Project Acronym: STORM CLOUDS

Grant Agreement number: 621089

Project Title: STORM CLOUDS – Surfing Towards the Opportunity of Real Migration to CLOUD-based public Services

Deliverable D3.4.2

Best practices for cloud-based public services deployment

Work Package: WP3

Version: 1.0

Date: 29/07/2016

Status:

Nature: Report

Dissemination Level: PUBLIC

Editor: Panagiotis Tsarchopoulos (Aristoteleio Panepistimio Thessalonikis - AUTH)

Authors: Panagiotis Tsarchopoulos (Aristoteleio Panepistimio Thessalonikis - AUTH)

Contributors: Komninos N., Kakderi C., Papargyriou C.

Reviewed by: Agustín González Quel, Alkiviadis Giannakoulias

Version Control

Modified by	Date	Version	Comments
Panagiotis Tsarchopoulos	25/7/2016	0.5	Ready for review
Agustín González Quel	26/7/2016	0.6	Comments and suggestions
Alkiviadis Giannakoulias	27/7/2016	0.7	Comments and modifications
Panagiotis Tsarchopoulos	29/7/2016	1.0	Final version

Executive Summary

Surfing Towards the Opportunity of Real Migration to Cloud-based public Services (STORM CLOUDS) is a project partially funded by the European Commission within the 7th Framework Program in the context of the CIP project (Grant Agreement No. 621089).

The project aims to define useful guidelines on how to address the process of moving towards a cloud-based solution for Public Authorities and policy makers. These guidelines will be prepared based on direct experimentation in at least 4 European cities, creating a set of relevant use cases and best practices.

Work Package 3 (WP3) of the STORM CLOUDS project aims to adapt the selected services and to integrate them to the STORM CLOUDS Platform (SCP) infrastructure created in WP2, while at the same time create the tools and procedures to facilitate services migration to the cloud. Once applications are successfully tested in the private cloud infrastructure, they will be deployed in the pilot cities public cloud infrastructure.

The D3.4.1 and D3.4.2 deliverables include several best practices related to the transferral of services into the cloud (service cloudification)". These best practices reflect the work done in WP1, WP2 WP3 and WP4. They facilitate Public Authorities to use the migration tools (developed in WP2) along with methodological approaches and software techniques (created in WP1, WP3 and WP4), in order to easily transfer their public services to the Cloud.

The aim of this document is to provide a practical reference to help Public Authorities (both decision makers and IT staff) to consider and execute application migration to Cloud Computing. It focuses on the migration of applications to the STORM CLOUDS platform. The document includes a list of best practices, closely related with the project's IT tools and methodologies, that take into consideration both business and technical requirements.

The D3.4.1 was the first iteration of the deliverable. It was incorporated the experience of the 1st stage of the cloudification process. In this stage the pilot cities have successfully cloudified a number of public services. The second iteration D3.4.2 covers the 2nd stage. In that stage, the pilot cities will make use of public services that have already cloudified in other pilot cities. Moreover, other European Cities will deploy the STORM Cloud services.

This document is an accumulated document as the contents of D3.4.1 has been evaluated based on experiences gained during the 2nd stage of the cloudification process. Moreover, the existing best practices have enriched with new ones about interoperability, privacy, Cloud provider selection, application re-architecting for the Cloud and containerization.

The section titled "Background" provides an overview of the Cloud Computing and its benefits for the Public Authorities. In addition, it presents the technical solutions, which the STORM CLOUDS project has created.

The section titled "STORM CLOUDS Best Practices" is the heart of the document and addresses the essential issues of the application's migration process.

The best practices evolve during the project lifetime and are available to the public at: http://stormclouds.urenio.org/resources/best-practices/

Table of Contents

List of Figures

List of Tables

Abbreviations

Acronym	Description
API	Application Programming Interface
CSLA	Cloud Service Level Agreements
CSP	Cloud Service Provider
HPE	Hewlett Packard Enterprise
IaaS	Infrastructure as a Service
ICT	Information and Communication Technologies
IMM	Interoperability Maturity Model
ISO	International Organization for Standardization
ITU	International Telecommunication Union
PaaS	Platform as a Service
SaaS	Software as a Service
SCP	STORM CLOUDS Platform
SMS	Short Message Service
SLA	Service Level Agreement
SSH	Secure Shell
SSL	Secure Sockets Layer
VM	Virtual Machine

1. Introduction

Over the past years, the term 'smart cities' has evolved to denote the cognitive processes combined with the deployment of ICTs, institutional settings for innovation and physical infrastructure, which taken altogether increase the problem solving capability of a city or a community [1]. The main features of a smart city include applications that connect, manage and optimize data from a complex set of devices, sensors, people and software, creating real-time, context-specific information intelligence and analytics, which aim to transform the urban environment and address its specific needs [2]. Managing such enormous amount of heterogeneous data requires, among others, high storage capacity and performance computing power [3]. For this, the latest developments in Cloud Computing and the Internet of Things are widely deployed in smart cities [4].

Although the private sector has already taken advantage of Cloud Computing technologies, Public Authorities are lagging behind in the utilisation of them, as they have not realised yet the tremendous potential that Cloud Computing holds to deliver public value by increasing operational efficiency and responding faster to constituent needs.

This document is a practical reference for Public Authorities (both decision makers and IT staff) in order to consider and execute application migration to Cloud Computing. It includes a list of best practices, which were derived from the experience gained by the migration of several Smart City Services into the STORM CLOUDS platform (SCP). The presented best practices take into consideration both business and technical aspects and requirements.

The section titled "Background" provides an overview of the Cloud Computing and its benefits for the Public Authorities. In addition, it presents the technical solutions that the STORM CLOUDS project has created.

The section titled "STORM CLOUDS Best Practices" is the heart of the document and addresses the essential issues of the application's migration process. The solutions presented in each issue came from the work done in WP1, WP2, WP3 and WP4.

The best practices evolve during the project lifetime and are available to the public at:
http://stormclouds.urenio.org/resources/best-practices/

2. Background

It is important for Public Authorities to understand the benefits of Cloud Computing. Moreover, it is good to have a clear view of the migration tools that the STORM CLOUDS platform offers. In this section, the most significant benefits but also the main concerns of Cloud Computing are presented. Also, the different components of the SCP are illustrated.

2.1 Benefits of Cloud Computing

Cloud Computing has emerged during the last years as a disruptive model, with the ability to transform the IT organizations to be more responsive and agile than ever before. This model represents a fundamental change in the way that information technology, hardware and software are invented, developed, deployed, scaled, updated, maintained and paid [5]. Cloud Computing serves as an enormous step towards delivering computing as a utility (like traditional utilities such as water, electricity and telephony) [6] by changing the traditional access model, where data and applications is fully contained in the same physical location (the users' computers), to a new one, where the users access their data and applications outside their own computing environment through the Internet.

The US Government's National Institute of Standards and Technology (NIST) defines Cloud Computing as "*a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction*" [7], while the IT research and advisory company Gartner, uses a simplified definition and defines cloud computing "*as a style of computing in which scalable and elastic IT-enabled capabilities are delivered as a service using Internet technologies*" [8]. A common analogy to understand Cloud Computing is renting versus buying. Actually, someone rent IT capacity (computing power, disc space, applications, etc.) from a cloud service provider and use them over the Internet, instead of buying his own IT requirements. Moreover, he pays only for the resources he uses.

Cloud computing promises economic benefits, speed, agility, flexibility, rapid elasticity and more innovation. The motivations of the organisations for the migration of their applications to the Cloud are closely related to the following four key Cloud Computing characteristics:

- On-demand self-service. A consumer can unilaterally provision IT resources (e.g., storage, processing power, memory, bandwidth, etc.), as needed automatically without requiring human interaction with the cloud provider.
- Broad network access. IT resources are available over the Internet and accessed through heterogeneous devices (e.g., mobile phones, tablets, laptops, and workstations). There is a sense of location-independence because the customer generally is not aware of the exact location of the provided resources.

- Rapid elasticity. A Cloud environment offers to the consumer the ability to rapidly scale up or down the IT infrastructure commensurate with demand. To the consumer, the capabilities available for provisioning usually appear to be limitless and can be reserved in any quantity at any time.
- Measured service. Cloud systems monitor, control and report the use of IT resources by leveraging a metering capability at some level of abstraction suitable for the type of service. This lead to a transparent relationship between the consumer and provider of the Cloud service.

The above-mentioned characteristics, create a highly efficient, scalable and elastic computing environment, which is available through a business model where consumers buy only the capacity and capabilities needed at any given time, instead of buying and deploying all the components of a computing centre.

Using this "pay for what you use" approach, an organization can significantly reduce the up-front costs by avoiding the procurement of hardware and software in advance, as well as the resultant infrastructure depreciation. Moreover, cloud adoption can decrease the operational costs, as the organization will maintain at a given time only the required resources, which could be scaled out/up and in through rapid elasticity. Also, other IT costs can be reduced by buying general purpose capabilities such as asset management, security, collaboration, etc. as a service, instead of maintaining a specialized in-house IT department. As organisations operate in an unstable economic environment, smart consumption-based procurement allows them to scale up to fulfil new demands and reduce spending, if necessary, to address changes in budgets and funding.

Besides IT cost reduction, the adoption of Cloud Computing increases significantly the effectiveness of an organisation in fulfilling its mission. The following benefits are clearly related to the ability of an organization to deliver more and better results [[9], [10]]: (a) Just-in-time infrastructure, (b) more efficient resource utilisation, (c) ability to respond to emerging needs, (d) increased reliability and performance through the adoption of new technologies, best practices and security enhancements, (e) better information sharing and collaboration through the heterogeneous, location-dependent access, and (f) evaluation and optimisation of business processes through enhanced real-time visibility and audit of applications and infrastructure.

Cloud Computing can also serve as an enabler for innovation in the organisations that adopt it in a lot of ways: (a) reduces time to market by allowing the scale-out/up of resources in a cost-effective manner, (b) gives constant access to commercial best practices and new capabilities, which can be incorporated to already existing services, instead of develop them in-house, (c) enables the launch of new initiatives based on new applications that are available from the cloud provider, and (d) connects to new and emerging technologies. Moreover, in government organisations, Cloud Computing can: (a) encourage entrepreneurial culture by reducing risk of lunching new initiatives and by allowing the low-cost experimentation in new applications and services, and (b) give them access to innovations developed in the private sector [[11]].

Besides the clear benefits, there are some concerns when Public Authorities thinking about the adoption of Cloud Computing [10]. In particular:

- Security. Moving the public data and code to an external provider creates some security risks. Although the technology to make cloud computing safe is available, securing cloud workloads usually requires new approaches and skills that may take time to internal IT staff to learn.
- Loss of control. For software-as-a-service (SaaS) and some platform-as-a-service (PaaS) and infrastructure-as-a-service (IaaS) solutions, the total control of hardware, software, security policies, etc., is placed in the hands of the external Cloud provider.
- Integration. Sometimes Public Authorities will need to integrate their internal systems with the external Cloud systems. This integration can be challenging.
- Availability and reliability of cloud applications. The combination of server performance, configuration errors, network design, and application architecture, may cause performance issues that are difficult to resolve.
- Cloud service provider lock-in. The concern is that the adoption of a Cloud Solution from a provider might make difficult the switching to a different provider.

Governments around the world have started to elaborate strategies in order to harness the benefits of cloud computing. In 2011, Vivek Kundra the U.S. Chief Information Officer, presented the Cloud First policy for Federal's Government bodies [10]. This policy is intended to accelerate the pace at which the government will realize the value of cloud computing by requiring agencies to evaluate safe, secure cloud computing options before making any new investments. In a similar way, Canada and UK Governments have adopted a Cloud First approach to public services provisioning [12,13]. In September 2012, the European Commission adopted a strategy for "Unleashing the Potential of Cloud Computing in Europe". The strategy is designed to speed up and increase the use of cloud computing across all economic sectors [14].

Despite the existence of the "Cloud First" policies, which mandates Government Agencies to take full advantage of cloud computing benefits to maximize capacity utilization, improve IT flexibility and responsiveness, and minimize cost, Public Authorities worldwide are not yet comfortable using Cloud Solutions. The STORM CLOUDS solutions aim to contribute, both methodologically and technically, to this problem.

2.2 The STORM CLOUDS Solution

The STORM CLOUDS project [15] aims to accelerate the pace at which Public Authorities move to the Cloud Computing. The project provides a methodology for the Cloud migration process, mainly from the point of view of the end-users, as well as the essential IT tools that will support this process. By taking the STORM CLOUDS approach, Public Authorities can take full advantage of the cloud computing model and provide citizens with highly reliable, innovative services quickly, despite resource constraints.

The main assets of the project are the following:

- STORM CLOUDS Platform (SCP). The platform provides the Cloud environment, which can host Smart City Applications.
- Migration tools. The tools ensure a high level of automation and data protection (in the form of automatic backups), including monitoring mechanism.
- Portfolio of Smart City Services. The portfolio includes many Cloud-based, open source applications, which are ready for use from Public Authorities.
- Roadmap for the migration of public services into the Cloud. The roadmap consists of guidelines that help Public Authorities to address the technical and business challenges in the adoption of cloud computing.
- Best practices for cloud-based public services deployment. The best practices include software techniques and methodological approaches, which facilitate the adoption of Cloud services in the Public Sector.
- Business models for the scalability and sustainability of the project's results. The business models cover the exploitation of the SCP and Smart City services, as well as the viability of the already cloudified services in the project's pilots.

The following figure presents the logical architecture of the SCP [16].

