COMPETITIVENESS AND INNOVATION FRAMEWORK PROGRAMME

CIP-ICT-PSP-2013-7

SERVICE DISTRIBUTION NETWORK AND TOOLS FOR INTEROPERABLE PROGRAMMABLE, AND UNIFIED PUBLIC CLOUD SERVICES

Deliverable D7.4

Best Practices and Policy Development Guidelines for Public Cloud Services

Workpackage	WP7 – Pilot Services Evaluation and Best Practices Elicitation
Editor(s):	G.Ledakis, I.Livenson, G.Ducatel, J.Daniel, A.Juan, J.C.Pérez, A. Toomsalu, R.Contri
Responsible Partner:	SingularLogic Information Systems & Applications SA
Quality Reviewers:	G.Ducatel (BT), I.Livenson (NICPB)
Status-Version:	Final – v1.0
Date:	30/11/2016
EC Distribution:	Public
Abstract:	This deliverable will document the best practices and policy development guidelines that will stem from the pilot operations and their evaluation. Emphasis will be put on the adoption of public cloud sector by public organizations (i.e. public bodies).

This project has received funding from the European Union's Competitiveness and Innovation Framework Programme under grant agreement no 621009

Document Revision History

Version	Date	Modifications Introduced	
		Modification Reason	**Modified by**
v0.1	24/03/2016	Draft Table of Contents (ToC)	SILO
v0.2	03/07/2016	First best practice ready for integration	SILO
v0.3	30/08/2016	More content added	SILO
v0.4	03/10/2016	Updates on the best practices topics at the Madrid meeting	All partners
V0.5	26/11/2016	Input on cross border guidelines	ATOS
V0.6	12/12/2016	Installation of IaaS using Mirantis	NICPB
V0.7	12/1/2017	Security and Marketplace related input	BT
V0.8	14/1/2017	More content in section 2	SILO
V0.9	17/1/2017	Integrating all inputs and new content, ready for review	SILO
V1.0	30/1/2017	Final version after quality review	SILO

Contents

List of Figures

List of Tables

Definitions, Acronyms and Abbreviations

Acronym	Title
API	Application programming interface
CBA	Cross Border Authentication
DOW	Description of Work
FP	Federation Proxy
IaaS	Infrastructure as a Service
ICT/IT	Information and Communications Technology / Information Technology (used interchangeably)
IDP	Identity Provider
ISV	Independent Software Vendors
G2C/G2B/G2G	Government to Customers / Government to Business / Government to Government
MAAS	Metal-as-a-Service
MoSG	Municipality of Stari Grad
PaaS	Platform as a Service
SaaS	Software as a Service
SEMIRAMIS	Secure Management of Information across multiple Stakeholders
SP	Service Provider

Table 1: Definitions, Acronyms and Abbreviations

Executive Summary

In this document, we are summarizing the most important findings we collected during project research, development, piloting and evaluation phases. The collection of these findings has been provided as aggregated and integrated knowledge in best practices and guidelines that are presented in this document and are also available on the project website[1]. The best practices and guidelines that have been created are targeted to many stakeholders but their main purpose is to help promoting and boosting the adoption of public cloud services within the public sector, always inline with the concept of STRATEGIC Service Store and the tools that the project has used or integrated.

[1] http://strategic-project.eu/guidelines-and-best-practices/

1 Introduction & scope of the deliverable

1.1 Scope and purpose of the document

This deliverable documents the best practices and guidelines that stem from our experience in working towards the cloud enablement of public bodies, the pilot operations and the evaluation process. Emphasis has been put on the creation of content that will be helpful to mainly to public organizations (i.e. public bodies), but also to other related stakeholders (cloud application developers, cloud providers).

These developed practices and guidelines cover a wide range of topics, issues and processes that were encountered throughout the project duration. The following table is based on the DoW [1] of the project and provides the suggested scopes and topics that should be covered by best practices produced during by the project.

# of scope	Topic/Area for Best Practices and Policy Development Guidelines	Target Group/Stakeholders
SC1	Cloud-Enablement of Public Sector On-line Distributed Services	Public Sector Organizations, Governmental Organizations, Solution providers ISVs
SC2	Guidelines for the development and sustainability of marketplaces of public cloud services	Cloud Services Providers, Policy Makers
SC3	Localization, Adaptation and Governance of Public Cloud Services – Relevant «last mile services»	(National/Regional Level) Cloud Providers, ISVs, Solution Providers, Policy Makers (local/regional/national/EU level)
SC4	Management and Brokerage Services for Public Cloud Services	(National/Regional Level) Cloud Providers, ISVs, Solution Providers, Policy Makers (local/regional/national/EU level)
SC5	Development of Novel Public Cloud Services for the Public Sector	(National/Regional Level) Cloud Providers, ISVs, Solution Providers, Policy Makers (local/regional/national/EU level)

Table 2: Scope of the Best Practices and Policy Development Guidelines to be produced by STRATEGIC

Also, according to the DoW Objectives, STRATEGIC would produce at least ten (10) best practices and policy development guidelines, so the list of topics has been refined and detailed as part of the project's workplan in order to produce best practices and guidelines that would cover all scopes and the same time will be useful to the stakeholders.

In the Table 3 that follows, produced best practices and guidelines are presented. Overall fourteen (14) separate best practices and guidelines have been created, exceeding the initial goal imposed by the DoW.

#	Title		Scopes covered	Chapter in the document	Responsible
1	Cloud Migration for Public Bodies: Benefits and Obstacles		SC1	2.1	SILO
2	Selection of appropriate cloud computing service model		SC1	2.2	SILO
3	Selection of Appropriate Cloud Computing Deployment Model		SC1	2.3	SILO
4	Selection of Cloud Provider and the Benefits of Cloud Brokerage		SC1, SC4	2.4	SILO
5	Preparation and Design of a Private OpenStack Installation		SC1, SC3	2.5	SILO
6	Installation of IaaS Using Automated Scripts	Installation of OpenStack IaaS Using MaaS and Juju	SC1, SC3	2.6.1	SILO
		Mirantis based installation of OpenStack IaaS	SC1, SC3	2.6.2	NICPB
7	Protection of Applications and Data on Cloud Environments Using STRATEGIC		SC1, SC5	3	BT
8	Challenges for public sector cloud services creation		SC1, SC3, SC5	4.1	Pilots / SILO
9	Guidelines for the Creation of Adaptable and Reusable Applications with Strategic Service Store		SC1, SC2, SC3, SC5	4.2	SILO/ BT / NICPB
10	Guidelines for the Creation and Sustainability of Application Marketplaces		SC2, SC4	5.1	BT
11	Connection of Private Clouds to a Public Marketplace		SC1, SC2,	5.2	SILO / BT

#	Title	Scopes covered	Chapter in the document	Responsible
		SC3		
12	Accessibility of a Private OpenStack Installation Through STRATEGIC Service Store	SC1, SC2, SC3	5.3	SILO
13	SEMIRAMIS usage by public bodies for the exchange of cross-border information	SC1, SC3, SC5	6.1	ATOS
14	STORK usage by public bodies for cross-border authentication	SC1, SC3, SC5	7.1	ATOS

Table 3: Best Practices and Guidelines produced by STRATEGIC

However, the most important achievement is that the produced material covers all aspects that a public body moving to the cloud should be aware of, and at the same time illustrate the added value of STRATEGIC, either through the STORK and SEMIRAMIS integrations or through the STRATEGIC Service Store usage.

