaster recovery process is usually highly stressful, as there will be extreme pressure on the operational team to restore the affected service as quickly as possible. The current Estonian e-government disaster recovery procedures are not fully documented nor have they been tested, a challenge likely to delay the execution of any disaster recovery failover. The following actions would be essential for a smooth, rapid, and low-risk execution of the disaster recovery process:

- Creation of a disaster recovery and failover design template, describing which systems failover, where they failover to, and the technical mechanism for the failover;
- Creation of fully documented disaster recovery procedures;
- Documenting roles and responsibilities for managing the disaster recovery procedures;
- Executing periodically planned disaster recovery tests to rehearse the procedures.

Moreover, once the above elements are in place, the technical mechanisms and procedures can be assessed for automation opportunities. For example, once the procedure and authorizations for switching a DNS entry to point to the new cloud location is defined, the procedure could be automated through additional software, scripting, etc. Such automation would further reduce the risk footprint. For example, currently the State Gazette has a backup site working in parallel with the production site. The backup site is in a different location and data there is automatically updated. In case of interruptions or failures of the production site, the State Gazette users are directed to the backup site. The backup site is read-only and it does not allow for publication. In the future the backup site could be operated from the cloud, which would allow it to support the full functionality of the State Gazette. However, this approach requires further analysis and an evolution of the current State Gazette environment to make it more compatible with the cloud.

2. Expected load characteristics and required capacity plans needed to be able to scale the cloud

A key benefit of cloud computing is its capacity to handle an increased load in comparison to on-premises solutions. Indeed, the cloud platforms can be configured to automatically scale up or down specific application instances when certain thresholds are triggered.

During the course of the research project, it became clear that to determine the most effective auto-scale settings to apply, more research is required. For example, when streaming media, there are fewer network connections, yet more network and disk Input/Output (IO) traffic, but when serving up web page content, there are a higher number of connections, but lower IO per session. Thus, just triggering the scale up or scale down on network connections alone would not address scenarios where serving media was bottlenecked by central processing unit (CPU) or IO load. Particularly beneficial would be a better understanding of the e-government systems’ current infrastructure load characteristics in terms of the CPU, memory, and disk consumption over time. In addition, an understanding of the capacity usage would be useful to be able to forecast how the load characteristics are likely to evolve.

3. Standardization across e-government services can ease operational management

Standardization and unification can help reduce operational management costs. For example, the standardization of the architectural landscape can help reduce the overall quantity and complexity of

operational support needed. It can also lower the infrastructure platform costs, for instance as a result of a reduced datacenter footprint. Moreover, unifying and standardizing operational procedures themselves also reduces the amount of effort needed to support the systems. This in turn reduces the overall operational support risk because fewer people and fewer diverse systems are involved. Consequently, this standardization has a positive impact on the security and availability of the systems, since both people and systems that can be utilized across multiple different e-government services. Finally, this approach also provides a sound foundation for migrating government services to the cloud by normalizing the on-premises environment prior to cloud migration.

4. Operation support roles benefit from use of RBAC

The use of RBAC, as described in the previous section, controls access to system functions and data. It can also be applied to controlling access for operational support of the systems, similarly reducing risk, for example across operational support activities, such as restricting who can change DNS, IP addresses, and security permissions.

Following from that, the research project found that operational role definition should be driven by and aligned to the operational procedures. The operational roles would then manifest through physical security groups within the system, split by environment, with the appropriate user account membership. By doing so, the risk of an unauthorized user making changes to critical roles is mitigated.

RECOMMENDATIONS

1. Governments should keep data governance and security models up to date across the data lifecycle and required of government entities.
2. The trusted relationship between the government, private sector company, host country, and the country where the provider is headquartered should be deepened for success of the Initiative.
3. Operational procedures should be prepared and tested in advance rather than in a crisis.

3.3.3.3. Application Architecture Findings

The application findings in this section are derived from specific issues encountered across the two workstreams despite them having different architectures and primary purposes. While the sections that follow attempt to highlight technology specific elements related to core compute, storage, and networking, this section will highlight findings that are derived from the two applications themselves.

One of the most critical observations regarding the application findings was the dependency that each Estonian e-government service has on other key services. As shown below, a mesh network of service dependencies emerges, in part due to the mandate that the government only collect information from a citizen once. This single master data model means that services inherently depend on one another, when any data is required that has already been collected, e.g. home address.

Another key observation of the project team was that the versions of the application development framework software used by the State Gazette were older and therefore not compatible with cloud technology. As a result, it was important to upgrade it both for cloud compatibility and to ensure that the most secure and supported versions were used. It is important that in the future, the application development framework software is regularly updated to ensure the State Gazette's reliability and readiness to move to the cloud. This would require the team to be familiar with the requirements of the

different cloud environments and upgrade cycles of the framework software. Additionally, this knowledge might lead to a change in the structure of the existing virtual machines, e.g. separate OS zone from other zones, as described in the compute findings section.