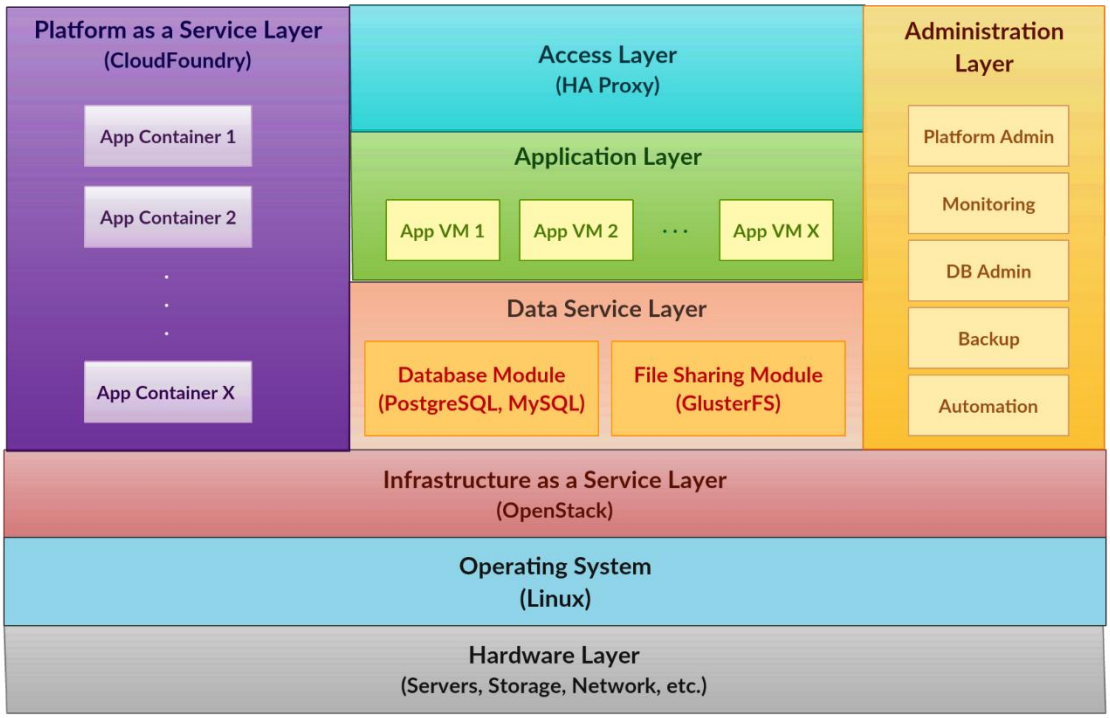

Figure 1 – The logical architecture of the STORM CLOUDS Platform

The following table presents the components of each layer.

Component	Description
Hardware Layer	Physical resources (i.e. servers, storage, network, etc.) that host the platform.
Operating System	A Linux distribution that supports OpenStack (Debian 7.0, openSUSE, SUSE Linux Enterprise Server, Red Hat Enterprise Linux, CentOS, Fedora and Ubuntu).
IaaS Layer	It is implemented using OpenStack [17]. OpenStack can be seen as a cloud operating system that controls large pools of compute, storage, and networking resources (the hardware). It gives administrators several visual tools to manage the infrastructure and deploy their applications, while empowering their users to provision resources through a web interface.
Data Service Layer	Contains the database and the file servers, both deployed on VM clusters and provide high-availability and scalability. The currently supported databases are the MySQL/MariaDB and PostgreSQL. The Data Service Layer provides each application with a private database and a private volume but multiple VMs of an application can share the same data allowing the implementation of H/A and scalability.
Application Layer	Contains the Virtual Machines that host the Smart City Applications.
Access Layer	It is the front-end for the Smart City services. It receives the http(s) requests from users and redirects them to the suitable application's VM. It also implements Load Balancing, which improves the overall performance of applications by distributing the workload to different

	VMs.
PaaS Layer	It is a sophisticated solution that allows developers to deploy their web applications to the cloud without having to take care of the underlying infrastructure. It is implemented using Cloud Foundry [18], a very popular open source system for managing the deployment of apps, services, and background tasks.
Administration Layer	Provides to the platform's administrators and to the application owners, the tools for managing and monitoring the components of the platform, as well as the Smart City applications. The tools include: • The Platform Administrator's Console, which allows the SCP administrator to have full control of the layers of the platform. Through the console, (s)he can manage the databases, the filesystem, the IaaS layer, and the PaaS layers. • The Monitoring Module, which monitors the resources (CPU load, disk space occupation, network traffic, number of processes, etc.) used by the platform's services or by the applications. • The Database Administration Module, which administers the supported databases. • The Backup Module, which takes backups from databases and file systems. • The Automation console allows application owners to deploy Smart City applications to the IaaS layer automatically using a number of predefined scripts.

Table 1- The components of the STORM CLOUDS platform

During the 2nd stage of the cloudification process, a second alternative architecture was developed in order to simplify the management of the platform for those cities that haven't cloud ready applications. The following figure presents the alternative, simplified, logical architecture of the SCP.

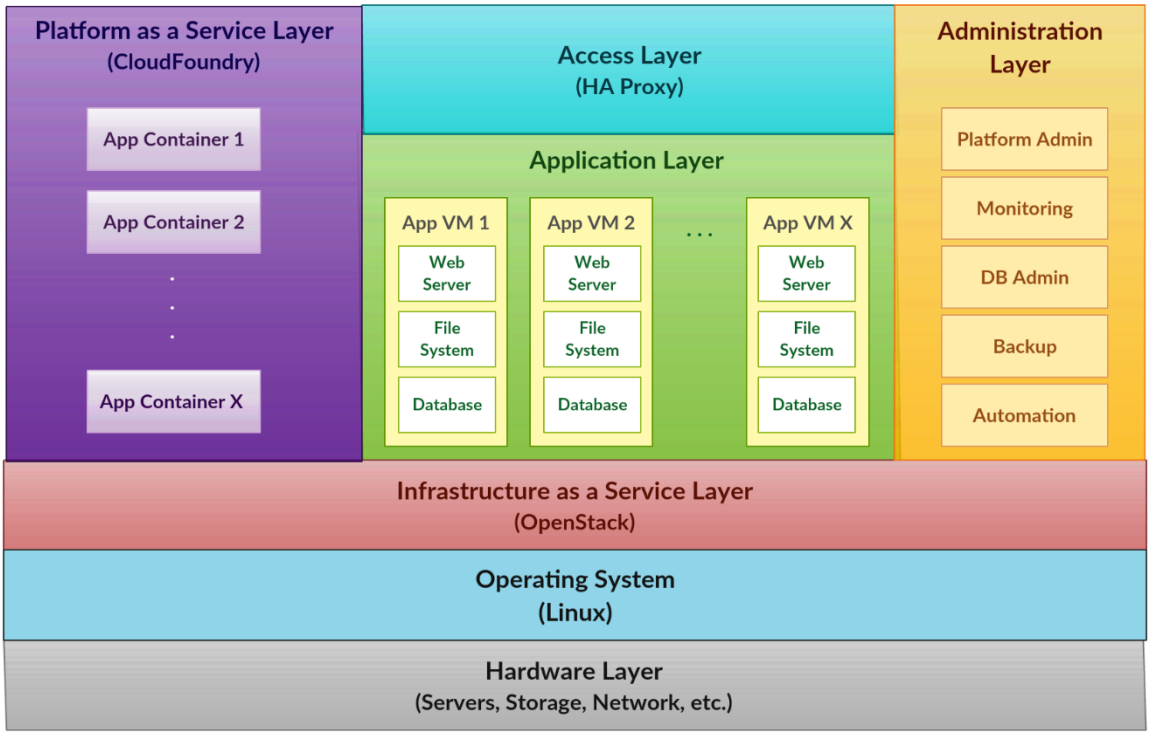

Figure 2 – An alternative, simplified architecture of the of STORM CLOUDS Platform

In this simplified architecture the components of each application (i.e. webserver, filesystem and database) are entirely contained into a single VM.

Two instances of the SCP have been implemented: one hosted on Hewlett Packard Enterprise's IT infrastructure (SCP@HPE), the other hosted on a public cloud-computing operator (SCP@Operator). SCP@HPE is used mainly for testing purposes and for supporting the technical activities for porting the selected applications to cloud (e.g. adaptation, configuration, automation, etc.). SCP@Operator, on the other hand, is used as the "production environment", for making the migrated applications available to the end users through the Internet. SCP@HPE is a private cloud providing services for exclusive use of the STORM CLOUDS project partners while SCP@Operator is hosted on a public cloud.

The STORM CLOUDS Platform has been implemented using open source technologies and solutions. This approach not only lowers the cost of the solution but also helps Public Authorities to avoid vendor lock-in, as it can be easily transferred to different Cloud providers. Another valuable feature of the platform is that it supports all Cloud deployment models (i.e. private, community, public and hybrid), as well as two of the Cloud service categories (Platform as a Service and Infrastructure as a Service).

3. STORM CLOUDS Best Practices

The selection of the best practices was based on the literature review, which revealed ten generally accepted guidelines that ensure the successful migration of applications to the cloud. These guidelines have been enriched with hands-on experience gained from a practical exercise, the deployment of Smart City applications to the SCP. The STORM CLOUDS approach, found in each best practice, offers practical knowledge that has been validated in four European Cities. Moreover, it presents the particular features and tools of the SCP that could support the implementation of that best practice. In that essence, the SCP and tools can be seen both as a valuable Cloud experimentation environment, as well as, a very useful asset for Public Authorities that aim to deploy their Smart City applications to the cloud.

3.1 Adopt an Open Innovation Methodology

STORM CLOUDS project adopted a User Driven Open Innovation methodology [19] to select the applications that would migrate on the Cloud. During this procedure, the involvement of several stakeholders was essential, not only as a methodological requirement but also because the participation of the parties concerned' would produce helpful feedback for the overall process to be closer to citizens and public employees. Moreover, another reason why this mechanism was adopted by STORM CLOUDS, compared to the selection of applications only by the Public Authority's technical staff, is the fact that this involvement would strengthen stakeholders' awareness regarding the efforts made by local authorities towards a modernisation of the public sector.

The open innovation methodology focuses on the idea of gathering external and internal knowledge to accelerate the process of innovation.

The STORM CLOUDS approach

The open innovation methodology can be achieved through the following three dimensions:

a) a user-driven innovation approach.

This approach implies that the source of the innovation relies on an intense understanding of the customer needs, which requires continuous interaction between the user and the developer of a new idea. It promotes direct involvement of the end user in the innovation process, reducing, thus, the chance of failure. At the end, user-driven innovation translates customer knowledge into unique products and experiences.

In this sense, involving citizens in the definition of the public services that the city will provide is a wise decision which is providing excellent results. Citizens that are engaged in the process of technological advancements of their city, normally perceive that their city is interested in covering their needs in the best possible way. Also, there are additional benefits when they are aware of the cost and other implication that are involved with their demands.

To achieve a user driven innovation approach, one has to identify 'lead users', defined as users ahead of the majority with respect to an important market trend or which are expected to benefit more from a solution. Lead users can be selected from different 'user groups' such as citizens, public servants, companies etc.

b) the treatment of innovation as an open system, allowing external actors to become key players in all parts of the innovation process.

A key aspect for the open innovation approach is the activation of stakeholders and users and their continuous engagement throughout the cloudification process. The activation can take place in five different phases: i) development of a communication strategy, ii) information disclosure about the services to be deployed and their benefits, iii) consultation, monitoring the stakeholders' response, iv) participation in the services' deployment, improvement and exploitation, and v) negotiation and partnerships aiming at future improvements and the sustainability of the deployed services.

Potential stakeholders can be derived from citizens, public authorities, local enterprises and organisations. These should be contacted with the use of different communication channels such as newsletters, social networks, personal meetings and working groups. Given that stakeholders' engagement is difficult to maintain, appropriate communication actions should take place from the very beginning. For example, public servants must be informed on how technological change will benefit them in order to avoid change resistance. Also, citizens should be an active part of the technological evolution of their city. Citizens must be informed that their Municipality is involved in a technological evolution process that will produce benefits for the city in terms of quick availability of new services for citizens, cost reduction, improved flexibility and transparency in procurement processes. This information will elevate the sense of transparency and accountability.

c) the use of a series of iterative innovation cycles.

The idea for these cycles lies on the interactive nature of innovation. It starts with a preparation stage, where the innovation environment is established, describing the scenario and the interaction between the key actors, followed by iterative innovation cycles in which the services are being evaluated by users, leading to improved versions.

3.2 Setup a Monitoring and Validation Process

The monitoring and validation process, for the successful migration of the selected applications to the Cloud, targets the business aspects of the applications rather than the technological ones. This approach is more holistic as the successful migration in business terms implies the success of the technical one.

The STORM CLOUDS approach

The process consists of three different steps: i) identifying the aspects to monitor and the specific indicators or criteria (depending on the task), ii) information gathering throughout the entire process of cloudification, and iii) analysis of the usage and acceptance of the new applications and/or variations on the usage patterns. The main indicators that usually apply to this process are the following: indicators monitoring the supply side of the service, indicators monitoring the demand side of the service, indicators related to dissemination, indicators related to validation of the service and, finally, indicators showing the financial benefits of migrating an application to the cloud.

As an example, the following two figures present the indicators for "Virtual City Market" and "CloudFunding" applications, which have been cloudified for the Municipality of Thessaloniki. Some of the indicators are common between the two applications, mainly those that are related to the dissemination and validation aspects. On the contrary, the indicators for the supply and demand sides are mainly different as they are tailored to the context of the applications.

Supply	Demand	Dissemination	Validation
Nbr of shops participating in the app	Total nbr of users – visitors	Total presence of the platform in third party websites	Number of users providing feedback for the application
Nbr of shops per category	Total nbr of registered users	Total e-mails/newsletters sent	Number of stakeholders providing feedback for the application
% of shops participating in the platform/shops in the area (total)	Mean nbr of visitors per shop		Number of modifications (new characteristics that have been modified based on the feedback received)
% of shops participating in the platform/shops operating in the area (category)	User demographics (area, age, education level)		
Nbr of shops that have extended their online presence in the platform			
Nbr of shops making online transactions through the platform			
Nbr of offers per shop			
Nbr of synergies between two or more shops			

Figure 3 – Monitoring and validation indicators for the Virtual City Market application

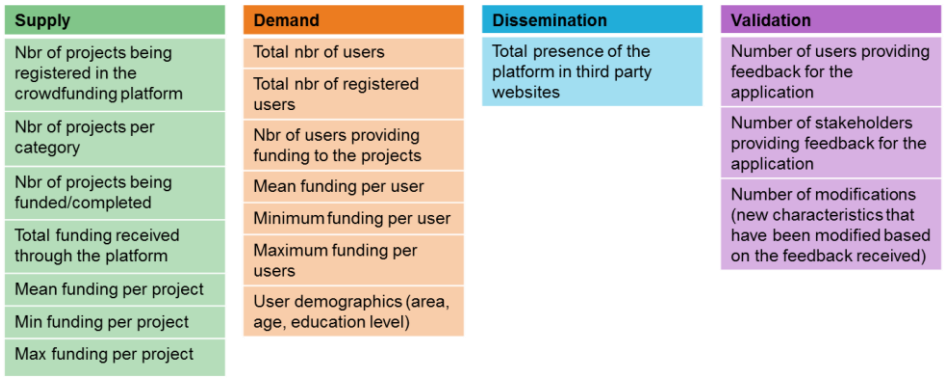

Figure 4 – Monitoring and validation indicators for the CloudFunding application

3.3 Prioritize Applications that Should Move to the Cloud

As Public Authorities start migrating their applications to the Cloud, it is important to determine which applications fit better into this environment. The best candidates are applications, which take advantage of the elasticity of Cloud Computing. In particular, the following type of applications will benefit from Cloud's ability to automate the dynamic of resources to match the current demand[[20]]:

- Applications that are designed to spread their workload across multiple servers.
- Applications that run occasionally but require significant computing resources when they run.
- Applications with unpredictable or cyclical usage patterns.
- Service Oriented Architecture (SOA) Applications.