1.2 Target audiences

The intention of this deliverable is to include all the created best practices and guidelines; therefore, the overall audience of this document is the audience of the best practises. This audience is external to the consortium and includes governmental and public sector organizations (IT people, administration, management), cloud providers, integrators/solution providers/ISVs, citizens, policy makers and actually anyone interested into cloud adoption.

1.3 Structure of the document

The whole document focus on the presentation of best practices and guidelines, therefore most of the document (chapters 2-7) consists of the actual text of the best practices. We have tried to organize and group together best practices with similar topics but each of the sections is a standalone best practice or guideline. Finally, in chapter 8 we provide information on how these documents are actually disseminated through the project website and social media. Chapters 1 and 9 serve as the introduction and conclusions of the document.

2 Public Sector and the Usage of the Cloud

2.1 Cloud Migration for Public Bodies: Benefits and Obstacles

An important decision for all public bodies is to evaluate the pros and cons of utilizing the cloud in their organization and decide on which cases the adoption of cloud is beneficial and in which case applications of a public organization have to be used as cloud native applications.

Starting with the benefits of using cloud we consider the following:

Using the Cloud is suggested when there is lack of the IT resources needed for the maintenance and management of the underlying infrastructure. Setting up and maintaining on-premise solutions can be time-consuming, and require lots of resources, especially IT experts, resulting in substantial technical and management challenges. Cloud resources simplify all of this by removing the on-site equipment, simultaneously making upgrades a less expensive, headache-free and more seamless experience [2].

Using the Cloud can help reducing overhead costs. A crucial aspect for the public bodies that can push towards the usage of cloud is the possibility to remove the substantial upfront expense of installing an on-premise solution and eliminating the need for dedicated hardware. So apart from having no equipment and infrastructural services to maintain, monthly payments based on usage might be more appropriate to specific public bodies.

A proper cloud installation provides reliable data backup. Especially when a public cloud is used, fail safes and backups mechanisms are built into cloud services, thus making the management of the deployed services and data a lot easier. And even if a private cloud has been constructed, built-in redundancies and backup mechanisms should be used

Cloud usage can help providing applications faster, thus fostering innovation and improving the overall state of public organizations. Using the Cloud allows public organizations to quickly and efficiently deploy services and make them available. In comparison to the in traditional deployments, by using cloud (both private cloud with appropriate resources or a public cloud) an organization is able to focus its resources on the development or refining of cloud services, rather than spending time with tactical, or technical issues.

However, it is there are also some facts that have been taken under consideration before moving to the Cloud. These facts apply not only to the organization level but also at application/service level, as a public organization can move only some of the applications on the cloud or make a gradual move to this direction.

If the application/service that is about to be migrated is performing to expectations. If an existing on-premise solution is working properly and truly meeting expectations, there might be no need to migrate the application to a cloud environment. However, it has to been analysed if the solution is optimized for cost aspect as well, as a common issue with on-premise solutions is that the resources are heavily underutilized.

[2] http://www.manageforce.com/blog/to-cloud-or-not-to-cloud-in-2017

In public organization, specific security measures and/or unique regulatory compliance have to be taken under consideration, thus making the usage of public cloud providers difficult or not possible. The ideal solution in this case is to setup and utilize a private cloud. STRATEGIC supports public organizations in this difficulty as its provides regional information of the infrastructure, supports private installations and the consortium has made the exercise of investigating the regulatory compliance of the pilot cases in three European countries (UK, Italy, Serbia).

Finally, it has to be mentioned that if **an investment in *on-premise infrastructure and services* has been done recently,** there is no rush for an organization to move to the Cloud, as it would be more difficult to reap the benefits of this change**.** One solution on this is to create a private cloud using this infrastructure, however even this should be done with proper planning, to avoid spending more resources than normally would be spent**.**

2.2 Selection of Appropriate Cloud Computing Service Model

In recent years, *cloud computing*, a term that broadly describes an emerging group of related technologies and business models, has become standard vocabulary in the private and public sectors who wish to harness the potential benefits of this technology for their organizations and businesses.[3]In particular governments are replacing their legacy IT systems with cloud computing technologies and implementing new cloud-based tools for collaboration and information sharing across agencies and units.[4]

However, an important decision for any organization that wants to join the cloud era is the selection of the appropriate cloud computing model. The basic types of cloud offerings, include infrastructure as a service (**IaaS**), platform as a service (**PaaS**), software as a service (**SaaS**). Each service is built on top of the other.[5]

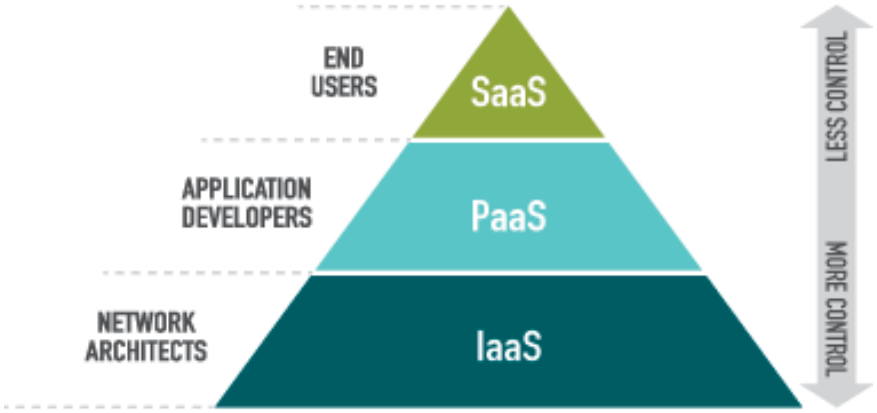

Figure 1: The layers of Cloud Computing
(source : http://www.cloudserver-saas.com/index.php/2016/03/02/what-is-saas-iaas-and-paas-cloud-computing/)

[3] Governments and Cloud Computing: Roles, Approaches, and Policy Considerations, Urs Gasser and David R. O'Brien, Working Paper No. 2013/23, August 2013
[4] http://www.imaginea.com/images/resources/white-papers/white%20paper%201.pdf
[5] https://www.ibm.com/blogs/cloud-computing/2014/02/how-does-cloud-computing-work/

SaaS provides all software application services through a Web browser and not a locally-installed application. SaaS eliminates worries about application servers, storage and application development. Consumer is usually an end user of applications.