Figure 5: X-Road Service Dependency Map for single data master

Key findings related to the applications include:

1. A DevOps approach to application building ensures a quicker response to threats

Most applications that have been developed prior to the broad deployment of cloud technologies, including the ones examined as part of the research project, do not follow a “DevOps” approach. The latter can be described as a software development method stressing communication, collaboration and integration between software developers and ICT professionals. Instead, the applications have been developed following the “waterfall methodology”, where progress is seen as flowing steadily downwards through the different development phases. This methodology does not appropriately incorporate operational and development activities, an omission that can become a challenge when applications need to be moved, updated, or fixed in time critical situations such as a cyber-attack.

The use of the collaborative DevOps approach is fundamental to making sure that the application developers and operations teams are working together to prevent disruption. It ensures a quicker response to cyber threats.

2. Understanding security threats for all applications is central to successful mitigation

Threat mitigation can take many forms, including changes to the application, systems or infrastructure, improved access control, or modifications to security and/or operational procedures. However, it is important that security features are built into the applications from the start to ensure their greatest effectiveness. To enable this, potential security threats for each application and use-case must be understood. This is referred to as “designing for availability.”

Availability design is fundamentally different from how the traditional on-premises applications were thought about, where the assumption had been made that the infrastructure is resilient. Given that the Internet, central to cloud transition, is inherently unreliable, the e-government applications examined as part of the research project should be re-designed based on the central premise that the on-premises infrastructure (across compute, storage, and network) will at some point fail. This would ensure that applications are resilient to infrastructure node loss by using techniques, such as the central persistence of state, which would allow any failed nodes to simply be re-instantiated.

It is also worth noting that the public cloud provides for application health monitoring, remediation and reporting, which allows failures to be detected and faulty nodes replaced automatically. Therefore, one of the significant benefits of cloud computing is the relatively short time required to repair any faults. Applications may still fail, but if failure can be detected and auto-repaired, the impact on availability is likely to be negligible.

3. A modern design, which allows for portability, gives the government more choice

Organizational needs change over time. Consequently applications may initially be required, for example, to run on-premises within a private cloud and at some point move to the public cloud and then potentially back again. Portability of applications across the different cloud types (public, private, service provider) should be incorporated into the design of government applications from the outset. These so called “modern” applications can run on different cloud types and can provide the government with the flexibility to respond rapidly to changing demands.

One of the key findings from the State Gazette team was that the preparations for moving and operating the State Gazette in the cloud took several days the first time but that later copying on and preparing the pre-set machines took significantly less time. It is important to note that the source machine must be closed (not public) in the process.

4. Public cloud can protect against DDoS attacks better than on-premises systems

Cloud computing can provide different functionalities, such as auto-scaling and geographically spread application instances, which can help mitigating DDoS attacks on the application level. To improve DDoS resilience, the solution would need to incorporate design changes to better utilize cloud capabilities to increase absorb capability and decrease detect time. While on-premises systems can be resilient to certain forms of application level attacks (layers 4 through 7), public cloud providers can address network level attacks at layers 2 through 3 that would typically cripple or disable most traditional ISP's or individual datacenters. Note that application level DDoS attacks were not tested as part of the research project.

3.3.3.4. Compute Architecture Findings

The workloads across both workstreams were running on older versions of different operating systems, FreeBSD and CentOS, with Workstream #1 running on physical servers, and Workstream #2 running in a virtual environment. Given that support for a hypervisor running in Azure™ only extends to recent versions of the operating system kernels, both workstream operating systems needed to be upgraded before they were able to be migrated to the cloud platform.

Key findings related to compute architecture included:

1. Standardized sizing of physical servers and virtual machines helps migration

All virtual machines in the cloud consist of standard-sized components, whether they are CPU, storage, memory, configuration, or operating system version. Standardization with one or two types of server hardware results in the advantage of having common hardware parts, technical familiarity, common firmware images and upgrade techniques, easier management and special pricing from vendors. Also, utilizing standardized infrastructure units, be it physical or virtual, helps reduce the operational costs and risk associated with maintaining a complex ICT landscape, and makes automation easier.

Furthermore, standardization is needed for virtual machines and hence application portability. For instance, live migration of virtual machines works only across the same family of processors from the same vendor (Intel or AMD), so uniformity of CPUs for virtualization host platforms is also necessary. In a similar way, standardization of virtual machine templates enables seamless porting of applications between clouds (private, service provider, public).

2. Master operating system images can drive standardization through virtual machine templates

Master operating system images can be used to drive environmental standardization. An operating system instance can be deployed in the cloud environment using templates, e.g. a master operating system image. Usage of this concept both in the cloud and/or on-premises can greatly contribute to environment standardization, which in turn increases security and supportability, whilst simultaneously decreasing the cost of maintenance and operations. This principle was used when building solution on Azure™ as part of this research project.

3. Complexity reduced if distinct functions are separated onto different compute nodes

If individual compute nodes need to perform multiple functions, it becomes much more complex to scale them and it means multiple functions are tightly coupled together, potentially impacting availability. These challenges arise because different functional roles in the cloud environment carry with them different resource requirements. If distinct functions are separated onto different compute nodes complexity is greatly reduced. An example of this is separation of media servers from static content servers.

4. Isolated content roles minimize risk

As per security best practices, isolation of any content administration and content publishing roles into different servers minimizes the possibility of content modification and rogue content publishing. While certain administrative users may have access to all site content, the research project found that in order to prevent unauthorized publication or distribution of content, different types of content and media should be handled by different user roles. For example, the text based content for the President.ee site, which may include news or speech material, should be considered at different security levels. Additionally, content, such as video footage originating from the government rather than third parties, should be identified and handled appropriately. A lower level authorization and user role could be made available to cite/reference third party publications, while only high level authorized users would have the ability to publish materials that are directly authorized by the president.