For these type of applications, the rapid elasticity combined with the pay-by-usage characteristic of the cloud can lead to significant financial savings.

On the contrary, the following types of applications are less suitable candidates for Public Cloud deployments:

- Applications that include extremely sensitive data, particularly when there is a regulatory or legal risk involved in any exposure. These applications can benefit from a hybrid deployment model, where application code can reside in a public cloud environment while sensitive information can be deployed in a private cloud infrastructure.
- Performance-sensitive applications.
- Applications that require frequent and/or voluminous transactions against an on-premises database that cannot be migrated to a cloud environment.
- Applications that run on legacy platforms that are typically not supported (or may not be supported in the near future) by the cloud providers.

The STORM CLOUDS approach

All the applications, which have been cloudified in STORM CLOUDS pilot cities, are good candidates for cloudification as they expected to take full advantage of the elasticity of Cloud Computing. In particular, some of them (Virtual City Market, City Branding, Live the City and Public Failure Reporting System) are used initially by a small number of citizens and therefore have low workloads, but it is expected to be used by the majority of city residents and in that case will have high workloads. The remaining applications (CloudFunding, Location Plans and Have your Say) have high workloads and occasionally require significant computing resources (i.e. when new plans are set for discussion, new crowdfunding campaigns are available, etc.).

3.4 Choose the Right Cloud Service Category

Public Authorities should consider, when they plan their Cloud strategy, the deferent service categories of Cloud Computing [7]:

- Software-as-a-Service (SaaS): The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a programming interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.
- Platform-as-a-Service (PaaS): The capability provided to the consumer is to use the provider's development platform (programming languages, libraries, services, and tools) in order to create, test and host new applications. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.
- Infrastructure-as-a-Service (IaaS): The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources in order to build a customized computing environment. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of selected networking components (e.g., host firewalls).

The figure on the right depicts the various control responsibilities that cloud customers and providers have in IaaS, PaaS, and SaaS environments [21].

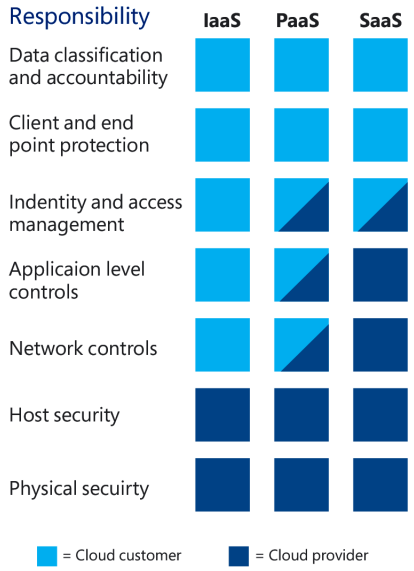

Figure 5 – Considerations for cloud service

choice

SaaS is not applied in case that the Public Authorities want to deploy their existing applications to a Cloud Environment. In that case, they have to select between IaaS and PaaS. To decide which of the two options (IaaS or PaaS) they will follow, the Public Authorities should evaluate the pros and cons of each solution. On the one hand, the IaaS offers excellent flexibility, as it does not require architectural changes to the applications, and full control of the resources used for the deployment. However, it increases the deployment complexity, as the application owners must take care of installing and configuring all the components for high availability and scalability. On the other hand, the PaaS "hides" the complexity of the underlying infrastructure and allows developers to deploy their web applications to the cloud without having to take care of the infrastructure. However, the applications may require significant changes to comply with the PaaS principles and take full advantage of high availability and scalability features. In particular, as application instances are ephemeral and can be started, stopped or fail at any time, they must be stateless and share nothing. All persistent data must go to external services (e.g. databases, file storage, message queues and caches) [22].

The STORM CLOUDS approach

The STORM CLOUDS Platform enhances the IaaS solution with two modules that provide the high-availability and scalability features in a way that is transparent to the application owners. By accompanying the IaaS layer with the Data Service layer and Access layer, the data and the HTTP traffic management are delegated to the platform while the application's business logic is still contained on the VM(s). This approach offers great flexibility as it does not require architectural changes to the applications but also keeps the deployment complexity low because the application owner "leverages" the high-availability and scalability features of the platform. The only drawback of this solution comparing with the SaaS is that the application owners are not entirely independent from platform administrators as the later should configure the high-availability and scalability features per application.

The following table summarises the different application migration options supported by SCP [16].

Option	Description	Pros	Cons
Full IaaS	All the application components are deployed on VM(s) explicitly managed by the application owner	+ No architectural change of the application + Full control on the resources used for the deployment	- High deployment complexity because the application owner must take care of installing and configuring all the components for high availability and scalability
IaaS + Data Service	Data and HTTP traffic management are handled by the	+ No architectural change of the application + Less deployment	- Because of the centralized administration of the

Layer + Access Layer	platform, while the application business logic is still deployed on VM(s)	complexity because the application owner 'leverages' the high-available and scalable features of the platform layers	shared functions (e.g. data service layer), application owners cannot deploy their applications in full autonomy
PaaS + Data Service Layer	Applications are hosted on the PaaS Layer and use the Data Service Layer for storing data	+ No infrastructure management required by the user (the platform does it for her)	- Applications may require significant changes to comply with PaaS principles

Table 2 - Different application migration options supported by STORM CLOUDS Platform.

3.5 Chose the right Cloud Deployment Model

Public Authorities should consider, when they plan their Cloud strategy, the different Deployment Models of Cloud Computing [7]:

- Private Cloud: The cloud infrastructure is used exclusively for internal applications within an organisation comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
- Community Cloud: The cloud infrastructure is used exclusively by multiple organizations that have shared concerns (e.g., mission, security requirements, policy and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
- Public Cloud: The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
- Hybrid Cloud: The cloud infrastructure is a composition of two or more distinct cloud deployment models (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g. cloud bursting for load balancing between clouds).

The following table summarises the pros and cons of the different deployment models [23, 24].

Option	Pros	Cons
Private Cloud	+ More control and reliability: IT can control the security of data, set compliance requirements, and optimize networks more effectively with cloud. + Customizable: IT can customize storage and networking components so that the cloud is a perfect fit for the specific organization and its needs.	- Requires IT expertise: A high-level of IT expertise is required to ensure maximum effectiveness and optimal configuration of the deployment. - Costlier: The long-term costs may be higher due to increased management responsibilities and smaller economies of scale.
Public Cloud	+ Ease of management: Organisations IT departments do not manage their public cloud; they rely on Cloud provider to administer the cloud. + Ease of deployment: With the public cloud, there is low barrier to entry, so you can quickly configure and stand up	- Can be unreliable: Public cloud outages are quite common, leading to headaches for users. - Less secure: The public cloud often has a lower level of security and may be more susceptible to hacks. In some

	a cloud. + Flexible: Users can add or drop capacity easily. Moreover, the environment is typically accessible from any Internet-connected device, so users don't need to jump through many hurdles to access.	cases, cloud providers may not be able to meet the strict constraints mandated by government institutions.
Hybrid Cloud	+ Flexible and scalable: Organisations are able to combine and match for the ideal balance of cost and security. + Cost effective: Organisations can take advantage of the cost-effectiveness of public cloud computing, while also enjoying the security of a private cloud.	- Complexity of management: Moving parts between public and private clouds can be a challenge. - Requires IT expertise: A high-level technical staff is required to guarantee security vulnerability on all aspects is decreased.

Table 3 - Pros and cons of private, public and hybrid deployment Cloud models

In short, when choosing a specific cloud deployment model, it comes down to a series of trade-offs related to cost, management and security. While public clouds may be the best option for small organisation from a cost perspective, organizations that require more control and/or security may opt for a private or hybrid cloud — providing they have the manpower and budget to manage those deployments effectively.

The STORM CLOUDS approach

The STORM CLOUDS Platform offers great flexibility to the cities, regarding the Cloud deployment model that they wish to follow. During the STORM CLOUDS project, the SCP operates both in a Private Cloud environment (hosted on Hewlett Packard Enterprise's IT infrastructure), as well as in a Public Cloud environment (hosted on a commercial cloud provider).

Consequently, SCP supports Public, Private, as well as Hybrid Cloud deployment models. In case of the Private Cloud, the cities can install the SCP in their IT infrastructure and use it as the Cloud environment for their Smart City applications. Alternatively, cities can use an instance of the platform that is offered by an external provider and deploy their applications in the Public Cloud. A hybrid approach is also supported, as they can combine a Private Cloud, which will host the high-risk applications—those with high privacy and security requirements (i.e. applications that contain customer data and other sensitive information)— with a Public Cloud for the rest of them.

SCP can also be the foundation for the creation of a Community Cloud that will support not only a Municipality but also other Public Organisations, which operate in a city and offer services to citizens and local business.

3.6 Embrace the Power of Open Technologies

Being open is about adopting the technology decisions that organizations have made and giving them the freedom to move across technologies, models and cloud providers. Systems composed of open technologies provide the freedom to change environments and deliver a robust and secure experience extending existing IT to the cloud. They enable customers to do more work with less infrastructure, deliver a broader range of services, incorporate new technologies and boost greater innovation around the cloud [25]. The majority of existing cloud offerings are implemented in proprietary and highly standardised form. What presents advantages for the provider – technological knowledge, economies of scale, etc. – creates troubles and frustration for the customer. Users complain of "vendor lock-in", where they are dependent on a given vendor with no freedom of choice. Embracing an open cloud means there is no technology lock-in, no contractual lock-in and no service lock-in. It means providers don't dictate technologies and that competition is embraced [26]. New, emerging standards will increase the portability and interoperability of systems across cloud service providers, and will reduce or eliminate this current barrier to cloud adoption.

The STORM CLOUDS approach

The STORM CLOUDS Platform is built upon widely accepted open source technologies. Moreover, its architecture is a baseline for future extensions and modifications with the objective to allow developers to improve the way functions are implemented or to add new features not currently available.

The SCP was developed using the following open source solutions:

- ✓ OpenStack for the implementation of the IaaS Layer. OpenStack is the most popular and most adopted open source IaaS solution [27].
- ✓ Cloud Foundry for the implementation of the PaaS Layer. Cloud Foundry was chosen because had the best combination of usability, open-source community, developer experience, and relation to SCP's needs. It is supported by the Cloud Foundry Foundation where EMC, HP, IBM, Intel, Pivotal, SAP and VMware are platinum members.
- ✓ LAMP (Linux, Apache, MySQL and PHP) for the implementation of applications' VMs.
- ✓ MySQL/MariaDB and PostgreSQL database engines for the implementation of Database Services Module.
- ✓ Gluster [28] for the implementation of file Sharing Service Module.
- ✓ HAProxy [29] for the implementation of Load Balancer Module.
- ✓ Zabbix [30] for the implementation of the Monitoring Module
- ✓ phpMyAdmin [31] for the implementation of the MySQL Database Administration Module
- ✓ phpPgAdmin [32] for the implementation of the PostgreSQL Database Administration Module
- ✓ Duplicity [33] for creating the backups.

The implementation of SCP on open source technologies will not lock the organisations that use it into a proprietary ecosystem and thus made it extremely hard to move their application to another provider. The development of custom APIs was avoided and all the components, including the database tier, storage tier, and any micro service endpoints, were created using open source tools. This approach encourages organisations to continue to use the platform because they want to, not because they have to. Another significant advantage of the use of open source solutions is the fact that the platform will continue to benefit from the improvements in the operability and security of OpenStack, Cloud Foundry and all the other tools. It can scale without significant development effort. Also, the selection of broadly adopted software packages guarantees the long-term support of the solution.

3.7 Plan Carefully and Automate the Migration

Application migration is the process of redeploying an application, typically on newer platforms and infrastructure. Comprehensive planning, driven by a disciplined migration process will contribute greatly to a successful redeployment of the applications to a new cloud environment. Successful initiatives have developed sophisticated, multi-phased migration methodologies to reduce implementation risk and speed-up the migration process. The diagram on the right illustrates a typical migration project life cycle [34]

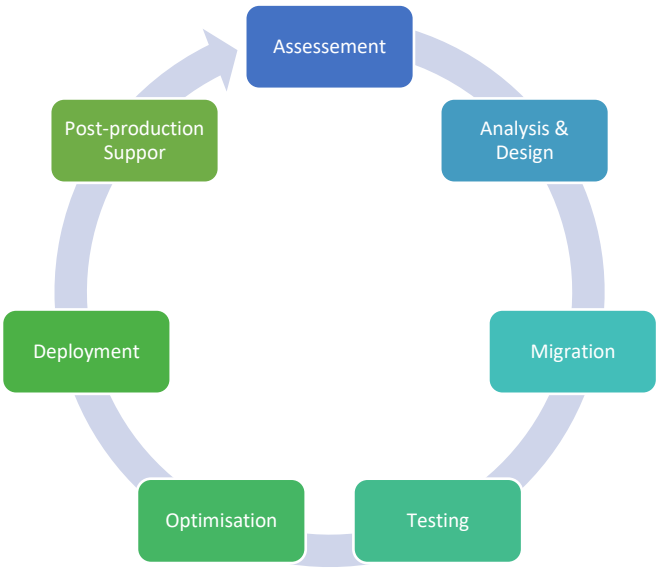

Figure 6 - A typical migration project life cycle

During the migration process the following technical considerations must be taken into account [10,20]:

- The creation of a detailed inventory of the current application portfolio really helps in terms of understanding the scope of the migration effort. This includes capturing information regarding the number of software modules, scripts, and external interfaces involved. It also includes hardware and software configuration information, including operating system versions, database versions, features/functionalities in use, and similar information.
- A security audit of the application and its data is vital. The cloud service's security features may be very different from those of the in-house environment, and the security risks and the measures applied to counter them must be assessed carefully.
- Temporary subsystems can be established to facilitate migrations.
- Standardization and automation can help reduce the risk of migration errors. Virtual machine templates can be rapidly deployed and bring an environment online in a day. Automated data integrity and validation methods can be used to verify and validate data, databases and files during the initial synchronization.
- The creation of migration tools, which ensure a high level of automation along with accuracy in migration can result in less time spent in migration and testing.