In the PaaS model, a hosted environment is provided to the consumer to run his applications. The consumer controls the applications that runs in the environment, but does not control the operating system, hardware or network infrastructure on which they are running. The platform is typically an application framework. Typical customer is an application developer or DevOps teams.

Infrastructure as a Service (IaaS) is the cloud model where consumer uses "fundamental computing resources" such as grids or clusters or virtualized servers, networks, storage and system software designed to augment or replace the functions of an entire data center. Typical customer must be familiar with and prepared to spend time and efforts on infrastructure management.

According to our experience from the pilot execution and the questionnaires of the project that were filled by more than 110 public stakeholders, we identified the deployment model of Software as a Service appears to be the most favourite approach, due to the ease of starting. There is also however a big part of the public sector using IaaS and PaaS or even a combination of the models (IaaS, PaaS and SaaS). This is common especially if there are cloud services that serve specific needs of the public body or there is a need for public agency to host the governmental cloud(G-cloud) provided resources (at the IaaS or PaaS level to other governmental agencies).

Trends in service and deployment models of G-Cloud solutions – Private, SaaS and "G2all" (i.e. G2C, G2B, G2G) seems to be dominating. The national G-Cloud early adopters have taken different paths in choosing for their cloud solutions. Although there are a few countries that have opted for public solutions, it appears that most countries currently working or planning towards a G-cloud lean towards Private cloud (i.e. a governmental agency hosting the G-cloud services with own-resources). Regarding the deployment model Software as a Service again appears to be the most favourite approach, although a combination of other models (IaaS, PaaS and SaaS) is also common (especially if there is a need for the agency hosting the G-cloud to provide resources (at the IaaS or PaaS level to other governmental agencies). Regarding the target audiences of G-Cloud early adopters or under planning the trend appears to be clearly towards servicing all players, i.e. citizens, business and governments, i.e. G2C, G2B, G2G.

However, a public organization that is now considering to move to the Cloud should decide the appropriate cloud model that should use. In this direction, the following parameters should be taken under consideration, in order to elect the ideal cloud model that best suits organization.

The difficulty to implement each cloud model should be taken under consideration. The decision is not easy as each model suits better different situations. If a public body would like to use an online application without extended configuration and SaaS approach is offered, this might be an ideal solution, due to the simplicity to adopt. However, specificities like imposed regulations might be tough to be accomplished in this scenario. In these cases the IaaS paradigm might be more appropriate as it offers the greatest amount of

control in of the deployment environment. IaaS adoption however means that the organization is responsible creating a cloud ready application, deploy it and properly secure the server. STRATEGIC Service Store is a tool that offers the best of both worlds; it allows the configuration of both the deployment environment and the application through easy to use menus, and the same time it offers security mechanisms out of the box.

The maturity of the different cloud models can vary. Although a general conclusion it is not easy to be provided, SaaS and IaaS are the most dominant models, while PaaS reflects a smaller market with lots and often changes.

The focus of each model is different and should be taken under consideration. SaaS is commonly adopted for the usage of online services that are provided ready for use. PaaS in used for the deployment of databases and web applications created developers. IaaS are used to provide storage, server and networks that should be configured properly and then used of the deployment of applications.

The added value that each model provides is also different. The ease of deployment that SaaS provides can be considered an added value as it allows the deployment of a product ready solution with a very short turnaround time.

2.3 Selection of Appropriate Cloud Computing Deployment Model

Apart from differences that exist in the cloud paradigms and should be taken under consideration by pubic bodies, a major decision is the deployment model of the cloud. Cloud hosting deployment models are mainly distinguished by the proprietorship and accessibility of the infrastructure, but also by the purpose and the nature of the cloud. Therefore, before investing on cloud, public administrations should be well aware of the deployment models and select the appropriate based on the requirements of the services to be deployed.

There are two main types of cloud; public and private, but there are three more models that are also possible to be used by public bodies; virtual private cloud, hybrid and community cloud.

Public Cloud: is a type of cloud hosting in which the cloud services are delivered over a network which is open for public usage. This model is a true representation of cloud hosting; in this the service provider renders services and infrastructure to various clients. The customers do not have any distinguishability and control over the location of the infrastructure.

Private Cloud: is also known as internal cloud; the platform for cloud computing is implemented on a cloud-based secure environment that is safeguarded by a firewall which is under the governance of the IT department that belongs to the particular corporate.

Virtual Private Cloud: is an on-demand configurable pool of shared computing resources allocated within a public cloud environment, providing a certain level of isolation between the different organizations using the resources[6].

[6] http://blog.unitedlayer.com/cloud-computing-terminologies

Hybrid Cloud: can be an arrangement of two or more cloud servers, i.e. private, public or community cloud that is bound together but remain individual entities.

Community Cloud: is a type of cloud hosting in which the setup is mutually shared between many organisations that belong to a particular community

Based on our experience we have concluded on the following best practices and factors that should be taken under account.

It is important to find the right balance between control and convenience. With public clouds, there is no need to worry about certain things like maintenance, but controllability is limited. With private clouds, more freedom is offered and a public administrator is responsible on how everything works.

Public cloud is better suited for applications that have more frequent peak times. A public cloud benefits on the decreased capital overheads that a service with spikes on the usages needs.

Scalability and flexibility of public clouds is considered to be greater. Private clouds can achieve high scalability and flexibility standards, as public clouds do, but a lot of investment is needed.

Private cloud gives the organisation greater and direct control over their data, as the organization is managing everything and as data can be more safe as it is not publicly available.

Highly sensitive materials usually are stored to private cloud. This is because as much as public clouds are somewhat secure; your data could still fall in the wrong hands[7].

Although private cloud is generally preferred regards to security and data privacy, in case of natural disaster and internal data theft the **private cloud may be prone to vulnerabilities**[8].

Private cloud is better suited for public organizations that may change their needs, organizations that have restricted management demands, host mission critical services, and in general **prefer to have fully control of their services and infrastructure**.