5. Cloud design patterns using small scale-out units are preferable

Achieving acceptable availability and performance in the cloud follows patterns different to those in traditional on-premises environments. For example, the cloud has a built-in load balancing capability and the default performance and availability pattern is "scale out".

Historically, the on-premises architecture was designed for the maximum availability/performance scenario with a considerable buffer incorporated. In reality, for the vast the majority of the time, the solution never got close to this limit, which resulted in a large percentage of unused capacity. For

cloud hosting, the solution is built for the minimum load, yet architected to scale-out horizontally as needed. This is achieved by building small stateless servers that can be added or removed on an hour by hour basis, resulting in much higher utilization levels, and therefore significant efficiency savings.

6. Auto-scaling configuration requires further development

Auto-scaling is an important cloud functionality, which allows the cloud infrastructure to add additional resources to the service at peak times and conversely shut down resource usage when it is not needed. For an understanding of how a particular application behaves under stress, the input/output patterns and CPU requirements are vital to any effective configuration.

During the research project's tests it was found that the response of the scale-up operation was not within requested boundaries (20 minutes) for non-Windows virtual machines. Moreover, while copying and scaling virtual machines in the cloud is typically supposed to be easy and fast, it worked too slowly in the State Gazette workstream to gain any additional benefit under the load and stress testing scenarios. Microsoft is working on improvement to align the capability with Windows™ based workloads.

7. A server patch management strategy can help improve security

The cloud environment is dynamic and fast evolving and requires that services run in the cloud are kept up-to-date. An appropriately defined, approved and tested patch management strategy can improve security and reduce attack vectors.

Currently a patch management strategy for servers across all Estonia's e-government system environments does not exist, and should therefore be developed. When critical zero day security flaws are identified in base operating systems, e.g. Centos/SUSE Linux, in application frameworks, e.g. Java, and in application platforms, e.g. Apache, it is critical for all affected e-government systems to be rapidly updated. For each e-government system, a patch management strategy may be a complete application redeploy on the new base operating system, or the application of a binary patch to a running system. Regardless of the approach, each system should have the policy defined and tested so that it can be applied as rapidly as possible.

3.3.3.5. Storage Architecture Findings

One of the most significant tests conducted included the migration of the storage architecture from a traditional direct or network attached storage to a cloud storage subsystem, which provides built-in high availability and redundancy with no additional cost or effort.

While a full utilization of the Windows Azure Storage (WAS) subsystem would have required substantially more effort in solution reconfiguration, as part of the research project two storage mechanisms were identified that were already supported by the most current versions of CentOS and FreeBSD: Hypertext Transfer Protocol (HTTP) and Common Internet File System (CIFS) / Server Message Block (SMB) storage.²⁰ As a result, by shifting away from local physical storage and utilizing the Network File System (NFS) storage, the workstreams were able to benefit significantly from the WAS properties and increase redundancy and performance under load. However, due to limitations discovered during testing, it was not

²⁰ About SMB/CIFS: <http://msdn.microsoft.com/en-gb/library/windows/desktop/aa365233%28v=vs.85%29.aspx>

possible to use CIFS as envisioned, therefore the team used direct attached VHDs that were hosted on Microsoft Azure™ storage and exposed as an NFS share.

Key findings related to storage architecture included:

1. Standard file system disk structure for database, log, and application data improves performance

Traditional on-premises operating system deployments tend to combine application and base operating system files on the primary disk partition. In the cloud, this impacts the ability for the application instance to be rapidly replicated or have content updated independently from the base operating system image. This is because all data files have to reside on the appropriate file system disks. When required, additional file system disks can be added: up to 32 extra disks depending on the virtual machine size. This approach ensures that when higher performance is needed, it can be easily achieved by configuring the data volumes using software redundant arrays of disks and virtual machine sizes with Solid State Drive (SSD) storage. For both workstreams the original file system layout needed to be changed in order to accommodate differences in cloud storage principles.

2. Any modification of the overall host, disk, and file system layout is dependent on boundaries

The cloud infrastructure consists of standardized virtual disks, which have size and performance boundaries. When planning the migration, especially from physical servers, the boundaries have to be taken into account and planned for. For example, the portable version of the State Gazette was based on a single disk drive, which was larger than the cloud platform's limitation of the operating system's drive size (127 GB). To be able to port the solution to the cloud, the operating system and data had to be split onto two disk drives. Similarly, the President.ee website was also based on a single disk drive, which was larger than the FreeBSD instance created from the template. In this case the operating system disk was extended rather than just adding an additional data disk.

3. The primary operating system disks cannot be used for application or service storage needs

The research project confirmed that the default operating system disks are optimized for a fast virtual machine boot and are unlikely to perform strongly in other roles. The performance and capacity is not sufficient when used for application or service storage needs. To overcome this limitation additional data disks need to be added to virtual machine to be used for application or service storage needs.