The STORM CLOUDS approach

STORM CLOUDS followed a multi-phase migration process, which included all the necessary steps that ensured the smooth deployment of the selected Smart City applications to the SCP. The process started with the assessment of each application regarding its readiness for the new Cloud environment, its architecture and its functional and non-functional requirements. This analysis leaded to some necessary improvements in order the application to be optimised for the SCP. Afterwards, the code and data deployed in the platform's Application and Data Service Layers, respectively. The process was completed with the validation that the application was fully operational in the new Cloud environment.

The following diagram presents the STORM CLOUDS migration process.

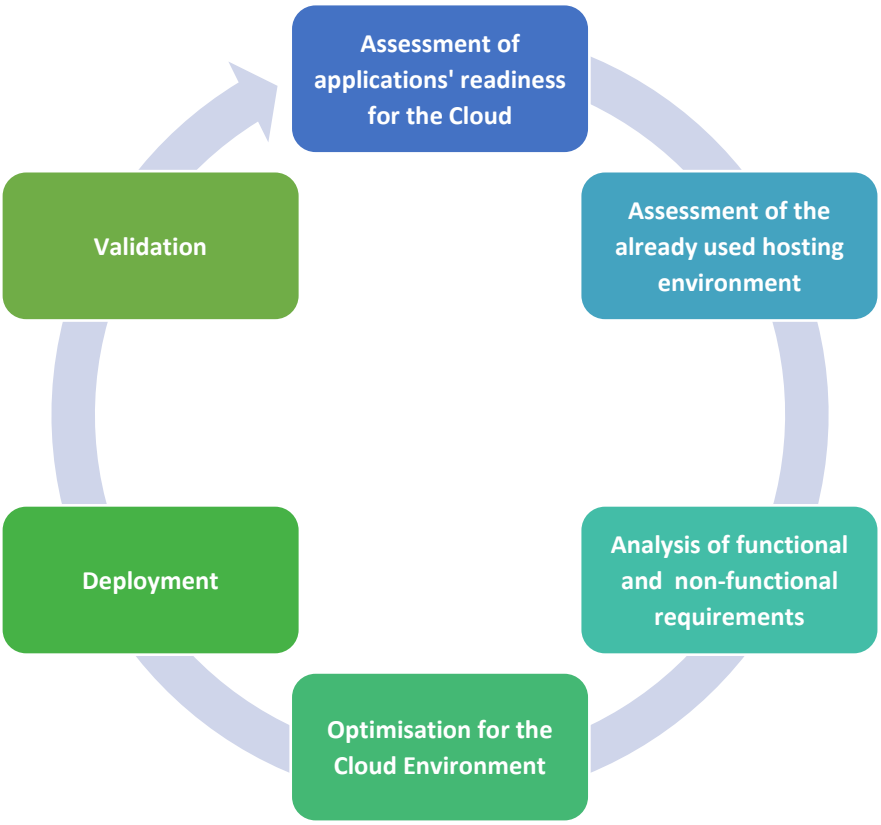

Figure 7 - STORM CLOUDS migration process

The STORM CLOUDS migration process includes the following six steps:

Step 1 - Assessment of applications' readiness for the Cloud

The 1st step aims to evaluate if the services are ready for the cloud environment. Aspects such as customization, regulatory compliance, complex service architectures and service maturity are carefully investigated, as they would negatively impact the cloudification process. A crucial aspect is the availability of both the application's source code and documentation (installation manual, code dependencies, required software packages, etc.). Finally, the commitment of the application's development and support team should be ensured.

Step 2 - Assessment of the already used hosting environment

The 2nd step aims to analyse the environment used to host the services. The analysis covers both the network (e.g. configuration, connectivity requirements from the municipality premises to the cloud environment, and supplementary services such as SMTP, DNS and WWW) and architecture (e.g. use of resources, underlining technologies, licenses, and security mechanisms) of the service.

Step 3 - Analysis of functional and non-functional requirements

The 3rd step aims to define the technical characteristics of the Virtual Machines that will host the applications on the new Cloud Environment. The analysis of the functional requirements covers technical details (e.g. Operating System, Scripting Language, Database, Web/Application Server, Data Formats, Frameworks/Libraries and External Services used), interoperability issues, and static characteristics such as hard-coded IP address and directory paths. Furthermore, the analysis of the non-functional requirements addresses issues related to the proper functioning of the application such as security, regulatory compliance, performance, availability, backup; privacy, reusability, and interoperability. An estimation of the use of resources regarding RAM, Disk Space, CPUs, Bandwidth, Hits/Month, Registered Users, Max On-line Users, and Average On-line Users contributes to the calculation of the expected workload per application. An important characteristic that should be examined in this step is if the application's design supports its deployment in multiple servers. In that case the application will take full advantage of the performance benefits that cloud offers.

Step 4 - Optimisation for the Cloud Environment

The 4th step aims to solve the problems identified in the previous step, so the application to be ready for deployment in the new environment. Moreover, it includes modifications that enable the application to support natively the most prominent Cloud characteristics (e.g. high-availability and scalability). The latter is closely related to the available budget or the internal IT capabilities.

Step 5 – Deployment

The 5th step aims to transfer the ready to be cloudified applications to the new Cloud environment. The deployment process includes the following actions:

a) setup of the cloud environment that will host the selected services;

b) launch the VM instances that will host the applications and their data (e.g. database and file sharing modules).

c) migrated both the applications and their data to the Cloud environment

Step 6 – Validation

The final step aims not only to ensure that the deployed applications are operational but especially that they meet the initial set of requirements regarding cloudification. The validation is made in collaboration with the municipalities and includes functional tests ensuring that the deployed application performs as designed.

Automation

The cloud encourages automation because the infrastructure is programmable. To ensure a high level of automation along with accuracy in the migration of the applications to the cloud, a set of tools

and procedures have been designed and developed. Automatic deployment is implemented using Heat, the OpenStack "orchestration engine to launch multiple composite cloud applications based on templates in the form of text files that can be treated like code" [32]. The aim of orchestration is to create a human -and machine- accessible service for managing the entire lifecycle of infrastructure and applications within the SCP Cloud environment.

The 1st step in the automation process is to prepare the bash shell scripts that will configure the VM hosting the application, and install and configure the application and its dependencies.

The 2nd step is to create the Heat scripts (Heat Templates) that describe the infrastructure (servers, floating IPs, security groups, ports) of the cloud applications and to integrate with them the application's installation and configuration scripts made at the previous step.

The available Heat Templates allow interested cities to automatically deploy the selected applications from the cloud-based service portfolio, as well as the municipalities to re-deploy their services in another instance of STORM CLOUDS Platform.

3.8 Use the Right Tools to Manage and Monitor the Cloud Environment

Cloud Computing imposes new concepts and challenges for the role of monitoring and management of the Cloud environment and the smart city applications. The System Administrator no longer needs to provide servers, install software and wire up network devices since all this work is replaced by few clicks and command line calls. Nowadays, most of the daily tasks performed by system administrators are related with the applications. One of the characteristics of the Cloud, which facilitates the agile deployment of the applications, is the fact that administrators don't have to master the art of capacity planning because they can create an automated elastic environment [9]. If they can understand, monitor, examine and observe the applications' load and traffic patterns, they will be able to manage this elastic environment more effectively. Moreover, by measuring and monitoring the performance of the cloud applications, the application developers will have the opportunity to identify proactively any performance issues and to diagnose the root causes, so they take appropriate actions.

The Cloud environment should offer both to system administrators and application owners the necessary tools required to manage and maintain the platform and the deployed applications. Using these tools, they can focus on how to optimize the cloud-based application in order to increase cost savings. The "pay for what you use" approach of the Cloud, leads application owners to strive to optimize the system whatever possible. Even a small optimization might result in thousands of euros of savings.

The STORM CLOUDS approach

The STORM CLOUDS Platform includes features that both the platform administrator and the application owners can use for managing, monitoring and administering the platform's components as well as the applications running in the cloud. The actions that a user can perform, depend on his/her role: the platform administrator has full control on all the components deployed in the cloud while application owners have full control of their applications and can perform only some actions on the platform components. For instance, application owners have full control over databases and shared volumes used by their applications but they do not have any control on databases and shared volumes used by other application owners. The following tools management and monitoring tools are available:

- The Platform Administrator's Console, which allows the SCP administrator to have full control of the layers of the platform. Through the console, (s)he can manage the databases, the filesystem, the IaaS layer, and the PaaS layers.
- The Database Administration Console, which allows administrators and applications' owners to administer the supported databases. The module includes phpMyAdmin for MySQL administration and phpPgAdmin for PostgreSQL. Both tools implement very similar functions for the corresponding database engines like creating, modifying and deleting

database users, databases and database objects (e.g. tables, indexes, etc.), submitting queries, importing/exporting data, managing database accounts, etc. The platform's administrator has full control of all databases and configures database accounts for the application owners, giving them the rights of managing only the database objects created for their applications.

- The Monitoring Console, which monitors the resources (CPU load, disk space occupation, network traffic, number of processes, etc.) used by the platform's services or by the applications. The module, implemented using Zabbix [30], continuously gathers information from the servers under control and, in case one or more parameters reach a threshold value, it notifies the operator by e-mail, Instant Message or SMS. Zabbix offers several monitoring options ranging from simple checks for verifying the availability/responsiveness of a server, to sophisticated measurements of parameters like CPU load, disk volume occupation, network traffic, number of processes, etc. Zabbix provides several ways for representing monitoring data in both graphical and textual/tabular format.

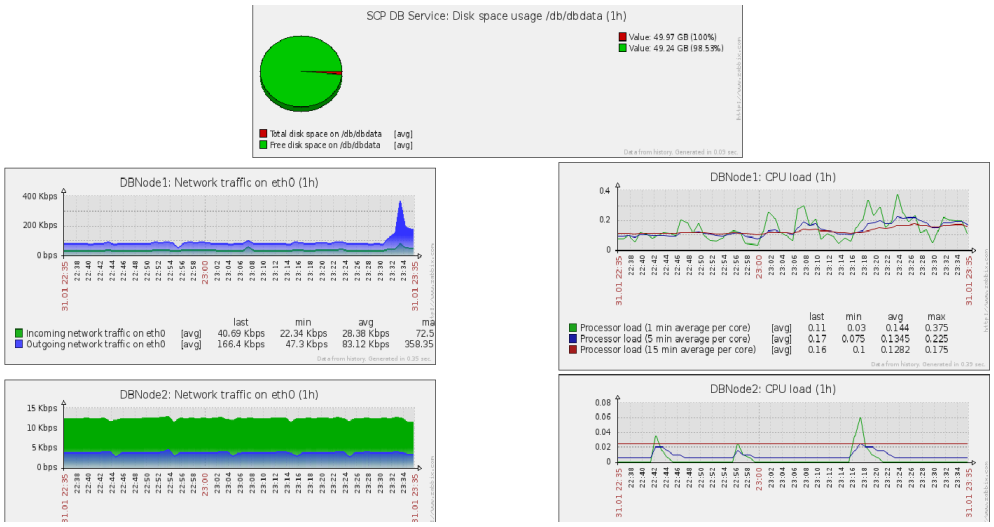

Figure 8 – Zabbix Monitoring Pages

- The Backup Tool, which takes backups from databases and file systems. This module is implemented using Duplicity, an application that creates encrypted bandwidth-efficient backups using the rsync algorithm.
- The Automation Tool, which facilitate the application owners to deploy the Smart City applications to the IaaS layer automatically using a number of predefined Heat scripts.

3.9 Focus on Security

Cloud computing security is an evolving sub-domain of information security and refers to a broad set of policies, technologies, and controls deployed to protect data, applications, and the associated infrastructure [35]. There are a number of security concerns associated with cloud computing, which can be broadly classified in two categories: (a) issues faced by Cloud Service Providers (CSPs) and (b) issues faced by their customers. Providers must ensure that their infrastructure is secure and clients' data and applications are protected; customers, on the other hand, must ensure that their provider has taken appropriate security measures to protect their information. The security expectations and obligations of both supplier and user are described in Service Level Agreements (SLAs) [36].

Organisations need to understand the specific security requirements, regarding data protection, audits, etc., and any regulations that are applicable to a particular application that they are looking to move to the cloud. To achieve this, they should map every application that is a candidate for migration to cloud computing to a set of security, governance, and compliance issues that are specific to that application. Thus, they have the ability to understand the application requirements, and how the migration and re-development effort to the cloud should impact application operations.

The UK's National Technical Authority for Information Assurance, which provides advice on Information Assurance Architecture and cyber-security to UK government and the wider public sector and suppliers to UK government, published 14 security principles to consider when evaluating cloud services, and why these may be important to an organisation [37].

Cloud Security Principle	Description
1. Data in transit protection	Consumer data transiting networks should be adequately protected against tampering and eavesdropping via a combination of network protection and encryption.
2. Asset protection and resilience	Consumer data, and the assets storing or processing it, should be protected against physical tampering, loss, damage or seizure.
3. Separation between consumers	Separation should exist between different consumers of the service to prevent one malicious or compromised consumer from affecting the service or data of another.
4. Governance framework	The service provider should have a security governance framework that coordinates and directs their overall approach to the management of the service and information within it.
5. Operational security	The service provider should have processes and procedures in place to ensure the operational security of the service.
6. Personnel security	Service provider staff should be subject to personnel security screening and security education for their role.
7. Secure development	Services should be designed and developed to identify and

	mitigate threats to their security.
8. Supply chain security	The service provider should ensure that its supply chain satisfactorily supports all of the security principles that the service claims to implement.
9. Secure consumer management	Consumers should be provided with the tools required to help them securely manage their service.
10. Identity and authentication	Access to all service interfaces (for consumers and providers) should be constrained to authenticated and authorised individuals.
11. External interface protection	All external or less trusted interfaces of the service should be identified and have appropriate protections to defend against attacks through them.
12. Secure service administration	The methods used by the service provider's administrators to manage the operational service should be designed to mitigate any risk of exploitation that could undermine the security of the service.
13. Audit information provision to consumers	Consumers should be provided with the audit records they need to monitor access to their service and the data held within it.
14. Secure use of the service by the consumer	Consumers have certain responsibilities when using a cloud service in order for this use to remain secure, and for their data to be adequately protected.

Table 4 - Cloud Security Principles (Source http://goo.gl/mUf5c2)

Consumers of cloud services should decide which of the principles are important, and how much assurance they require in the implementation of these principles, while providers of cloud services should consider these principles when presenting their offerings to public sector consumers. This will allow consumers to make informed choices about which services are appropriate for their needs.