Hybrid cloud is ideal for utilizing the benefits of both the private and public deployment models. A typical example is increasing the capacity of the cloud by aggregating different cloud offerings or services.

In hybrid cloud solutions, **resources that are non-critical like development and test workloads can be housed in the public cloud that belongs to a third-party provider**, while the workloads that are critical or sensitive can be housed internally. During the piloting period of STRATEGIC, we utilized hybrid based scenarios, as applications were deployed initially in public cloud of Amazon.

If high demand on resources is needed, **hybrid cloud can be utilized in order to support cloud bursting** and allow public cloud to serve specific instances.

[7]https://www.esds.co.in/blog/how-to-choose-and-implement-the-correct-cloud-model/#sthash.sA8lZCg3.WT10k52N.dpuf
[8]https://www.ibm.com/developerworks/community/blogs/722f6200-f4ca-4eb3-9d64-8d2b58b2d4e8/entry/4_Types_of_Cloud_Computing_Deployment_Model_You_Need_to_Know1?lang=en

Community cloud can be used by public bodies in order to share common resources and have great savings on the cost of the cloud. This is possible as public bodies generally have similar privacy, performance and security concerns. A community cloud may be internally managed or it can be managed by a third-party provider. It can be hosted externally or internally[9].

For Virtual Private Clouds, the isolation of the organization can be achieved through allocation of a private IP subnet and a virtual communication construct such as a VLAN. This case was tested by STRATEGIC for the pilot of MoSG.

2.4 Selecting Cloud Provider and the Benefits of Cloud Brokerage

After the selection of the cloud model (IaaS, PaaS, SaaS), a public body has to select not only deployment model (private, public, etc.), but also **must select the solution or provider that will use.** We consider the IaaS model the most relevant for the usage by public bodies, so the IaaS market had been extensively analysed in the project duration.

IaaS is a diverging market with IaaS providers that range from small start-ups with products based on open source solutions, like OpenStack or CloudStack to big companies like Amazon, Rackspace that follow their own approach and have big market shares. We have collected the following advices for the selection of the appropriate cloud provider or cloud platform.

Calculating the actual cost of a cloud provider is not always easy, as IaaS is more than just subscription fees.

For public administration, it is important to extensively **check SLAs and agreements prior to signing contracts**.

For ensuring that public cloud offers adequate security, **question about the disaster recovery mechanisms** that are used.

For private cloud installation both open source (OpenStack, CloudStack) and commercial (VMware VCloud, Microsoft Azure) can be used. Before selecting the platform to be used, an **evaluation regarding compatibility should be done to both infrastructure and software level**.

Complexity of installation can be a differentiation factor. Many companies provide support of the installation of both open source and commercial offerings, so it might be better to contact a specialized company for the installation.

There might be some **compromises to be made in order to find the ideal point between maturity and state of the ar**t. In our pilot examples we had to invest to the latest version of Ubuntu OS(16.04) in order to benefit from the Long Term Support and this also imposed a specific version of OpenStack.

Security of the cloud platform should be taken under consideration, including an analysis of the underlying technology that each platform uses.

[9]https://www.ibm.com/developerworks/community/blogs/722f6200-f4ca-4eb3-9d64-8d2b58b2d4e8/entry/4_Types_of_Cloud_Computing_Deployment_Model_You_Need_to_Know1?lang=en

Extensibility of the platform is also important, especially if planning to create an initial cloud setup that grows with the organization needs.

Cost analysis of the IaaS should be done and it should take under consideration both one off and running costs, including the pricing for the license of the solution, the installation resources cost, the resources spent for the hosting of the Cloud infrastructure services and the efficiency of the VMs that are created, in terms of the compute power available to the deployed service.

In many cases **using a Cloud Brokerage platform can be helpful**. Cloud users have to choose between a huge number of vendors and they would also have to manage the integration between the various platforms. The overhead involved discourages many businesses from building multi-cloud environments, obviating many of the benefits of the cloud[10]. In STRATEGIC we try to make governmental bodies confident with cloud usage and try to help on the diversity of IaaS market by **supporting multiple IaaS cloud providers** that can be used for the deployment of services of governmental bodies[9].

The **usage of Cloud Application Management solutions is highly recommended for entities interested in fast deployment**. Cloud Application Management is an emerging category[11] of solutions that is expressing the need to move and manage application workloads to the cloud, instead of the more traditional Cloud Management solutions that are just orchestrating the physical or virtual infrastructure level. STRATEGIC Service Store is a multicloud service with lifecycle management & a marketplace that enables governmental bodies to use both private and public clouds for the easy deployment of services. The multi-cloud coverage and the support for both private and public clouds allow STRATEGIC Service Store to be used by public bodies to create hybrid and federated clouds and exploit the benefits of cloud computing era with minimum effort.

2.5 Prepare and Design a Private OpenStack Installation

OpenStack is an open source cloud computing platform that provides an Infrastructure-as-a-Service (IaaS) solution through a variety of complemental services. Each service offers an application programming interface (API) that facilitates this integration, therefore the overall OpenStack installation is an integration of services.

One of the objectives of STRATEGIC is to help public bodies with low or without previous experience in cloud to prepare their own private cloud infrastructure and integrate this infrastructure with the STRATEGIC Service Store. For this reason, the following sections reflects experiences of our consultancy dully collected from the technical and pilot partners of the consortium and it is provided for the benefit of Administrators and IT experts of public bodies.

2.5.1 Identifying hardware needs

The OpenStack hardware requirements vary a great deal, depending on the desired target deployment type, storage backend choices and services co-location policies.

[10] https://www.computenext.com/blog/the-three-minute-guide-to-cloud-marketplaces/

[11] http://www.appcara.com/wp-content/uploads/2014/07/Cloud-Management-versus-Cloud-App-Management-v2.0.pdf

The absolute minimal setup for the suggested installation supported from STRATEGIC Service Store can consist of 3 physical servers: Production grade OpenStack deployments usually require at least 5-9 physical servers, depending on high-availability and storage backend choices, while a high-availability mode for all OpenStack services (ie N+1 resilience where possible) would need 28 service units (if no service co-location is imposed).

The Openstack **server nodes require at least two physical network interfaces** ports/trunks (1Gbit speed or better), together with VLAN capable core-switch and router hardware. Networked storage backends for production grade deployment – using iSCSI or FiberChannel SAN or Ceph backends – need additional high-speed and dedicated network trunks (FC 8Gbit or 10Gbit Ethernet) for better stability and performance.

For the preparation of OpenStack, **selecting storage hardware is an important decision**. Storage hardware and the architecture should be selected by evaluating possible against the following critical factors, the user requirements, technical, and operational considerations[12].