4. Data files can be shared between application servers

Data that is stored on a network file share accessed by application servers can be copied directly to the cloud platform storage and accessed through SMB file service, in much the same way as it can be shared on the network file services storage on the on-premises solution. However, it has to be noted that support for this functionality is built-in only to latest versions of operating systems.

The research project found that this functionality would be beneficial across both workstreams. Workstream #1 contains a media library, where the usage of this functionality could help with to save space and make synchronization faster and more efficient. Workstream #2 on the other hand, includes a large document cache, where the usage of functionality could help save space. Moreover the need for additional Virtual Machine simulating network file system (NFS) storage would be removed.

5. Data replication between on-premises and cloud systems is straightforward

Existing e-government solutions for data replication were able to be utilized for this research project, but were not necessarily optimized for cloud systems. The existing applications replicated data on a

one-to-one server basis, which did not take advantage of the cloud platform ability to scale-out multiple read-only instances of a given set of data.

At later stages the replication mechanism was optimized to replicate only once through a low bandwidth link to the cloud and then utilize high bandwidth network connectivity to replicate to additional virtual machines.

6. Cloud storage architecture can enhance existing application storage

The existing applications were designed and used in line with a more traditional model for file storage, specifically using a combination of NFS shares and local cache. As noted above, this does not allow for the applications to fully utilize the cloud platform to provide storage replication and scaling. To be able to fully access those cloud benefits, a cloud scalability pattern should be adopted for the storage architecture, including caches, and common data files, so that the storage tier can be horizontally scaled-out for load balancing, additional capacity, and improved availability.

7. Regular updates of cloud virtual machine operating systems needed for all WAS functionality

The public cloud is a very dynamic environment, where new functionalities are added on a frequent basis. To be able to use these new functionalities, it is critical that the operating systems are kept under the support boundaries, usually across the last two versions, although it can sometimes be solely the latest version. In certain circumstances, access to the latest and most advance cloud platform features such as CIFS and cloud file based storage require the most recent upstream version of both core operating systems and drivers. In these circumstances, it is essential for applications to be able to upgrade to, and run on, the most current operating system releases. By doing so, applications are able to take advantage and utilize capabilities that have been integrated into the current operating system releases.

3.3.3.6. Network Architecture Findings

The network architecture in the cloud environment differs significantly when from traditional computing, primarily across the features that enable increased load and capacity for the target applications. As a result, applications and their network configurations need to be updated to be able to fully utilize the cloud platform's network features, such as rapid DNS updates and low Time To Live (TTL) for high level addresses, load balancing, dynamic public and private virtual IP addresses, as well as selective routing based on client access location, server load, and response time.

For example, when the State Gazette servers were switched from one location to another, i.e. from local servers to the cloud, the DNS had to be adjusted manually. It became clear that user experience differed: some users were directed to the new server in ten minutes, while others ended at the old IP address even several hours after the DNS change. This means that DNS TTL and DNS caching can play a critical role in times of emergency or rapid failover scenarios. As a result, an analysis of how to make this process quicker is warranted. Similarly, an understanding of how to build up the data handling during the switch and the fallback between different locations must be developed. It is possible that the data might differ in between locations after the switch and therefore the data must be renewed in the prime location, ensuring that the data and indexes retain their integrity.

In this particular case, the existing applications used in the research project originally had conventional network configurations in place across both workstreams. This has meant that the changes to their settings were mostly required in the configuration files. The changes needed were minimal as the selected cloud platform allows for guest operating system agents, installation and boot time scripting. Nevertheless, they

were critical to enabling the test scale and load scenarios required to survive a DDoS attack. Other key findings related to network architecture included:

1. Statically configured internal DNS naming conventions are required in the cloud environment

The IP address space in the cloud is dynamic by default. Internal DNS servers and virtual machines need to be configured correctly to allow appropriate name resolution.

2. Ownership, configuration and maintenance of the DNS Start of Authority (SOA) is vital

The SOA record defines the best source of information for domain name queries when searching for website addresses. As a result, well defined and secure operational procedures, including management and update procedures, are critical to ensuring the services in question will present authorized information to the public.

3. Statically configured IP public and private addresses need to be carefully handled

Static IP addresses can be achieved through Dynamic Host Configuration Protocol (DHCP) reservation. However, it is important to note that statically configured IP public and private addresses could prevent application solutions from working if they are migrated to cloud environments without proper consideration of the IP and DHCP requirements.

4. Virtual network, IP address and name resolution function differently in the cloud

A data center hosted in the cloud can be considered a remote network location. As a result, virtual network, IP address and name resolution function differently in the cloud in comparison to on-premises environments, which needs to be taken into consideration in the preparatory phase.

This finding is directly related to the above findings and is critical to address how IP and DNS resolution work during the critical failover and failback time periods. When it is critical to failover rapidly to the cloud, DNS and the appropriate TTL intervals need to be adjusted so that users are properly and rapidly redirected to the cloud instances of the applications, thus avoiding possible scenarios where they are directed to a compromised on-premises site or server.

5. Applications need to be able to use the load balancer provided by the cloud platform


When moving applications to the cloud, there is no need to use application specific load balancers, as they are typically provided within the cloud infrastructure. Load balancers are important, as they act as a reverse proxy and distribute network or application traffic across a number of servers, thereby increasing capacity and reliability of applications.