The STORM CLOUDS approach

In order to achieve a clear understanding of the security requirements of both the SCP and the Smart City applications, the following vulnerability scanning tools were used to scan web applications to look for known security vulnerabilities:

- Zed Attack Proxy (ZAP) [38], an easy to use integrated web applications vulnerability scanning tool;
- OpenVAS [39], a framework of several services and tools offering a comprehensive and powerful vulnerability scanning and vulnerability management solution;
- SQL Inject Me [40], a Firefox Extension used to test for SQL Injection vulnerabilities
- Qualys SSL Server Test [41], an online service performing deep analysis of the configuration of any SSL web server on the public Internet;

- Vega [42], an open source scanner and testing platform to test the security of web applications.

The security testing identified a number of critical security issues resulting in applications' modifications in order to address them. In particular, the following issues were fixed:

- "Cross Site Scripting" (XSS) security risk, being the most prevalent web application security flaw, whereby an attacker's malicious content is supplied to our application as a result of that content not being properly validated or escaped. To address it we used the OWASP ESAPI reference implementation for HTML entity escaping and unescaping as well as JavaScript escaping and unescaping.
- "Directory listing" security risk, whereby an attacker can simply list directories to find files. To address it we have updated the Apache configuration file by removing the *Indexes* from the file.
- "Insufficient Transport Layer Protection" security risk, caused by not requiring SSL (at least for all sensitive pages) allowing an attacker that monitors our network traffic to obtain an authenticated victim's session cookie, itself then replayed to take over the user's session. To address it we have included an HTTPS certificate and requested that all traffic is forwarded to the secure connection (HTTPS). However, new vulnerabilities introduced by the HTTPS certificate, such as the RC4 cipher and the POODLE attack vulnerability, resulted in the disabling of:
    - TLS 1.0 compression and weak ciphers
    - SSL 3 in our browser and our servers
- "Clickjacking" security risk, caused by an attacker that is "hijacking" our clicks meant for the application page and routing them to another malicious page. To address it we send proper X-Frame-Options HTTP response headers that instruct the browser to not allow framing from other domains. This is done at the Apache configuration file.

Regarding the use of the above-mentioned tools, one should always check the terms of service of the selected CSP to determine whether running security tests on the CSP infrastructure is allowed, even if own machines are the target. If this is not so, either a CSP that allows penetration tests on own VMs must be chosen, or have tests run on a development or testing environment before deployment to the production environment.

In order to enhance the authentication process, applications have been updated to support session expiration, thus minimizing the time available to an attacker who uses a valid session identifier. In order to balance between security and usability, applications have properly selected the timeout values, allowing users to complete their operations without frequent session expirations.

The acquisition of SSL certificates is necessary for protecting *data in* motion. The Apache server was configured to forward all traffic to the secure connection. However, as applications don't deal with sensitive data no encryption is applied for data at rest, apart from the OpenStack object storage that is protected using LUKS (Linux Unified Key Setup or LUKS is the standard for Linux hard disk encryption).

Virtualization technologies have their own vulnerabilities such as those coming from the virtual switch, those coming from reallocation of resources from one VM to another and vulnerabilities coming from the remote administration port that is turned on by default on all VMs. To respectively address these attack vectors:

- We've used layered security mechanisms to increase the security of the system as a whole. This was achieved using OpenStack security groups in order to define a number of IP firewalling rules that describe what kind of network traffic is allowed to go to or come from the VMs. With this solution even if a VM is compromised the security group rules continue providing the required level of security because they are implemented in the host operating system.
- We've used OpenStack functionality for zeroing all data used by a virtual resource once the resource is released.
- We associated each VM with a valid SSH Keypair. This was then forwarded to the application owners in order to allow them to access the VM instances given the public IP address the VM is configured to use.

Finally, securing our cloud infrastructure means not only implementing controls for the layers we are able to do so, but also auditing our CSP regarding actions taken to lock-down the tenant instances. We must conduct our own analysis of our needs, and assess, select, engage and oversee the cloud services that can best fulfil those needs.

3.10 Protect your Data

We define customer data as "*all the data, including all text, sound, software or image files that a user provides, or are provided on the users' behalf, to the cloud provider through use of the online services.*" [43] Data protection is the process of safeguarding important information from corruption and/or loss [44]. Cloud providers should commit to protecting the data and limit the use of them. The data that that Public Authorities host in cloud services belongs to them—and should not be used by a cloud provider for purposes other than to provide the customer's service. Moreover, cloud providers should not use customer data for purposes unrelated to providing the service, such as advertising. Additionally, each service has established a set of standards for storing and backing up data, and securely deleting data upon request from the customer.

The best-designed and implemented service cannot protect customer data and privacy if it is deployed to an environment that is not secure. Customers expect that their data will not be exposed to other cloud customers. They also assume that the processes used at the datacentre, and the people who work there, all contribute to keeping their data private and secure.

The term data protection is also used to describe operational backup of data that usually comes in the form of incremental backups. The aim of the backup procedure is to keep data from being lost due to intentional or unintentional access.

The STORM CLOUDS approach

The STORM CLOUDS backup process is based on the data requirements of the services and the architecture of the SCP. The backup process aims to best exploit the features implemented by the IaaS cloud where the VMs are hosted and more specifically Swift, the Object Storage Service implemented by OpenStack. The main steps needed for backing up application's data are presented below.

1st Step: Design a Backup Strategy

During this step, several aspects related to the data and/or the application(s) managing the data were analysed in order to put together a list of what needs to be backed up, when to backup, how long to keep the backup data and how long it takes to restore. It includes the following tasks:

- Analysis of current data usage that reveals:
  - Types of data used.
  - Data locations, including folders and/or databases.
  - Approximate amount of data.
  - How often data changes, as this affects our decision on how often the data should be backed up.
  - Data sensitivity. For critical data, such as a database, we should have redundant backup sets that extend back for several backup periods. For sensitive data, we should ensure that backup data is encrypted, using public/private key-pair technology.

  - How quickly we need to recover the data.
  - What's the best time to schedule backups (scheduling backups when system use is as low as possible will speed up the backup process).

- Set an up limit for the backup volume as the amount of data we need to backup is only going to increase as time goes by.
- Identify the software tools that will be used
- Select the appropriate backup type/policy (Full or Incremental). Typically, one of the following approaches is used: (a) Full daily, (b) Full weekly + Incremental daily. The process of taking incremental backups following an initial full backup is known as data deduplication. The final choice depends on the required performance levels and data protection levels, the total amount of data retained and the cost associated with it, since cloud storage space comes at a cost that depends on the service provider.
- Choose where to store the backups. Using the cloud environment to store the backup data is arguably more resilient to disaster than other technology solutions because it is not physically located at the same place as the organisation. Moreover, since the applications are hosted in the Cloud we also save bandwidth and time taken to transfer the files needed to restore the application correctly. However, the cost associated with storing the backup data in the cloud is a significant factor in our decision.

2nd Step: Generate a Key-Pair on the Client Machine

Although we can create the key pair directly on the VM, it is good practice to keep a copy of the keys outside the VMs using them. The reason is that VMs are "ephemeral", meaning that once a VM is deleted, we are not anymore able to decrypt our backup data when restored. Moreover, creating key-pairs requires some level of "entropy" for ensuring randomness in the generation.

3rd Step: Prepare the VMs for Backup

Install and configure

4th Step: Implement the Backup Strategy

The backup scripts that address all the aspects of backup strategy are created and executed using the Duplicity tool.

5th Step: Validation tests.

Validation includes tests on the restore mechanism. More specifically both incremental and full backups were used to bring the applications to a previous operational state successfully. The backup solution should be tested many times after it has been implemented in order to ensure that it is working as intended. Moreover, the applications should be re-tested periodically to ensure they're functional, and data is being backed up appropriately. Validation not only will help us to identify problems in the backup process but will also train the Municipalities' IT personnel to recover quickly and efficiently the files if this becomes necessary.

After the initial setup the backup process is scheduled according to the backup strategy.

3.11 Assess and Improve the Interoperability Maturity of a Public Service

Interoperability is "*the ability of two or more systems or applications to exchange information and to mutually use the information that has been exchanged*" [45]. In the context of Cloud Computing, interoperability can be further described as the "*capability of public clouds, private clouds, and any other systems in the enterprise to understand each other's application and service interfaces, configuration, forms of authentication and authorization, data formats etc. in order to cooperate and interoperate with each other*" [46]. In our case, interoperability could be understood as how well a public administration service interacts with external entities in order to organise the efficient provisioning of its public services to other public administrations, businesses and or citizens.

The European Commission's ISA (Interoperability Solutions for *European* Public Administrations) programme [47] developed an Interoperability Maturity Model (IMM) [48] to provide public administrations insight into two key aspects of their interoperability performance:

- The current interoperability maturity level of a Public Service;
- Improvement priorities to reach the next level of interoperability maturity.

The IMM helps owners of a Public Service to enhance the quality of the service delivery, reduce costs and overcome integration issues by reusing available services and orchestrate services in an effective manner to maximise service outcome and benefits for citizens and public administrations [49].

In the context of interoperability maturity, the IMM measures how well a public service is able to interact with other organisations to realise mutually beneficial and agreed common goals through the exchange of information and reuse of services. Three different domains of interoperability are distinguished:

- Service Delivery – Providing end-users accessibility to the public service
- Service Consumption – Consumption of reusable services from other public administrations and businesses. This can include the consumption of functionalities, base registry information and security services
- Service Management – Controlling and monitoring the process flow related to external service interactions from trigger to outcome

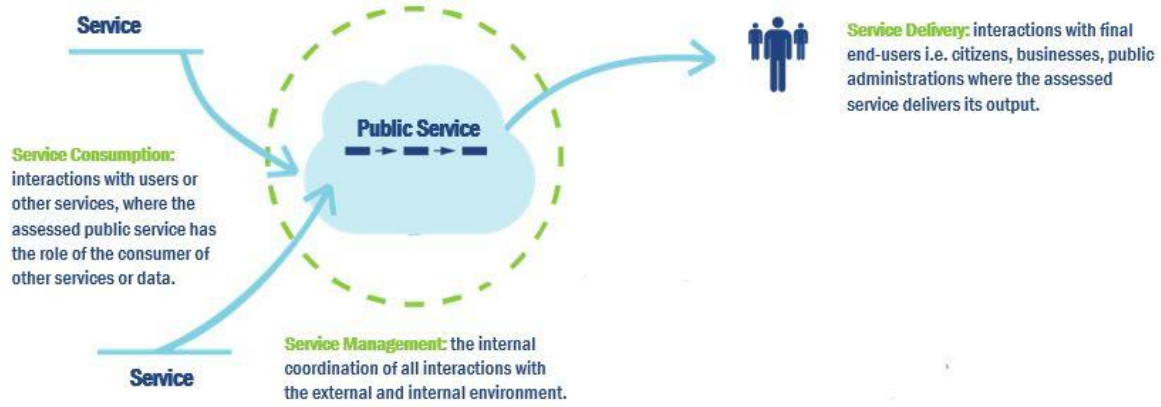

Figure 9 – The three different domains of interoperability in IMM (Source: European Commission)

The IMM uses a five-stage model to indicate the interoperability maturity of the public service. The reason for the usage of these various maturity levels is two-fold:

- To measure the interoperability maturity of the public service as a whole and of the underlying aspects;
- To indicate which capabilities and next steps are required to improve interoperability maturity.

The five maturity levels for the IMM are summarised in the table below:

Maturity level	Maturity stage	Interpretation
1	Ad Hoc	Poor interoperability – the service has almost no interoperability in place
2	Opportunistic	Fair interoperability – the service implements some elements of interoperability best practices
3	Essential	Essential interoperability – the service implements the essential best practices for interoperability
4	Sustainable	Good interoperability – all relevant interoperability best practices are implemented by the public service
5	Seamless	Interoperability leading practice – the service is a leading example for others

Table 5 - Five maturity stages of IMM (Source: European Commission)

The desired interoperability level for a public service is at minimum level 4: 'Sustainable'. At this level, the public service is considered to have implemented all relevant best practices.

The STORM CLOUDS approach

The STORM CLOUDS Smart City Applications were evaluated using the Interoperability Maturity Model Questionnaire. Based on the assessment a tailor-made set of recommendations was provided towards the service owner. The following five principles are applied to generate recommendations:

- Principle 1: Each interoperability attribute differentiates between at least two maturity levels;
- Principle 2: The improvement tables provide recommendations how to improve maturity step-by-step for a specific interoperability attribute;
- Principle 3: When a public service does not have the maximum level yet for a specific interoperability attribute, a recommendation is given to make the step towards the next interoperability level;
- Principle 4: When a public service does have the maximum level for an interoperability attribute, no recommendation is given;
- Principle 5: When the foreseen maturity improvement is a sliding scale (e.g. from less to more), a generic recommendation (not maturity level specific) is given to improve the maturity further along the sliding scale.

For each improvement step the provided recommendation tables show the next maturity level to be achieved through improvement and the general recommendation as to how to achieve the next maturity level [50].

3.12 Protect Users' Privacy

Privacy is understood as the right of a person to have his personal data properly secured. Moreover, it is related with the ability of a person to control, edit, manage and delete information about them and to decide how and to what extent such information is communicated to others [51]. [Any data that could uniquely identify a person or, which is not supposed to be known to any person other than its owner and/or her immediate family, without her consent is called Private Data [52].

Cloud services make it easier for Public Authorities to take advantage of opportunities to share information. For example, sharing personal information with another public Authority or Agency may be achieved by simply creating user accounts with the appropriate permissions within a SaaS solution rather than having to implement a system-to-system interface to exchange information. Although cloud services have the potential to lower the technical barriers to information sharing Public Authorities must ensure that they appropriately manage access to personal information and comply with the requirements of the European and National Privacy Legislation.

The main threats to privacy in a cloud computing environment are:

- Lack of User Control
- Lack of Training and Expertise
- Unauthorized Secondary Usage and Loss of Trust
- Complexity of Regulatory Compliance
- Transborder Data Flow
- Litigation
- Legal Uncertainty

In 2014, the International Organization for Standardization (ISO) adopted ISO/IEC 27018:2014, an addendum to ISO/IEC 27001, the first international code of practice for cloud privacy. Based on EU data-protection laws, it gives specific guidance to cloud service providers (CSPs) acting as processors of personally identifiable information (PII) on assessing risks and implementing state-of-the-art controls for protecting PII [53].