Selecting networking architecture determines which network hardware will be used. **Networking software is determined by the selected networking hardware**. The selection of certain networking hardware (and the networking software) affects the management tools that can be used. There are exceptions to this; the rise of open networking software that supports a range of networking hardware means that there are instances where the relationship between networking hardware and networking software are not as tightly defined[13].

The **operating system (OS) and hypervisor have a significant impact on the overall design** of the cloud. Selecting a particular operating system and hypervisor can directly affect server hardware selection and vice versa. Public administration should verify that the storage hardware, the topology support and the networking hardware selection will work with the chosen operating system and hypervisor combination[14].

As OpenStack is a modular platform comprised by many services/components, it is a **part of the design process to select the OpenStack components to use**. Some OpenStack components, like compute and Image service, are required in every architecture. Other components, like Orchestration, are not always required. Excluding certain OpenStack components can limit or constrain the functionality of other components, therefore it is **important to research the component interdependencies** in conjunction with the technical requirements before deciding on the final architecture[15].

For private installation of OpenStack that are connecting to external brokers (like STRATEGIC Service Store) API connection tunneling firewall and VPN devices/services are also needed.

Installation of OpenStack is possible also on virtualized infrastructure. However, there are difficulties to be tackled like the networking configuration of

[12]http://docs.openstack.org/arch-design/generalpurpose-architecture.html#selecting-storage-hardware
[13]http://docs.openstack.org/arch-design/generalpurpose-architecture.html#selecting-networking-hardware
[14] http://docs.openstack.org/arch-design/generalpurpose-architecture.html#operating-system-and-hypervisor
[15] http://docs.openstack.org/arch-design/generalpurpose-architecture.html#openstack-components

the virtual machines, while the underlying virtualization technology is also not always supported. Therefore, we suggest to use this option only if there is not possible to use physical machines. **Installing IaaS on physical infrastructure is preferred in terms of performance** and due to more straightforward installation.

OpenStack setups can have specific focus and use the nodes in appropriate way. The list below presents the most common setups of OpenStack

- Compute focused
- Storage focused
- Network focused
- Multi-site
- Hybrid
- Massively scalable

We found that the normal, **compute focused setup of OpenStack was ideal for the pilots of the project**.

For a production architecture deployment of OpenStack **security methods such as firewalls, encryption, and service policies should be always used**.

Extra nodes for core and optional services can be added in order to achieve performance and redundancy requirements.

2.6 Installation of IaaS Using Automated Scripts

One of the objectives of STRATEGIC is to help public bodies with low or without previous experience in cloud to prepare their own private cloud infrastructure and integrate this infrastructure with the STRATEGIC Service Store. The setup of an open source solution like OpenStack is the suggested way of creating a private IaaS that can be connected on STRATEGIC Services Store.

A possible installation approach of OpenStack is using automated scripts that help on the installation process. In STRATEGIC we have experienced two of the most popular OpenStack distributions(Ubuntu OpenStack, Mirantis OpenStack) in order to use create IaaS installations that were used for the deployments through STRATEGIC Service Store.

2.6.1 Installation of OpenStack IaaS Using MaaS and Juju

Ubuntu Metal-as-a-Service (MAAS) allows treating physical servers like virtual machines in the cloud. It allows connecting and commissioning physical servers easily even after the initial setup of an environment. In conjunction with the Juju service orchestration software MaaS will enable dynamic deployment of complex OpenStack services with ease and confidence. Juju[16] is a service orchestrator provided by Ubuntu that can be used on top of MaaS and supports the deployment of complex service like OpenStack. Juju orchestration software packages can be installed on MAAS server:

[16] http://www.ubuntu.com/cloud/juju

The OpenStack hardware requirements vary a great deal, depending on the desired target deployment type, storage backend choices and services co-location policies. Bare minimal OpenStack setup with MAAS and Juju can be deployed on three physical servers, yet high-availability mode for all OpenStack services (ie N+1 resilience where possible) would need 28 service units (if no service co-location is imposed).

Absolute minimal setup for an OpenStack installation that can be supported by STRATEGIC Service Store consists of 3 physical servers: MAAS server, single OpenStack controller node and one Compute node. Typical patterns for minimal OpenStack deployments are:

- Evaluation/testing setup with local storage: 1x MAAS node + 1x Controller node + 1x Compute node
- Production grade setup with SAN backed storage: 1x MAAS node + 3x HA Controller nodes + 2x Compute nodes
- Production grade setup with scale-out Ceph storage backend: 1x MAAS node + 3x HA Controller nodes + 3x Ceph storage nodes + 2x Compute nodes

The OpenStack server nodes require at least two physical network interface ports/trunks (1Gbit speed or better), together with VLAN capable core-switch and router hardware.

MAAS server depends on network access to servers out-of-band BMC IPMI hardware management interfaces, in order to control their power state - which is required for MaaS operation.

Networked storage backends for production grade deployments – using iSCSI or FiberChannel SAN or Ceph backends – need additional high-speed and dedicated network trunks (FC 8Gbit or 10Gbit Ethernet) for better stability and performance.

For gateway network security and Service Store API connection tunneling firewall and VPN devices/services are also needed.

Detailed installation guide for OpenStack using MaaS and Juju can be found in the project website.

2.6.2 Mirantis based installation of OpenStack

Mirantis OpenStack is a vendor supported OpenStack distribution provided by Mirantis, an US based B2B cloud computing services company. Mirantis launched its own OpenStack distribution in October 2013, in competition with OpenStack distributions from RedHat, Hewlett-Packard and others. One of the key features of Mirantis OpenStack is Fuel - an open-source software application that simplifies the deployment of highly available OpenStack environments, as well as enables you to manage your OpenStack environments after deployment. Mirantis OpenStack is available as a free download (for evaluation purposes) and Fuel project has become an official OpenStack Foundation governed project in 2015.

Mirantis Fuel and it's easy to use web-based user interface makes it simpler to deploy production grade highly-available OpenStack environments.

Fuel software provides network based installation and hardware inventory services out-of-the-box - allowing you to boot, configure and install OpenStack nodes from single web-based console.

For advanced use cases Fuel provides also CLI and REST API management interfaces, making the automation of whole process possible.

When comparing Mirantis OpenStack to Ubuntu Juju OpenStack we have found it to be better documented, easier to install and creating more production-ready OpenStack environment configuration by default - without the need for heavy customizations.

Also Fuel supports plugin system - where its functionality can be extended via third-party plugins.