Depending on application requirements, the appropriate load balancer distribution method needs to be selected in the cloud platform. This will ensure that the cloud load balancing capability functions correctly.

3.3.4. Technical Performance Findings

In this section, the research project team outlines the performance results, when comparing the applications' behavior in their original on-premises environments and the target cloud environment. All of the performance, load and stress test cases were built using Microsoft Visual Studio® 2013 and Visual Studio Online and run on the Azure™ cloud platform. The results were exported and analyzed by the testing team. While the tests were conducted across the three test cases, this report focuses on load testing in particular. Microsoft Azure™ was used to initialize a large number of distributed test agents that can successfully simulate a website load under different circumstances, such as different bandwidth, browser types, click pattern, etc. The following diagram depicts the load testing approach and architecture:

Figure 6: Load testing architecture



Workstream #1 Load Test Results

The load tests were performed on both the on-premises and cloud versions of President.ee. Overall, the website performed as expected and the tests showed that the response time for the home page load was in the proposed limit (goal) of 5 seconds. The tests we conducted that are particularly worth highlighting were:

- Test case 15: Basic performance test measuring load time of homepage and search results from the website;
- Test case 16: Media performance test measuring load time of random image and random video from the website;
- Test case 17: Auto-scaling test, verifying that website scales up automatically under heavy load;
- Test case 19: Mixed performance test (also known as a mixed feature test), simulating a typical scenario in which the website is under heavy load in different areas: e.g. 70% of the users on the home page, 16% on the search page and 12% on image content and 2% on video content.

The table below represents average response times for different load test scenarios. Response times are listed for both cloud and on-premises version of President.ee.

Load test	Number of users	Duration of test	Azure Avg. Response Time	On-Premises Avg. Response Time	Comments
HomePageLoadTest	1,000 users	30 minutes	3.14 seconds	3.38 seconds	Auto-scaling enabled, starting with 1 virtual machine instance
SearchResultsLoadTest	1,000 users	30 minutes	2.90 seconds	2.93 seconds	
MediaImageLoadTest	1,000 users	30 minutes	0.20 seconds	not tested	
TestMixLoadTest	500 users	30 minutes	3.19 seconds	4.08 seconds	Mixed tests will all features tested

Table 7: Overview of key tests conducted in Workstream #1

Azure Workstream #1 Load Test Results: The following Figure shows typical load test results, when run from Visual Studio® and deployed to Visual Studio Online. This is a sample load test for the image content (MediaImageLoadTest, test case 15) for 500 concurrent users during a period of 30 minutes. The test gave us the average page response time of around 0.2 seconds.




Figure 7: Azure™ version of MediaImageLoad test results

It is important to note that the first few spikes in Figure 7 were managed by the auto-scaling feature. This feature was demonstrated during the tests under heavy load by starting an additional virtual machine instance to offload the traffic. This can be seen in Figure 8 on the following page.

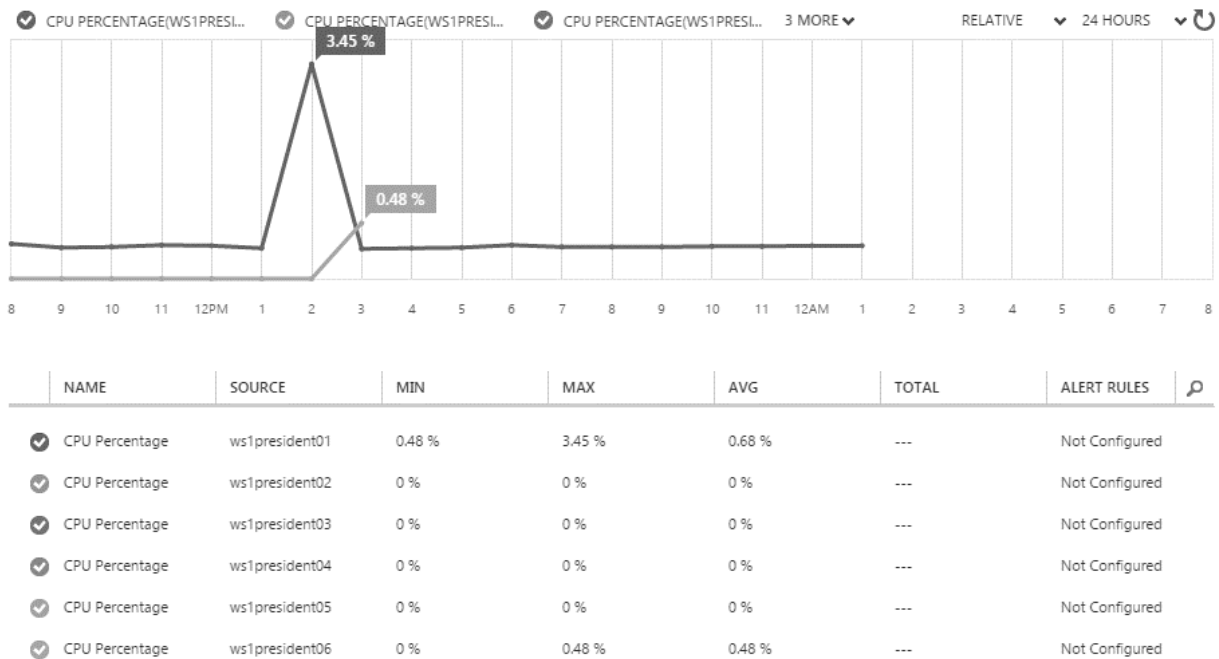



Figure 8: Azure™ Monitoring Dashboard demonstrates auto-scaling features

Furthermore, the following Figure depicts a longer period of activity, starting from six virtual machine instances, i.e. scenario without auto-scaling, then shrinking down to just one virtual machine instance, i.e.



with auto-scaling enabled, and finally with an second virtual machine instance, which started automatically using the auto-scaling feature. After the load tests stopped, the additional instance was also automatically shut down.