The new standard sets out best practices for public cloud service providers. It establishes security guidelines to protect personal data and provides a privacy compliance framework that addresses the fundamental obligations of a data processor under EU data protection laws. Any organisation that processes PII through a cloud computing service under a contractual arrangement can be certified under ISO 27018 – this means all types and sizes of organisations, including public and private companies, government entities and not-for-profit organisations, are eligible. To qualify for certification under ISO 27018, the applicant provider must agree to be audited by an accredited certification body and must also submit to periodic third party reviews.

Public Authorities can use this standard as an independent measure when evaluating and comparing privacy controls of potential public cloud service providers. An essential step is the signature of the service level agreement with the cloud provider. The agreement defines, among other things, a privacy policy prescribing where and how the organization's data is stored, processed and used (i.e. accepted and prohibited uses) by the cloud service provider. It should also define some privacy related measures and technical controls to be applied on the cloud side, such as the vetting of employees, breach notification, isolation of tenant applications, and the use of products certified to meet national or international standards.

Although the agreement covers a lot of privacy issues, the lack of physical control by cloud users over data storage, and the absence of standardised and mature techniques for monitoring how data is accessed, processed and used inside the cloud, it is harder to verify a cloud's compliance with such privacy policies.

In addition to the evaluation of cloud provider, Public Authorities should also assess their Smart City services to identify issues that may lead to infringing users' privacy. This applies mainly to applications that keep personal information or handle payments. In the first case the application must comply local laws about storing personal data, including any rules about the location of data centres, such as the EU Directive on data Protection [54] while in the second with any rules about safe payments, such as the Payment Card Industry's Data Security Standard (PCI DSS) [55].

However, there are many Smart City infrastructure management applications, such as applications related to public transport, street lighting or road traffic management that do not fall into any of the above categories, and for these data privacy is not such an issue.

Agencies planning to place personal information on a cloud service should perform a Privacy Impact Assessment (PIA) to verify that privacy requirements are adequately addressed.

The STORM CLOUDS approach

The STORM CLOUDS Smart City services have been evaluated regarding privacy issues. The involved Public Authorities in collaboration with the applications' developers perform a Privacy Impact Assessment (PIA) to ensure that they identify any privacy risks associated with the use of the services together with the controls required to manage them effectively.

The privacy impact assessment questionnaire, which was used for each application, contained the following 14 questions [56]:

Q1. Will the project involve the collection of new information about individuals?

Q2. Will the project change the way personal data, particular important to individuals, is being handled? Examples include racial and ethnic origin, political opinions, religious beliefs, trade union membership, health conditions, sexual life, offences and court proceedings.

Q3. Further important examples apply in particular circumstances. The addresses and

phone-numbers of a small proportion of the population need to be suppressed, at least at particular times in their lives, because such 'persons at risk' may suffer physical harm if they are found.

Q4. Will the project compel individuals to provide information about them?

Q5. Will the project perform any data processing at personal data on a large number of individuals? Examples include applications seeking to locate people, or to build or enhance profiles of them.

Q6. Will information about individuals be disclosed to organisations or people who have not previously had routine access to the information, or other third parties that are not subject to comparable privacy regulation?

Q7. Are you using information about individuals for a purpose it is not currently used for, or in a way it is not currently used?

Q8. Will the project involve new or significantly changed consolidation, inter-linking, cross-referencing or matching of personal data from multiple sources?

Q9. Will the project significantly contribute to public safety? Application dealing with critical infrastructure and the physical safety of the population, usually have a substantial impact on privacy.

Q10. Does the project involve using new technology which might be perceived as being privacy intrusive? For example, does the project use biometrics, facial recognition, radio frequency identification (RFID) tags, locator technologies (including mobile phone location, applications of global positioning systems (GPS) and intelligent transportation systems), profiling, data mining, and logging of electronic traffic?

Q11. Does the project involve new identifiers, re-use of existing identifiers, or intrusive identification, identity authentication or identity management processes? For example, does the project use digital signatures, presentation of identity documents as part of the registration scheme, or intrusive identifier such as biometrics?

Q12. Will the project result in you making decisions or taking action against individuals in ways which can have a significant impact on them?

Q13. Is the information about individuals of a kind particularly likely to raise privacy concerns or expectations? For example, health records, criminal records or other information that people would consider to be particularly private.

Q14. Will the project require you to contact individuals in ways which they may find intrusive?

Table 6 – STORM CLOUDS privacy impact assessment questionnaire

In order to drive consistent privacy practices during the development of new Smart City Applications, the Public Authorities should define a privacy framework, which will define standard privacy features and practices. Because security is critical to privacy, the alignment of complementary privacy and security processes helps minimise vulnerabilities in software code, guard against data breaches, and helps to ensure that developers factor privacy considerations into Smart City Services.

3.13 Select the Right Cloud Provider

The choice of a Cloud Service Provider (CSP) requires the evaluation of an extensive list of options. The principal elements to consider for almost every organisation are [57], [58], [59]:

- Service Levels: This characteristic is essential as the Public Authorities in most cases have strict needs regarding availability, response time, capacity and support. Cloud Service Level Agreements (CSLA) are an essential element to choose the right provider and establish a clear contractual relationship between a cloud service customer and a cloud service provider of a cloud service.
- Support: The support is a parameter to consider carefully. It could be offered online or through a call centre, and in some cases, it could be necessary to refer to a dedicated resource with precise timing constraints.
- Security: As already mentioned security is paramount. Although normally, the potential supplier should follow recognised security policies in line with industry best practice, Public Authorities have to formulate a number of relevant questions (i.e. what is the security level offered by the providers? which mechanisms are in place to preserve client's applications and data? Etc.) to evaluate this essential feature for the overall architecture.
- Privacy: Particular attention has to be reserved to legal requirements for the protection of the personal data hosted in the cloud service. Public Authorities should understand the data privacy and retention policies too, as well as where the CSP's data will be located, including any transborder data transfer, if applicable.
- Open Standards: In order to avoid getting locked-in to cloud infrastructure that has restrictive contracts or proprietorial technologies (technologies that are unique to the particular supplier), Public Authorities should prefer solutions that are implemented with fully open source technologies and open cloud standards. These technologies have an elegant escape hatch built into them by their design. Public Authorities can take the entire stack and host it on another CSP or in their premises without losing productivity or data. This backup plan protects them against legislative changes, company restructuring, and much more.
- Compatibility: The requirement of the cloudified applications have to fit into the CSP's existing pre-configured templates and may increase the cost of configuration. Moreover, the CSP's architecture should meet scalability, availability, capacity and performance guarantees and should be sufficient for agency requirements.
- Pricing: Although most cloud providers use the aforementioned "Pay per Use" model, each CSP has a different price system. Understanding how you pay for each service is essential for a meaningful comparison. Moreover, additional costs can still arise, for example through the use of extra features. Terms of the contract, payment methods and payment dates can be deciding factors as well.

- Redundancy: The provision of duplicate or backup equipment that takes over the function of equipment that fails should be discussed at an early stage. The redundancy process and timeframe have to meet the agency's requirements and especially its obligations to the citizens Thus, adequate backup procedures and robust disaster recovery plans must be incorporated into the cloud offering.
- Easy to use administration environment. Make sure your potential provider has a user-friendly client portal. It should allow you to conduct admin tasks or add storage space or services quickly. Ask for a demonstration before you choose one CSP over another.

Most of the considerations mentioned above have already be analysed in separate sections of this document. Given this, we will put more emphasis into the evaluation of two essential elements: Cloud Service Level Agreements and pricing.

Regarding the CSLAs, a prescriptive series of steps should be taken by Public Authorities to evaluate them when comparing multiple cloud providers. [60]:

1. Understand roles and responsibilities
2. Evaluate business level policies
3. Understand service and deployment model differences
4. Identify critical performance objectives
5. Evaluate security and privacy requirements
6. Identify service management requirements
7. Prepare for service failure management
8. Understand the disaster recovery plan
9. Develop an effective governance process
10. Understand the exit process

Regarding pricing, Public Authorities should validate the cost model against the CSP's pricing considering the following [61]:

- Assure pricing is transparent, e.g. subscription or pay-as-you-go pricing, upgrades, maintenance and exit costs
- Costs for unexpected peaks in demand
- Require service price for upgrade and maintenance fees appropriate to the services being procured, some upgrades may be automatic and included in the service
- Confirm the cost model is suitable and allows for scaling and changes to service
- Look for commitment requirements, such as minimum use
- Confirm setup, training and integration fees

- Request references to clarify ongoing cost of service

The STORM CLOUDS approach

STORM CLOUDS Platform is based on OpenStack. There are OpenStack-powered public clouds all over the world. The OpenStack Foundation maintains a Marketplace to help Public Authorities make an informed decision. The most essential OpenStack details of each provider, like which components are included, the versions used, and the APIs exposed, are presented. The Foundation also implemented interoperability testing to validate OpenStack-Powered products, and the results are now available in the Marketplace for public clouds, hosted private clouds, distributions & appliances.

The following map presents the locations of the OpenStack Cloud Providers across Europe. In most cases in each location there are more than one Cloud Providers.

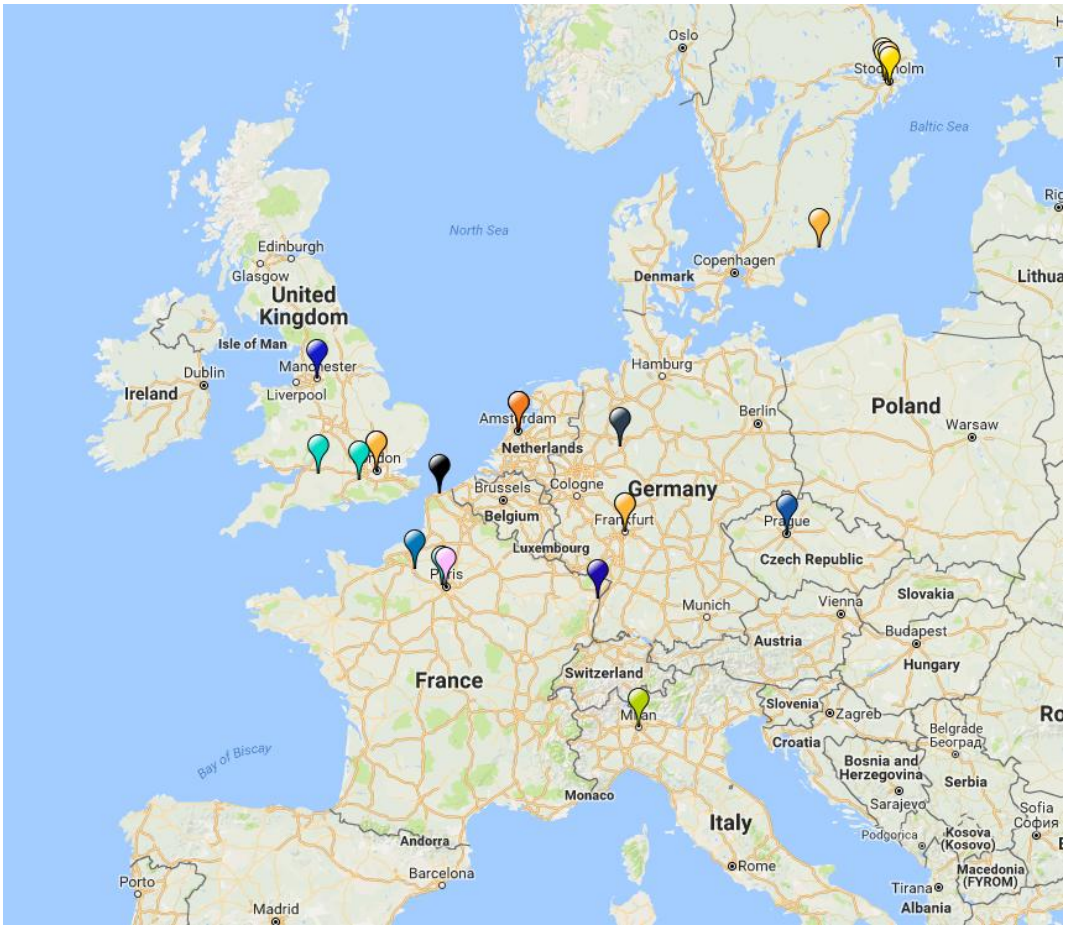

Figure 10 – Locations of the OpenStack Cloud Providers in Europe

The OpenStack Marketplace is available at: https://www.openstack.org/marketplace/

Hewlett Packard Enterprise (HPE), the STORM CLOUDS partner responsible for the SCP, select Enter SLR (http://www.entercloudsuite.com/) as the project's Cloud Service Provider. The company offers Public Cloud IaaS built using the OpenStack. Enter's infrastructure is multi-region: you can deploy resources in Italy (Milan), Germany (Frankfurt) and Netherlands (Amsterdam). The three

locations are interconnected with a proprietary optical ring, which guarantees the lowest possible latency.

The company has an easy to understand pricing policy (https://goo.gl/ddfwRt) and offers a monthly calculator (https://goo.gl/lE8fsi) that allows Public Authorities to evaluate different hosting options. Using this calculator, the STORM CLOUDS partners have estimated the cost of operating the SCP alternative architectures.

HPE is also leading the Cloud28+ initiative*, which started as the project was underway. Cloud28+ (http://www.cloud28plus.eu/) is an open community of Cloud Service Providers, Cloud Resellers, ISVs, Systems Integrators and government entities dedicated to accelerating enterprise cloud adoption across Europe, the Middle East and Africa. Cloud28+ maintain a catalogue of trusted, business cloud services that matches in-country or cross-border buyer and regulatory workload requirements. The initiative offers the following benefits to the Public Authorities: [62]

- Find the right cloud service for your needs based on location of datacentres, price, SLA, certification level, or other workload criteria
- Enable your business to transform to fast, agile Hybrid IT
- Access the largest cloud services community and software developer network in the EU
- Learn about best practices and implementation success stories
- Maintain data sovereignty and feel secure with trusted certification
- Avoid proprietary technology lock-in, thanks to an open source service provider community

European Public Authorities can use the Cloud28+' Catalogue [63] to find public and private cloud services providers, software products, and system integrators across Europe.

* The Cloud28+ initiative is not affiliated with the STORM CLOUDS project.

3.14 Re-Architecting Applications for the Cloud

The applications that have been selected to be migrated to the Cloud may require significant changes to take full advantage of Cloud' characteristics such as high availability and scalability. In particular, as application instances are ephemeral and can be started, stopped or fail at any time, they must be stateless and share nothing. All persistent data must go to external services (e.g. databases, file storage, message queues, caches). Applications should be re-architected in order to take full advantage of the Cloud's features (Figure 11) [64].