Fuel based OpenStack deployment process is pretty straightforward, involving the following steps:

- Downloading Mirantis OpenStack ISO image from Mirantis website
- Configuring network hardware according to your VLAN networking plan
- Deploying and configuring Fuel master server
- Booting server nodes for hardware discovery and inventory
- Creating target OpenStack environment configuration and mapping hardware in Fuel
- Deploying target OpenStack environment with Fuel

Mirantis has been used by NICPB in order to make an IaaS installation that was connected to the STRATEGIC Service Store and was used for deployment of test services. Detailed Mirantis OpenStack installation guide can be found in the project website.

3 Protecting Applications and Data on Cloud Environments Using STRATEGIC

STRATEGIC is applying a *click to secure* policy to the Service Store. *Click to secure* is an approach that is designed to allow STRATEGIC customers to own security and protection of their applications. This solution is adapted to environment, such as STRATEGIC, that can target multiple clouds.

Conventionally, Cloud Services are divided into three categories which are IaaS, PaaS, and SaaS. They reflect the degree to which the customer is responsible for the environment stack. STRATEGIC solution proposes a more flexible framework where customers can benefit from all the functionalities of a PaaS but also, STRATEGIC allows multiple cloud targets to be accessible. This additional level of flexibility singles out STRATEGIC from conventional PaaS providers. However this also has two collateral effects. The first one is the requirement to manage deployment packages (services, or applications), and as a result, security cannot be managed in a single virtualisation environment. Therefore, the second effect is the decoupling of the application level security with the virtualisation platform. **Figure 2** highlights the differences between conventional IaaS, PaaS, and SaaS, and STRATEGIC.

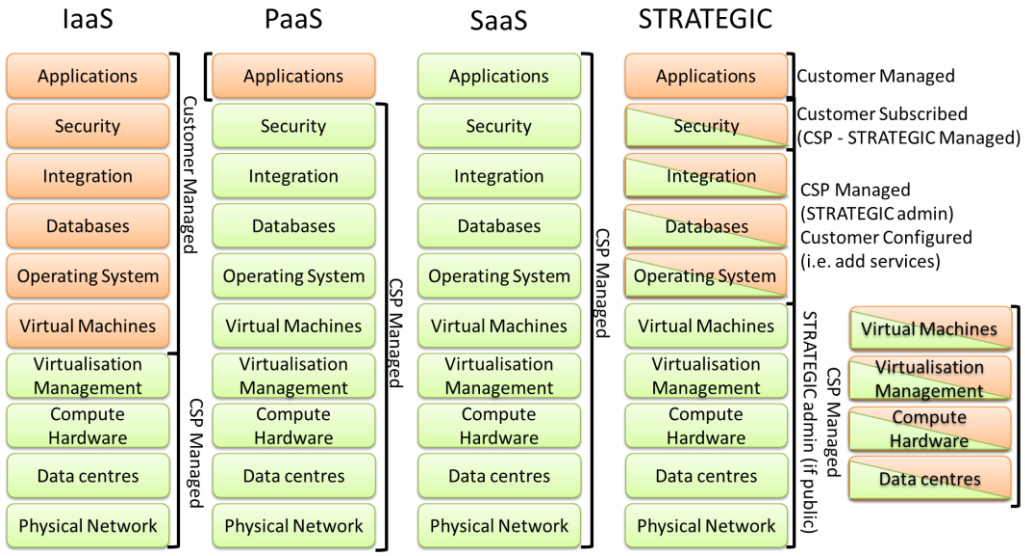

Figure 2: How STRATEGIC differs from conventional IaaS, PaaS, and SaaS

The security layer provided by STRATEGIC cannot be controlled in the same way as a conventional PaaS. This de-coupling in effect transfers the security burden over to the customer. STRATEGIC is addressing this problem by providing a subscription based model for security.

The customer experience has been simplified to subscribe to security components including protection of the Virtual Machines and applications. This includes patching, anti-malware, and anti-intrusion. The same experience is available for data encryption. The customer experience is referred to as click to secure because there is no other interaction required. Equally, customers can click to cancel, or update. In effect subscribing to the BT security features procures a managed security service which is itself highly configurable.

Compared to conventional PaaS, this approach requires customers to select security. This is one additional step that did not exist in single cloud PaaS however, this extra functionality comes with benefits too. It is possible to customise security and choose to either strengthen or cut back on the assurance levels which has an impact on the cost. Also, a level of security can be designed to achieve a particular compliance requirement. Not only have we seen that security services can be deployed to multiple cloud environments, but, they also work as a compliance management tool.

Security Services in the Service Store are integrated by the STRATEGIC administrator. Customers benefits from the security services available without the need to be security savvy. Currently, security services that have been integrated include BT Intelligent protection, and BT Data encryption. These services are scheduled to be implemented in production within 2017 in commercial implementation of BT.

4 Development of cloud services using service store

4.1 Challenges for public sector cloud services

STRATEGIC aims to help public sector with cloud adoption, and during the piloting period, we had three different public bodies, dissimilar size and level of cloud adoption. All pilots tried to increase cloud adoption and extend their portfolio of cloud services that use internally or provide to their citizens. These are some of the main findings from their experience.

When considering public administrations of medium-size, it is preferable at least when possible, to draft a **clear list of complete requirements both in term of material resources and tasks to be accomplished** given in the least possible chunks instead of many little tasks. This is easily explained by considering that each task to be accomplished follows a first-in-first-out ticket queue processing by PA personnel. Also, regarding the material resources to acquire either a tender or some kind of internal agreement needs to be made and that requires time.

We also found that **medium-large cities seem to have skilled personnel to carry on the required tasks** with the necessary professionalism in order to deploy software or even in some cases to prepare an IaaS environment.

The **decision of the IaaS to be used is an important step** for public bodies, regardless the size. Most public bodies preferred to use a private cloud for their services.

In some countries finding cloud hosting providers is not an easy task, and the **choices for public bodies interested to store their data with location based restrictions can be very limited** or non-existent. In one of our pilot cases, using a public cloud was preferred option but due to regulation rules that forced the storage of data to be within the country, a private cloud installation was created and used.

For public bodies making a tender is usually needed for the purchasing of hardware. This can cause some delays and difficulties that have to be resolved in order not to affect the cloud adoption. So, a **dillema regarding the options for the IaaS implementation** for a public administration entity can be created, with the following options representing the findings of one of our pilots:

1. Purchassing own physical machines, install IaaS and maintaining it internally – in own premises
2. Purchassing own physical machines, install IaaS and maintaining it in a provider’s premises
3. Renting VMs and install OpenStack and applications
4. Renting VMs with some cloud infrastructure and install applications
5. Renting physical machines and install OpenStack and applications

For our pilot case, the opinion was that the options 2 and 5 were the most acceptabe. A problem is that it is that hosting providers don’t always offer cloud infrastructure or physical machines.