Figure 9: Auto-scaling demonstrated

On-premises Workstream #1 Load Test Results: The on-premises version was running on the following hardware: CPU: Intel Xeon X3350 and RAM: 8GB. The following diagram shows typical load test results when run from Visual Studio®(and deployed to Visual Studio Online). This is a sample load test for the image content (HomePageLoadTest, test case 15) for 500 concurrent users during 30 minutes. The first couple of spikes for the average page response time could be explained by “cold boot”; when the server/database is warmed-up, the page response time gets normalized.




Figure 10: Load test results on premises, Workstream #1

Workstream #2 Load Test Results – Portable version

In workstream #2, tests were performed against a portable version of the website running in Azure™ first. The focus was on the read-only mode with limited write/modify functionality. As with workstream #1, the website for home page load performed as expected and the tests showed that the response time was in the proposed limit (goal) of 5 seconds. The tests conducted that are particularly worth highlighting were:


- Test case 30: Basic performance test, measuring load time of homepage and search results from the website;
- Test case 31: Auto-scaling test, verifying that the website scales up automatically under heavy load;
- Test case 33: Mixed performance test (also known as a mixed feature test), simulating a typical scenario in which the website is under heavy load in different areas, e.g. 50% of the users on the home page, 30% on the search page and 20% on law/act details page.

Load test	Number of users	Duration of test	Avg. Response Time	Comments
HomePageLoadTest	1,000 users	30 minutes	5.50 seconds	Auto-scaling enabled, starting with 1 VM instance
SearchResultsLoadTest	1,000 users	30 minutes	30.83 seconds	Performance level to be evaluated. (possible cause – browser certificate)
LawDetailsLoadTest	1,000 users	30 minutes	6.07 seconds	
TestMixLoadTest	500 users	30 minutes	20.60 seconds	Mixed tests will all features tested

Table 8: Overview of key tests conducted in workstream #2, portable version

The following figure shows a sample of load test results, when run from Visual Studio® and deployed to Visual Studio Online. This is the load test for the search results page, which was part of the basic performance test for workstream #2, test case 30, for 500 concurrent users during 30 minutes. The spikes are caused by “heavy” search operations that perform full text search in the database according to the search string.

Figure 11: Sample load test results in Visual Studio®, Workstream #2



During the performance and load testing, it was found that the auto-scaling feature was not behaving as expected for a non-Windows operating systems running on virtual machines. The performance data lagged by around 20 minutes from real time on the Azure™ portal, which was one of the reasons for the slower than expected response to scale-up operations. After some additional tests it was realized that auto-scaling behavior is performing as expected on the Windows operating systems running on virtual machines.

Workstream #2 Load Test Results – Full version

The second set of tests was performed against a full version of the website, which had read, write, and modify functionalities enabled. The website ran both in the cloud and on-premises. As with the first workstream, the website performed as expected and the tests showed that the response time was in the proposed limit (goal) of 5 seconds. The tests conducted that are particularly worth highlighting were:

- Test case 30: Basic performance test, measuring load time of homepage and search results from the website;
- Test case 31: Auto-scaling test, verifying that website scales up automatically under heavy load;
- Test case 33: Mixed performance test (also known as a mixed feature test), simulating a typical scenario in which the website is under heavy load in different areas, e.g. 50% of the users on the home page, 30% on the search page and 20% on law/act details page.

Load test	Number of users	Duration of test	Azure Avg. Response Time	On-Premises Avg. Response Time
HomePageLoadTest	1,000 users	30 minutes	3.21 seconds	3.01 seconds
SearchResultsLoadTest	1,000 users	30 minutes	156.74 seconds	90.64 seconds
LawDetailsLoadTest	1,000 users	30 minutes	1.25 seconds	3.91 seconds
TestMixLoadTest	500 users	30 minutes	72.49 seconds	25.28 seconds

Table 9: Overview of key tests conducted in Workstream #2, full version

Azure Workstream #2 full version load test results: The following diagram shows a sample of load test results when run from Visual Studio (and deployed to Visual Studio Online). This is the load test for the

search results page (part of the basic performance test for WS #2, test case 30) for 500 concurrent users during 30 minutes.




Figure 12: Sample load test results, full version Azure™ workstream #2

It is worth noting that during the performance/load testing, the auto-scaling feature for the full version of workstream #2 government services on the cloud platform was not enabled. The following virtual machine sizes formed the basis of the solution:

- Application server: 1 A5 virtual machine instance
- Database server: 1 A5 virtual machine instance

On-premises workstream #2 full version load test results: The following diagram shows a sample of load test results when run from Visual Studio (and deployed to Visual Studio Online). This is the load test for the search results page (part of the basic performance test for WS #2, test case 30) for 500 concurrent users during 30 minutes. The following hardware was used for the on-premises solution: the application server: 4vCPU, 12 GB RAM and the database server: 4vCPU, 8 GB RAM.