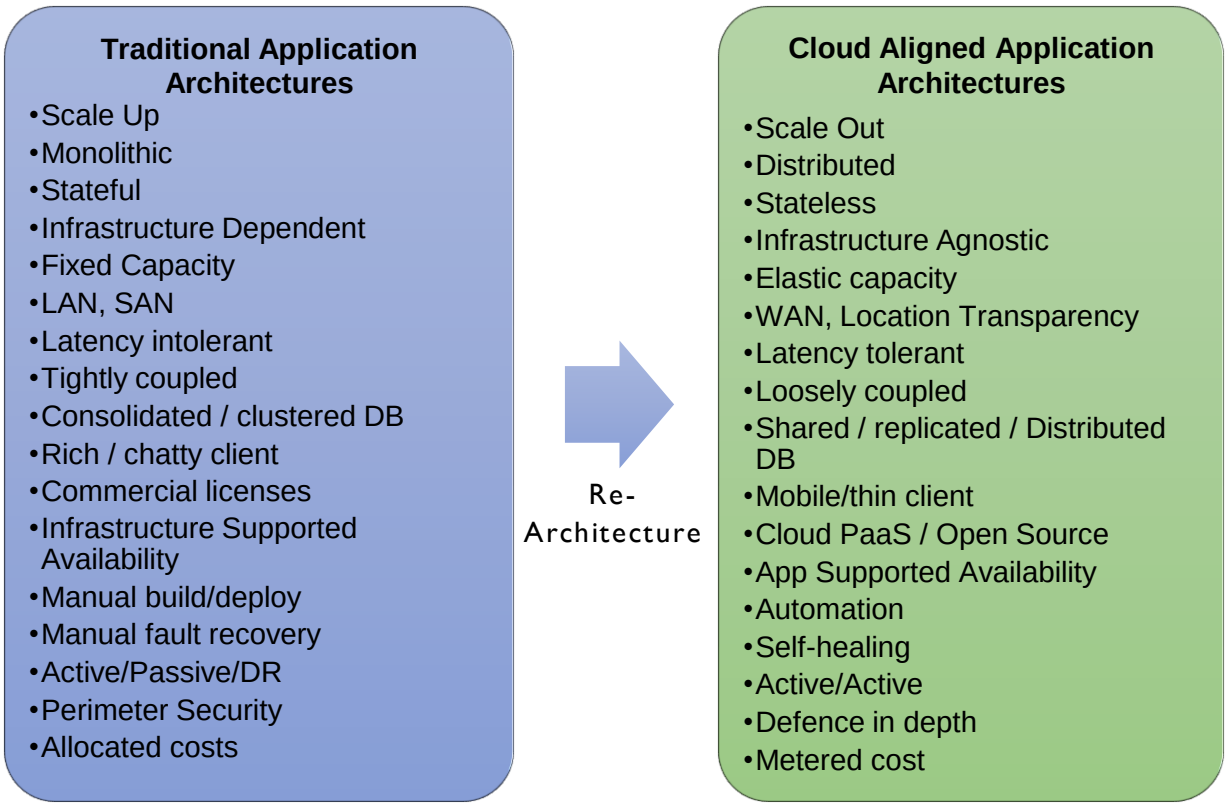

Figure 11 – Traditional vs Cloud Aligned Application Architectures (Source: New Relic)

Moving an application to the cloud will require some work under the hood. However, choosing to re-architect an application is ideal when: [65]

- Hardware cost is substantial. Re-architecting an application for the cloud means access to world-class hardware without needing a world-class budget. Companies are able to pay as they go and avoid the investment required for more hardware.
- IT staffing levels are low. Moving to the cloud automates a lot of the server and application management, as well as maintenance tasks that would otherwise be performed by in-house IT staff.
- Geolocation is a requirement. The cost to do geolocation on the cloud is miniscule since many data centres are located in central regions.

- The application needs to scale for predicted, but infrequent, uptime. The cloud allows systems that have occasional spikes, such as an e-commerce application that sees a lot of activity on Black Friday, to quickly and easily scale servers on demand without an expensive hardware investment or footprint.

Determining the right migration strategy for an application depends on its level of cloud alignment, cloud readiness, potential benefits achieved from migrating, and risks. There is several application migration common methods and approaches, which should consider the Public Authorities for their existing applications (Figure 12) [64]:

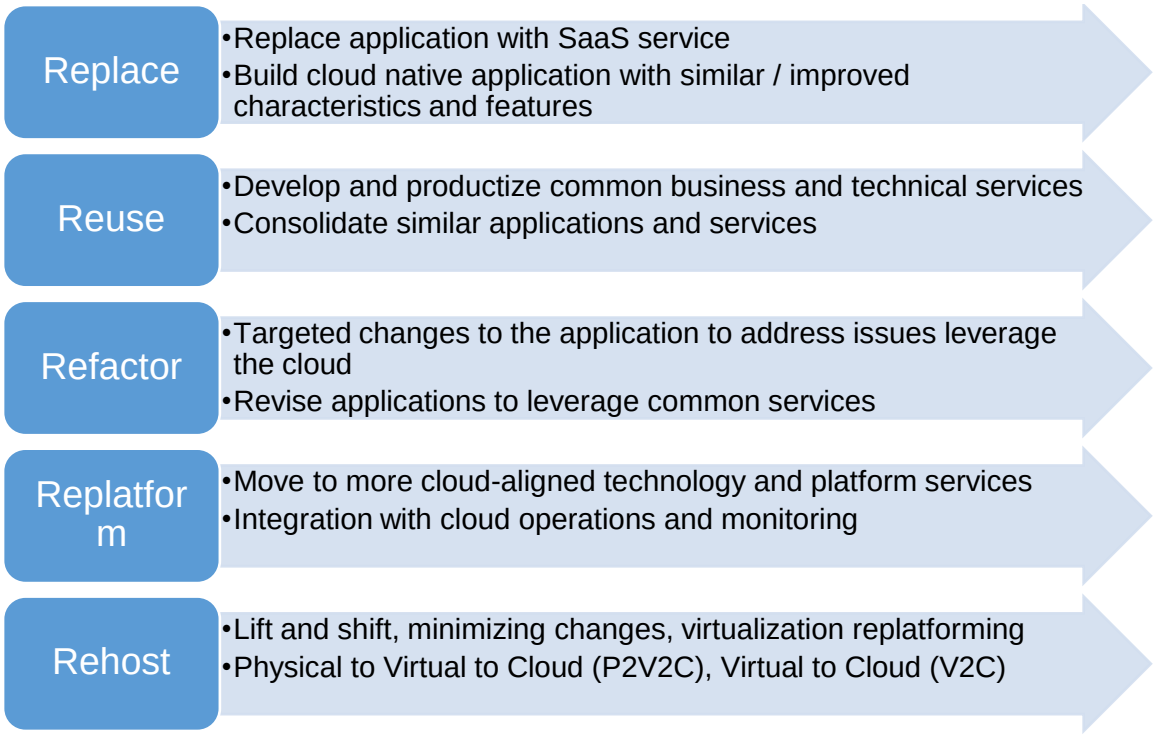

Figure 12 – Application Migration Common Methods and Approaches ((Source: New Relic)

Depending on the changes the candidate applications reach different Cloud maturity levels. The characteristics of each level are presented in the following table: [64]

Maturity Level	Characteristics
Cloud Washed	- Force fit to run in cloud environment - Resources not optimize – no horizontal scaling - Minimal modification done to be cloud compliant
Cloud Adopted	- Resources not optimize – no automatic elasticity-instance manually started - Some modification done to be cloud compliant
Cloud Optimized	- Resources being optimized – horizontal scaling possible

	- Elastic on instance level – cloud management layer determines when to start/stop additional instances - Major modification done to be cloud compliant
Cloud Native	- Fully cloud aware - can communicate with the cloud management layer to start-up or shutdown instances of itself - Designed for failure and self-healing - Elastic and resource efficient

Table 7 – Cloud Application Maturity (Source: New Relic)

The STORM CLOUDS approach

The STORM CLOUDS platform supports two different architectures; a "scale-up" architecture for applications with traditional architectures (Figure 13) and a "scale-out" architecture for applications with Cloud-aligned architectures (Figure 14).

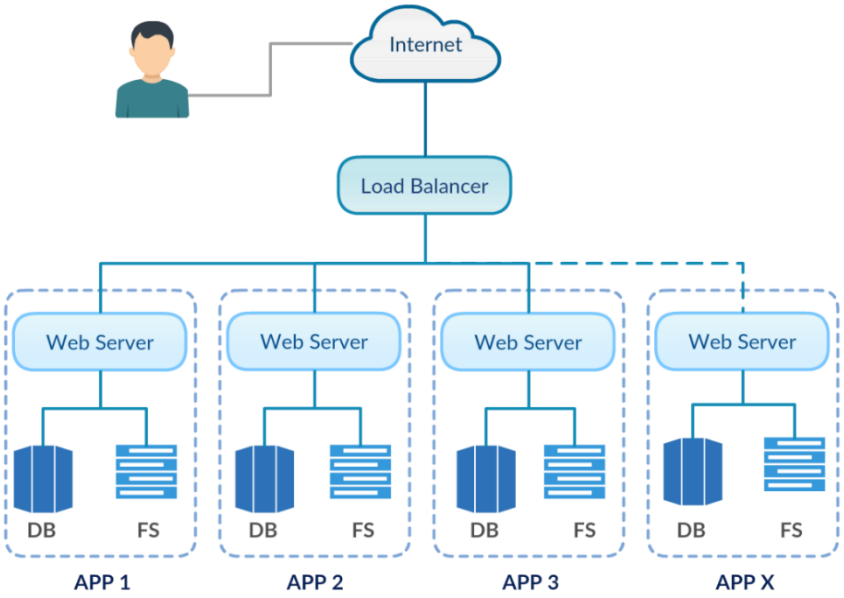

Figure 13 – SCP "Scale-up" Architecture for traditional applications

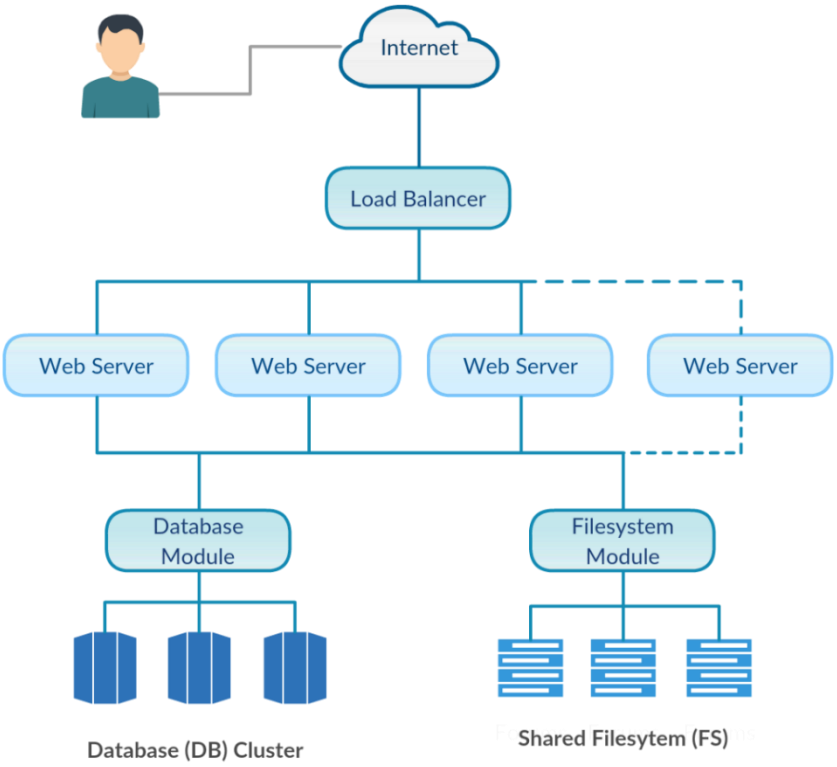

Figure 14 – SCP "Scale-out" Architecture for Cloud-ready applications

The "scale-up" architecture is the one where the application can benefit from more resources, such as CPU and memory, being added to a single server or node. In contrast, the "scale-out" architecture is the one where an application scales by additional nodes being made available for the workload; that is, it scales horizontally.

Scale-out applications can take advantage of the pay-by-usage cost model of the cloud. When there are increased requests for an application, more nodes can be deployed to handle the increased load. When the requests slow down, the additional nodes can be powered off to reduce costs.

Although virtualisation technologies have started to support the dynamic scale-up and scale-down of a single VM's resources (memory, CPU and disc space), this action requires, at the moment, the use of custom scripts. On the contrary, the systems that support the scale-out architectures provide native support to the dynamic increase of resources.

3.15 Explore the Containerization Technologies

Application containerization is an operating system level (OS-level) virtualization method for deploying and running distributed applications without launching an entire virtual machine (VM) for each app. Instead, multiple isolated systems are run on a single control host and access a single kernel. The application containers hold the components such as files, environment variables and libraries necessary to run the desired software. Because resources are shared in this way, application containers can be created that place less strain on the overall resources available [66].

The following figure (Figure 15) compares application deployment using a hypervisor and a container. As the figure shows, the hypervisor-based deployment is ideal when applications on the same cloud require different operating systems or OS versions. The abstraction must be at the VM level to provide this capability of running different OS versions. With containers, applications share an OS (and, where appropriate, binaries and libraries), and as a result, these deployments will be significantly smaller in size than hypervisor deployments, making it possible to store hundreds of containers on a physical host (versus a strictly limited number of VMs). Because containers use the host OS, restarting a container doesn't mean restarting or rebooting the OS [67].

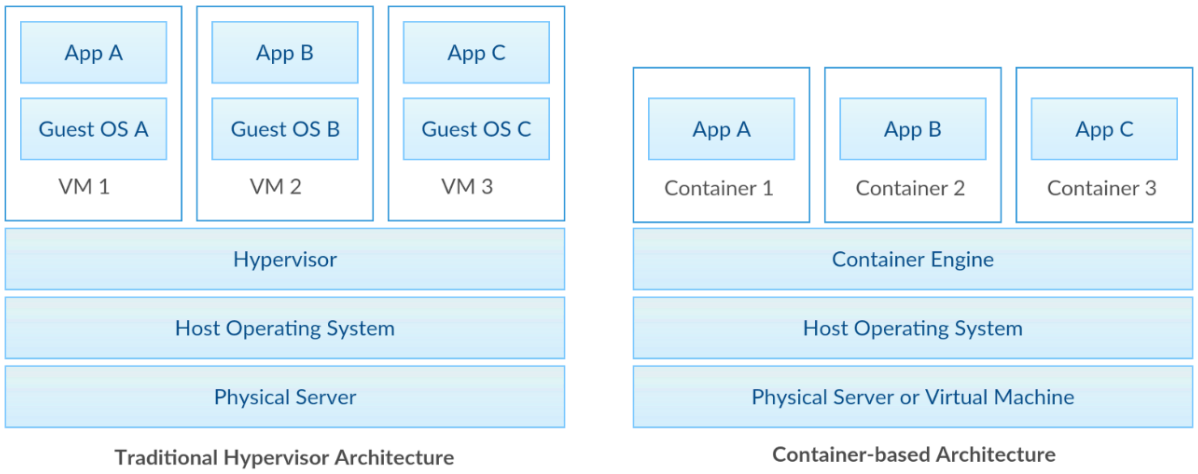

Figure 15 - Comparison of (a) hypervisor and (b) container-based deployments.