Finally we found that the usage of the STRATEGIC Service Store on top of a private cloud infastructure provided a positive experience, as public bodies found

that there were many applications that could choose, easily adapt and deploy to their premises. A recommendation of one of our pilots is is that **STRATEGIC Service Store should be a mandatory part of any cloud project for the public administration** body no matter if this is implemented as the private, public or hybrid cloud. In other words, the **Cloud Application Management solution** that is offered by STRATEGIC provides many advantages (security, easy replication, etc.) that should be used in any cloud based project for the public administration body.

4.2 Guideline for the Creation of Adaptable and Reusable Applications with Strategic Service Store

Cloud Application Developers use STRATEGIC Service Store for the preparation and creation of applications. As Cloud Application Developers have technical background they have to use advanced capabilities of STRATEGIC Service Store to prepare their applications through the user interface and also by creating scripts when needed.

In order to create a new application and publish it in the STRATEGIC Service Store, a developer start by providing a descriptive application name, application icon and also Description and License type. A "How to get started" section should contain general steps on how a user can deploy and access the application.

It is important to be set accessibility setting to "public" when application is ready to be fully published.

Once the application profile is completed and updated/saved, we need to create new application topology to edit:

After its creation, **the topology can be edited** in order to set up the desired service through the following editable sections:

Parameters – used for defining user filled input variables which can be used in the service provisioning script(s)

Servers – defining the nodes/VMs that form the application, together with their firewall settings and the provisioning steps/scripts

Scripts – defining provisioning scripts that are used to install and configure services inside VMs

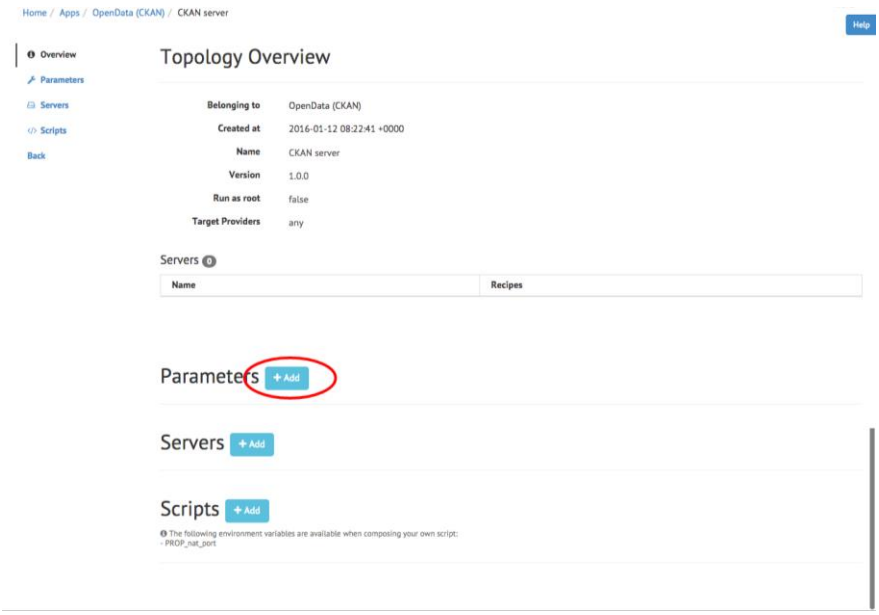

Figure 3: Defining the application topology through STRATEGIC Service Store

Requirements can be set to specify the needs of the application in OS and hardware level. Specify VM minimal requirements – like "Min. Memory", "Min. VCPU" – also which VM operating system templates are supported (for example: Ubuntu 12.04). It is important that VM OS templates listed here have to be available from target cloud image catalogue.

Firewall rules can also be added in order to secure the application.

Adding scripts-recipes is the core part of the packaging. Service provisioning inside the VM requires at least a single Recipe to be executed. The scripts can use the input variables defined earlier by the developer and are filled by the users. We have concluded in the simple rule that **if the installation process of an application can be executed with scripts, it can be packaged and added to STRATEGIC Service Store**. Most common script languages are supported.

More details about the creation of an application that is deployable through STRATEGIC can be found at the project website.

5 Application Marketplaces as Part of Public Services Deployment

5.1 Guidelines for the Creation and Sustainability of Application Marketplaces

The Service Store is a multi-cloud target environment which allows the launch and in-life management of applications. Customers can also be allowed to insert their own application into the service store. Application packaging allows customers to create a default configuration which is agnostic of which cloud target customers choose to deploy to. STRATEGIC as a multi-cloud broker and application marketplace, offers the following capabilities.

- Support for integration with private IaaS clouds
- Support for public IaaS clouds
- Support for shareable application templates
- Support for reselling of the applications
- Support for reporting on resource usage
- Matchmaking of resources based on EU location
- Matchmaking of resources based on EU legal system
- Federated authentication support

Pilots have been interested in migration projects. The process for a migration requires knowledge about the application installation. A step by step process needs to be established, and automated. There should not be any manual intervention required even when there are multiple scripts, and multiple servers involved.

Our approach consisted in reproducing the installation process manually in order to establish the sequence of events are installations steps required. This process reveals the following:

- Step by step installation commands
- Dependencies
- Working Operating Systems
- Cross server configuration requirements

Once this is established the installation script can be written which will execute the following steps: load configuration, install all dependencies, install application. However, fully automated **installation scripts also require debugging**. This process is done mostly by log analysis. Unfortunately, during long installation processes logs took a long time to be available leading to slow and inefficient debugging sessions. **Two main strategies are available** to work around this problem. Primarily the **breaking up of the installation script** in smaller independent instructions. There might be no need for the full script to be run if a failure happens at the end. An entire deployment script can be cloned and renamed appropriately so as to run tests and fix installation issues. The other strategy consists of **moving installation scripts in public repositories** (e.g. AWS S3, or GitHub) and downloading those scripts on the fly. This has the benefit that errors can be corrected and application launch can be tested without having to access the Service Store package directly.

Some applications can be time consuming to migrate and the process may require expertise from different people in one organisation. **Clear governance should be established at the start of the migration process**. The migration process requires access to the virtualisation platform enough provisioning should be planned so that accounts being used have all the necessary access rights before working sessions so as to avoid un-necessary last minute cancellations of working sessions.

Applications that need **multiple servers and heavy installation processes** (e.g. interconnected web servers and large resilient database systems on Windows servers) **demand a lot more time and resources**. This should be taken into account when resourcing the migration process. As a rule of thumb, a migration process can take between two weeks and two months for bigger applications even if man-power has been resourced correctly.