Figure 13: Sample load test results, on-premises, workstream #2

4. Conclusions and Recommendations

The research project confirmed the majority of the initial hypotheses. It became clear early on that both the Estonian President’s website and the State Gazette website were able to successfully migrate to and operate in the public cloud for the duration of the project (Hypothesis 1). Moreover, while certain issues were encountered, it also became clear that cloud computing can be leveraged to enhance the performance and resilience of the government services, given cloud capabilities, such as DDoS protection and auto-scaling (Hypothesis 5). Indeed, virtual machines in the cloud environment can be hosted and stored in numerous locations, which may be situated in different countries, or even continents, consistent with European Union Safe Harbour rules. The project identified that using this diversification approach provides additional security for the data and services used.

#	Hypotheses Tested	Result
Hyp 1	Services migrated to the public cloud and are able to run successfully.	
Hyp 2	Minimal architectural changes required to migrate the services.	
Hyp 3	A number of areas for improvement identified given that existing services were originally developed to run on-premises.	
Hyp 4	Minimal time and effort required to modernize the government services after migration.	
Hyp 5	Cloud platform built-in capabilities, such as auto-scaling, DDoS protection, etc., leveraged to enhance the performance, stability, security of, and public trust in the services.	
Hyp 6	Design of the operational procedures, such as failover and fail-back processes, key to the services running in the cloud.	
Hyp 7	Estonia’s ISKE security standard should be updated to address public cloud.	
Hyp 8	Legal basis exists for asserting international law protections against compelled disclosure of data stored in Virtual Data Embassy.	
Hyp 9	Under Estonian domestic law, non-restricted data can be migrated to the cloud.	

It also emerged that while most applications, originally designed for on-premises use, can be moved to the cloud “as is”, this might result in difficulties with scaling and achieving full functionality. At present, although most applications can be ported over to a cloud platform, without modernization to address key limitations such as denial of service or load balancing, such migration provides limited benefits. For e-government applications to truly benefit from a migration to a cloud platform, they should be thoroughly evaluated, e.g. undergo a risk assessment, to ensure that the applications meet the current threat landscape and have a defined switching procedure and procedures for running services in cloud in place.

During the research project, the underlying cloud platform was also able to demonstrate the ability to handle DDoS attacks (layer 1-3) regardless of the application used. However, the DDoS scenarios underlined the need for application modernization on both workstreams to be able to properly handle

application level DDoS attacks (layer 4-7). Due to the limits of auto-scaling for Linux workloads under DDoS attacks, it might be prudent to consider utilizing PaaS to mitigate any attacks. A State Gazette requirement for cloud services was that auto-scaling would work within an expected timescale, however this was not met as auto-scaling features had longer latency on non-Windows operating system virtual machines. In addition, the research project also confirmed that the applications work well in the cloud and users did not experience any difference with respect of servers operating locally, in the cloud, or on servers physically located outside of Estonia. Overall, the project demonstrated that from a technology perspective, the *Virtual Data Embassy Solution* is feasible, as only limited architectural changes would be required (Hypothesis 2).

Despite the relative ease of the technical work, several areas for improvement have been identified (Hypothesis 3), in particular across the management and operations processes. It was found the research project's multi-disciplinary approach, combining legal, technical and risk management aspects, was particularly beneficial. It clearly revealed that projects such as this one need substantial preparation and testing across these areas and that textbook readiness is near impossible to achieve. This confirmed that the testing of failover to public cloud should not be done in an actual crisis (Hypothesis 6). By testing and migrating data during non-critical situations, i.e. not "fragile" or "no control" scenarios, the government is able to verify the operational procedures are working to ensure a seamless transition from the on-premises master data systems to the cloud based systems. However, the learning processes that achieving this understanding required also meant that the time and effort needed for the migration and testing was substantially heavier than expected (Hypothesis 4).

A critical challenge identified was that many existing information systems are poorly documented. Detailed documentation on information system architecture and functionality specifications, interfacing, or customizations is often missing and frequently only a small number of experts understand the workings of the system. This could lead to gaps in digital continuity, in particular when quick reactions are required. Moreover, as this is a new concept, website and system migration to cloud suffers from lack of government personnel experience and expertise. While the application teams were eager to learn, the gaps mean that the government would be almost fully reliant on the selected vendor for the migration.

A very concrete lesson that emerged was the need for the government to be able to request changes regarding the operation of DNS zones and records, which is critical to usage of the DNS by end users. The DNS system and certificate authority system will play a very significant role in establishing citizens' trust and their ability to get to the appropriate services during a crisis or cyber-attack. Presently, this is currently not the case, and the DNS and critical name resolution systems rely heavily on manual updates by a collection of civil servants and private companies.