Containerization, provides additional benefits to existing, traditional, virtualisation technologies which open up new use cases that threaten to revolutionize the IT industry. In particular they provide: [68], [69]

- A lightweight system; containers start instantly and use less RAM. Images are constructed from layered filesystems and share common files, making disk usage and image downloads much more efficient.
- An improved process for Continuous Integration and Delivery; developers are able to develop, test, and deploy applications to a large number of servers.

- An easy to migrate solution; each container is an isolated instance that doesn't hold a guest operating system, it's very easy to migrate to from one deployment to another. The container stack allows for great portability.
- A secure environment; containers isolate applications from one another and the underlying infrastructure, while providing an added layer of protection for the application.

In order to overcome the biggest obstacle to the adoption of the container-based deployments, the Open Container Initiative (OCI) formed under the auspices of the Linux Foundation in 2015. OCI is a lightweight, open governance structure (project) for the express purpose of creating open industry standards around container formats and runtime [70]. The OCI currently contains two specifications: The Runtime Specification (runtime-spec) and the Image Specification (image-spec).

Many of the organisations that work on Cloud technologies are part of the Open Container Initiative, so the Public Authorities can be confident that we will increase cross-container support and flexibility going forward.

The STORM CLOUDS approach

The STORM CLOUDS project will follow the containerization approach in order to create a faster, more resource-effective, and more secure setup for the SCP components and services. The Docker Engine will be used in order to create containers to deliver some of the Smart City Services. A developer will be able to use Docker locally, or on a different cloud infrastructure, and then deploy into the selected production environment with no major issues.

Docker is an open-source engine which automates the deployment of applications as highly portable, self-sufficient containers which are independent of hardware, language, framework, packaging system and hosting provider [71]. With Docker, you can separate your applications from your infrastructure and treat your infrastructure like a managed application.

Docker is the most popular and advanced platform today, which provides access to a set of high-quality core tools and additionally to a community and a whole ecosystem of third-party products and services that help developers from the very first steps of developing their application through every incremental deployment of it.

At its core, Docker provides a way to run almost any application securely isolated in a container. The isolation and security allow you to run many containers simultaneously on your host. The lightweight nature of containers, which run without the extra load of a hypervisor, means you can get more out of your hardware.

For both developers and operators, Docker offers the following high-level benefits, among others:

- Deployment Speed/Agility – Docker containers house the minimal requirements for running the application, enabling quick and lightweight deployment.
- Portability – Because containers are essentially independent self-sufficient application bundles, they can be run across machines without compatibility issues.

- Reuse – Containers can be versioned, archived, shared, and used for rolling back previous versions of an application. Platform configurations can essentially be managed as a code.

Although Docker is a way of managing multiple containers on a single machine, the capability to be used behind Nova (OpenStack's Hypervisor Engine) makes it much more powerful since it's then possible to manage several hosts, which in turn manage hundreds of containers. The current Docker project aims for full OpenStack compatibility. [72]

As containers don't aim to be a replacement for VMs but they are complementary in the sense that they are better for specific use cases the Public Authorities should evaluate both solutions to find which fits better to their needs and requirements.

4. Conclusions

This report provides guidance to Public Authorities (both decision makers and IT staff) for considering and executing the migration of their applications to a Cloud Computing environment. Fifteen general guidelines for migrating applications to the cloud have been enriched with hands-on experience gained from a practical exercise, namely the deployment of Smart City applications to the STORM CLOUDS platform in four European Cities (i.e. Agueda, Miscolc, Thessaloniki and Valladolid).

By focusing on concepts and best practices – like adopting an open innovation methodology, setting up a monitoring and validation process, prioritizing the applications that should move to the Cloud, choosing the right Cloud service category and Cloud deployment model, embracing the power of open technologies, planning carefully and automating the migration, using the right tools to manage and monitor the Cloud environment, focusing on security, protecting your data, ensure interoperability, protect users' privacy, select the right Cloud provider, re-architecting you applications for the Cloud and explore the containerization technologies – Public Authorities can understand how to successfully deploy Smart City application on the Cloud. Moreover, they became familiar with the SCP set of tools and features, which facilitate the migration process.

References

1 Komninos N, 2013, Smart Cities and the Future Internet: Innovation ecosystems of embedded spatial intelligence in *Proceedings of International Conference for Entrepreneurship, Innovation and Regional Development*, ICEIRD, 2013

2 Mitchel S, *et al.* (2013) *The Internet of Everything for Cities: Connecting People, Process, Data, and Things to Improve the 'Livability' of Cities and Communities*, Point of View, Cisco

3 Fu, Y, Jia, S and Hao, J. (2015) A Scalable Cloud for the Internet of Things in Smart Cities in *Journal of Computers*, 26(3), pp.

4 Schaffers H, *et al.* 2011, Smart Cities and the Future Internet: Towards Cooperation Frameworks for Open Innovation' In *J. Domingue, et al. (Eds.) The Future Internet: Future Internet Assembly 2011: Achievements and Technological Promises*, pp. 431-446

5 Marstona S, *et al.* 2011, Cloud computing — The business perspective, *Decision Support Systems,* Volume 51, Issue 1, Pages 176–189.

6 Buyya R, *et al.* 2009, Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility, *Future Generation Computer Systems,* 25, 599-616.

7 U.S. Department of Commerce - National Institute of Standards and Technology, 2011, *The NIST Definition of Cloud Computing*, viewed November 5, 2015, <http://goo.gl/6MgjdS>

8 *Gartner IT Glossary: Cloud Computing,* viewed November 5, 2015 <http://goo.gl/SQHw7O>

9 Jinesh Varia, 2011, *Architecting for the Cloud: Best Practices*, Amazon Web Services, viewed November 6, 2015, https://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

10 *Best Practices for Federal Agency Adoption of Commercial Cloud Solutions*, 2015, Professional Services Council, viewed November 6, 2015 <https://goo.gl/j5NbfF>

11 U.S. Chief Information Officer, 2011, *Federal Cloud Computing Strategy,* White House

12 *Cloud Computing: Transforming the Government of Canada for the Digital Economy,* 2015, Information Technology Association of Canada, viewed November 6, 2015 < http://goo.gl/G7Qmu2>

13 Government adopts 'Cloud First' policy for public sector IT, 2013, viewed November 9, 2015 <https://goo.gl/ARRaIT>

14 European Commission, 2012, *Unleashing the Potential of Cloud Computing in Europe*, viewed November 9, 2015 < http://goo.gl/VXfLic>

15 *STORM CLOUDS Project – Main Page,* viewed November 7, 2015 <http://stormclouds.eu/>

16 Battarra M, et al., Storm Clouds Platform: A Cloud Computing Platform for Smart City, Applications in *Smart Cities and Cloud Computing Special Issue*, Journal of Smart Cities (submitted)

17 *OpenStack – Main Page,* viewed November 7, 2015 <http://www.openstack.org/>

18 *Cloud Foundry – Main Page,* viewed November 9, 2015, <https://www.cloudfoundry.org>

19 H. Chesbrough, W. Vanhaverbeke, and J. West, Open innovation: Researching a new paradigm. OUP Oxford, 2006.

20 *Migrating Applications to Public Cloud Services: Roadmap for Success*, 2013, Cloud Standards Customer Council, viewed November 5, 2015 <http://goo.gl/EGHN8l>

21 *Transforming Government – Cloud policy framework for innovation, security, and resilience*, 2015, Microsoft, viewed November 9, 2015 <http://goo.gl/TJz1Jh>

22 *The twelve-factor app – A methodology for building software-as-a-service apps*, viewed November 9, 2015 <http://12factor.net/>

23 *Finding Your Right Cloud Solution: Private & Public Clouds,* 2015, Oracle, viewed November 9, 2015 <http://goo.gl/KRolV6>

24 *The pros and cons of public, private and hybrid clouds*, 2015, viewed November 9, 2015 <http://goo.gl/D6jBYX>

25 *Open cloud: Just a buzzword or the future of infrastructure?*, viewed November 9, 2015 <http://goo.gl/1QRJCY>

26 Open Cloud Alliance: Openness as an Imperative, 2014, Crisp Research AG, viewed November 9, 2015 <http://goo.gl/db8fZr>

27 *The Top Open Source Cloud Projects of 2014*, viewed November 9, 2015 <https://goo.gl/HQutA2>

28 *Gluster – Main Page*, viewed November 2, 2015 <https://www.gluster.org>

29 *HAProxy – Main Page*, viewed November 10, 2015 <http://www.haproxy.org>

30 *Zabbix – Main Page*, viewed November 10, 2015 <http://www.zabbix.com>

31 *phpMyAdmin – Main Page*, viewed November 10, 2015 <https://www.phpmyadmin.net>

32 *phpPgAdmin – Main Page*, viewed November 10, 2015 < http://phppgadmin.sourceforge.net/doku.php >

33 *Duplicity – Main Page*, viewed November 10, 2015 <http://duplicity.nongnu.org/>

34 Laszewski T, Nauduri P, 2012, *Migrating to the cloud: Oracle client/server modernization*, Syngress, Elsevier Inc.

35 *What does the Commission mean by secure Cloud computing services in Europe?,* 2013, European Commission, viewed November 10, 2015 <http://goo.gl/M0Rqia>

36 Giannakoulias A, Cloud Computing Security: Protecting Cloud-Based Smart City Applications, Applications in *Smart Cities and Cloud Computing Special Issue*, Journal of Smart Cities (submitted)

37 Cloud Security Guidance: Summary of Cloud Security Principles, viewed June 24, 2016 <http://goo.gl/mUf5c2>

38 *OWASP Zed Attack Proxy Project*, viewed November 10, 2015 <https://goo.gl/A46kLn>

39 *OpenVAS*, viewed November 10, 2015 <http://www.openvas.org>

40 *SQL Inject Me Firefox Add-on*, viewed November 10, 2015 <https://goo.gl/2VYKoM>

41 *SSL Server Test*, viewed November 10, 2015 <https://goo.gl/J9MuUl>

42 *Vega*, viewed November 10, 2015 <https://subgraph.com/vega/>

43 Microsoft, 2014, *Protecting Data and Privacy in the Cloud*

44 Data protection, viewed November 10, 2015 <http://goo.gl/MipM9c>

45 ITU-T, 2002, *Global Information Infrastructure terminology: Terms and definitions*

46 ISO, *ISO/IEC 19941 standard: "Information Technology – Cloud Computing – Interoperability and Portability"*

47 European Commission, *Interoperability Solutions for European Public Administrations Programme,* viewed June 2, 2016 <http://ec.europa.eu/isa/>

48 European Commission, *Interoperability Maturity Model Documentation (Guideline & Definitions)*

49 European Commission, 2013, *Interoperability Maturity Model*, viewed June 23, 2016 <http://goo.gl/ydFPHs>

50 European Commission, 2016, *Interoperability Maturity Model Full - Recommendations*

51 Conducting privacy impact assessments code of practice, 2014, UK Information Commissioner's Office

52 Privacy Issues and Measurement in Cloud Computing: A Review. International Journal of Advanced Research in Computer Science, Volume 4, No. 4, March-April 2013.

53 International Organization for Standardization, 2014, *ISO 27018: Cloud Computing Privacy Standard*, viewed June 23, 2016 <http://goo.gl/4dBG2h>

54 European Commission, *Protection of personal data,* viewed June 2, 2016 <https://goo.gl/xwemTY>

55 PCI Security Standards Council, viewed June 2, 2016, <https://www.pcisecuritystandards.org>

56 STORM CLOUDS Project, 2015, *Deliverable 4.3: Privacy and security measures*

57 IDG Enterprise, 2014, *Best practices for moving workloads to the cloud,* viewed June 5, 2016, <https://goo.gl/D21nMS>

58 OpenSource.com, 2013, *Do cloud right: Four critical steps to selecting the provider for you*, viewed June 5, 2016, <https://goo.gl/bwVJdY>

59 IT Lab, 2013, *Cloud Migration Guide*, viewed June 5, 2016, <https://goo.gl/u8YRjW >

60 Cloud Standards Customer Council, 2015, *Practical Guide to Cloud Service Agreements*, Version 2.0, viewed June 5, 2016, <https://goo.gl/82BhCE>

61 Australian Government, 2012, *A Guide to Implementing Cloud Services*

62 Cloud28+ Europe's Cloud of Clouds, viewed June 5, 2016, <http://www.cloud28plus.eu/>

63 Cloud28+ Catalogue, viewed June 5, 2016, <https://member.cloud28plus.eu/catalogue>

64 New Relic, 2015, Cloud Migration Cookbook: A Guide to Moving Your Apps to the Cloud, viewed June 15, 2016, < https://goo.gl/DYAqX5>

65 Headspring, 2014, *Migrating to the Cloud: Re-Platforming Legacy Enterprise Applications*, Best Practices Guide, viewed June 15, 2016< https://goo.gl/shTHBi >

66 Definition: Application containerization, viewed June 23, 2016 <http://goo.gl/hOzurx>

67 D. Bernstein, 2014, Containers and Cloud: From LXC to Docker to Kubernetes, IEEE Cloud Computing, vol. 1, no. 3, pp. 81–84

68 Claus Pahl, 2015, *Containerization and the PaaS Cloud*, IEEE Cloud Computing, vol.2, no. 3, pp. 24-31, May-June 2015, doi:10.1109/MCC.2015.51

69 CodeShip, 2016, *Why Containers and Docker are the Future*, White paper

70 Open Container Initiative, viewed June 23, 2016 <https://www.opencontainers.org/>

71 What is Docker, viewed June 23, 2016, <https://www.docker.com/what-docker>

72 Using Docker with OpenStack, viewed June 23, 2016, < https://wiki.openstack.org/wiki/Docker>