It is also worth noting that the research project sought to identify any technical or operational requirements that the government could oblige the selected cloud provider to comply with. Some of these requirements were clear from the start, such as the need for data centers to be located in secure, stable countries that abide by international law and are located in the European Economic Area or other safe harbors. Others only emerged in the process of implementation and could for example include the ability to: support multiple diverse operating system and application runtime environments; provide scale up and scale out computing and storage resources; deliver geo-redundant storage for backup and mirroring of critical data; implement multiple layers of network security for prevention of DDoS attacks; provide secure network connections to on-premises computing facilities; deliver a common virtual machine environment for cloud and on-premises operations; provide high scale network load balancing to handle traffic surges; and, move and restart applications across multiple data centers. The research project made clear that these requirements are not out of the ordinary and the major public cloud solutions available today are likely to meet the majority of them.

When it comes to the legislative environment, the research project found there are no clear legal restrictions under existing Estonian domestic law to the migration of government services to a Virtual Data Embassy in a public cloud, with limited exceptions, confirming Hypothesis 9. These include data relating to “critical” services, as defined by the Estonian Emergency Act. To enable migration of critical government services, changes to laws must be addressed, e.g. to clarify digital continuity. The amendments would need to consider data protection requirements, as well as developing a detailed understanding of Emergency Act processes. Moreover, cloud security requirements, as well as a clear division of responsibilities between the cloud provider and the government should be considered. It is also pivotal that data migration to a public cloud platform be highly trustworthy with operational and security measures in place to prevent unauthorized tampering. This requires third party auditing reassurances, as well as secure digital encryption key management to ensure that data has not been altered. To this end, information security standards, ISKE in this case, should be updated, building upon international standards, as appropriate, to address cloud computing.

When it comes to international laws, the research project found there are reasonable grounds for arguments against compelled disclosure, although there is a lack of precedent for this exact set of circumstances (Hypothesis 8). While existing international law can be applied to new technologies there is a need for further discussion at the international level concerning modern storage methods of government data, based on principles of protecting public safety, individual privacy, transparency and due process.

The results of the research project have led to eight high level recommendations, listed below, to be taken into account when considering migrating and operating government services in a public cloud. These recommendations fundamentally point to the fact that, as technology inexorably advances, governments must be flexible and adaptable in order to anticipate and benefit from the latest technological advances, as well as being able to create new law or find new ways to apply existing laws to ensure digital continuity.

Recommendations

1. Information security standards, ISKE in this case, should be updated, building upon international standards, as appropriate, to address cloud computing.
2. For the migration of restricted data, countries should consider developing digital continuity legislation and a strategy to increase assurances of diplomatic and other international law protections.
3. Governments should keep data governance and security models up to date across the data lifecycle and required of government entities.
4. Cloud computing should be utilized to increase security and resilience of government infrastructure.
5. An overarching cloud strategy and government action plan facilitating cloud migration should be developed to enable technical and operational agility and increase cost-effectiveness.
6. Comprehensive risk assessments should be conducted to establish acceptable risk tolerance and associated requirements.
7. The trusted relationship between the government, private sector company, host country, and the country where the provider is headquartered should be deepened for success of the Initiative.
8. Operational procedures should be prepared and tested in advanced rather than in a crisis.

5. Research Project Team

Taavi Kotka, Ministry of Economic Affairs and Communications for Estonia, Project Lead

Bruce Johnson, Microsoft, Project Lead

Tyson Storch, Microsoft, Project Lead

Chris Brown, Microsoft

Kaja Ciglic, Microsoft

Amanda Craig, Microsoft

Tomaz Cebul, Microsoft

Dejan Cvetkovic, Microsoft

Kaspar Hanni, Microsoft

Jüri Heinla, Ministry of Justice of Estonia

Laura Kask, Ministry of Economic Affairs and Communications for Estonia

Andres Kütt, Information System Authority of Estonia

Rain Laane, Microsoft

Mikk Lellsaar, Ministry of Economic Affairs and Communications for Estonia

Simon Liepold, Microsoft

Luka Lovosevic, Microsoft

Theo Moore, APCO Worldwide

Rebecca Radloff, Microsoft

Jüri Raidla, Raidla Lejins & Norcous

Raimonds Rudmanis, Microsoft

Karoliina Raudsepp, Ministry of Economic Affairs and Communications for Estonia

Mehis Sihvart, Centre of Registers and Information Systems, Estonia

Additional Technical Team

Martin Parik, Ministry of Justice of Estonia

Raivo Oravas, Centre of Registers and Information Systems, Estonia

Mait Makkar, Centre of Registers and Information Systems, Estonia

Rome Mitt, Centre of Registers and Information Systems, Estonia

Martti Allingu, Centre of Registers and Information Systems, Estonia

Aldo Vaikre, Centre of Registers and Information Systems, Estonia

Evar Sömer, Centre of Registers and Information Systems, Estonia

Ivo Vellend, Office of the President of the Republic of the Estonia

Taavi Meos, Office of the President of the Republic of the Estonia

Ardo Birk, Office of the President of the Republic of the Estonia

Tarmo Hanga, Information System Authority of Estonia

Kirkka Kivilo, IT and Development Centre of Ministry of the Interior, Estonia

Alvar Haug, IT and Development Centre of Ministry of the Interior, Estonia

Ants Nõmper, Raidla Lejins & Norcous

Andres Ojaver, Estonian Data Protection Inspectorate

Urmo Parm, Estonian Data Protection Inspectorate