Some **considerations to be taken into account when resourcing include**: what is the governance to access to existing application, **who is knowledgeable** about the software stack deployed, **who owns security of the application**, **who is technically capable** of migrating the application, and the security around it. It is also critical to **involve the administrators of the virtualisation layer**, i.e. who can provision virtual resources. STRATEGIC administrators may also be required if a new Operating System template is required, or for the connection of a private cloud (see next section).

The Service Store keeps a record of application deployment and usage. It is possible for an application owner to obtain reports to measure the success of their application. This includes usage data, and download frequency. **There is no billing management associated with the Service Store** therefore this functionality is subject to integration with new or legacy application.

Inevitably, **changes in Cloud Service Providers APIs** require software updates of the Service Store. This has **occasionally disrupted deployment** especially if the Cloud Service Provider has immature APIs. For instance AWS is extremely reliable and API changes are managed with sufficient warning to minimise disruption. On the other hand Microsoft Azure is more volatile in its functionality and less stable with their public APIs which can create launch issues. Equally, a private, or, virtual cloud update has to be validated by the STRATEGIC administrator ahead of its deployment to ensure compatibility (see next section).

5.2 Connecting Private Clouds to a Public Marketplace

The STRATEGIC Service Store is a multi-cloud proposition. The ability to manage and deploy applications to a public cloud infrastructure at the same time as to a customer's Private or managed infrastructure was considered a requirement. The Service Store enables this by accepting a credentials from customer for a particular target cloud and then using its cloud connection SDK to communicate with the cloud infrastructures. The method is the same for all cloud infrastructures be it private (Openstack, Cloudstack) or public (AWS, Azure, BT Cloud, etc.).

Many Public cloud infrastructures have the necessary APIs exposed with sufficient security, such that a customer would generate an access key/password. This process of access generation is supported by an effective management policy that supports the management of its lifecycle. **The customer's responsibility, hence, is to ensure an approved and usable access key/password is**

available to the Service Store to communicate with the respective infrastructure.

In comparison, the private infrastructures are either based on Openstack or Cloudstack, that are primarily intended for data centre management, and lack the public API exposure management systems. In STRATEGIC, **the solution was to adopt a heightened security policy with restricted access control list (ACL)** adhering to the security policies of the customer to whom the infrastructure belonged to. There were three types of access that were recommended in combination or on their own; there were given as options to choose from:

a) STRATEGIC Service Store IP address to be white listed by the Infrastructure provider in order to allow IP restricted access to the infrastructure APIs.

b) Creation of a site-to-site VPN between the Service Store provider and the infrastructure provider, with an ACL such that only a predefined set of Service Store servers are allowed to communicate with infrastructure APIs.

c) Creation of a IPsec VPN server by the infrastructure provider, to which the Service Store provider creates a persistent VPN connection via a proxy server hosted by themselves to gain access to the infrastructure APIs.

A few things to be considered while adopting the above solutions are to **ensure a secure communication channel to share pre-defined VPN keys** and to ensure the API key management policies by the infrastructure provider are made aware to the customer such that they can update the Service Store with newer keys whenever necessary. Finally, **the passwords and keys needs to be of sufficient strength and rotated sufficiently enough** that advanced persistent threats can be detected or averted on a periodic manner.

5.3 Accessing a Private OpenStack Installation Through STRATEGIC Service Store

After an OpenStack installation, has been successfully created, it shall be configured in order to be used. Although there are many tutorials online presenting these steps, we have collected the steps that are needed in order to connect an OpenStack of a public organization on STRATEGIC Service Store.

5.3.1 Creating Users and Projects

At first, appropriate users and projects should be created. A project is created by the Admin menu of OpenStack, as shown in Figure 4.

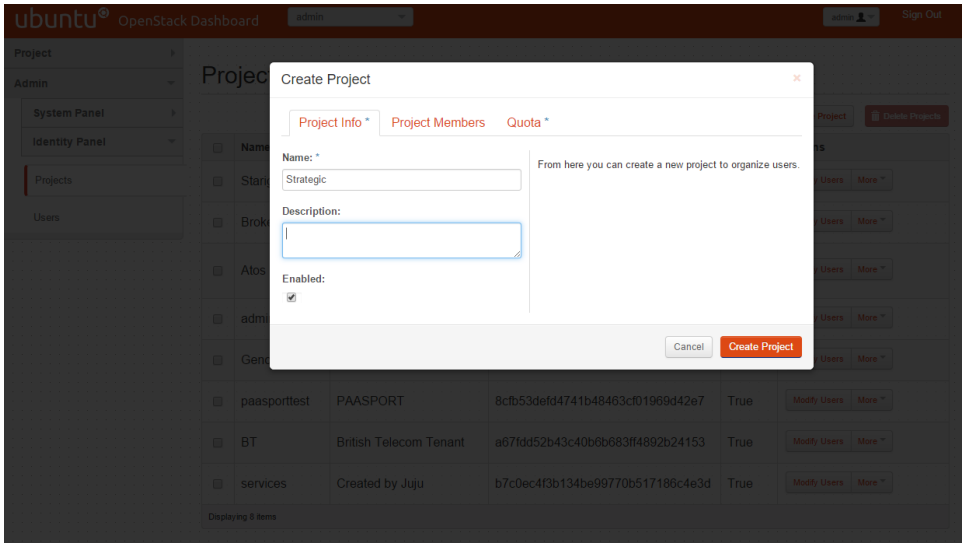

Figure 4: Create a new project

If users exist already they can be added directly to the project and if a new user is created, the main project of the user can be assigned. The newly added users can use their credentials in order to authenticate though the STRATEGIC Service Store.

5.3.2 Creating OpenStack Flavors

On OpenStack IaaS, these templates are called flavors[17] and the specific flavors reflecting the needs of a specific organization can be created. These flavors will be available in the STRATEGIC Service Store as option for the deployment of the service.

5.3.3 Creating OpenStack networks

In order to deploy an application on OpenStack the creation of appropriate networks is needed. The networks can be internal or external and for privacy issues a network should be assigned to a single project. After their creation, networks of the project/tenant are displayed along with their status .

5.3.4 Verification of OpenStack API endpoints

When installing OpenStack, it should be noted that the API endpoints should be exposed to STRATEGIC Service Store. OpenStack installer (FUEL) usually deploys both public and internal API endpoints within the system - as demonstrated below.

STRATEGIC Service Store should be seeing a localized IP (in the network of STRATEGIC Service Store) for the primary API. However, since the primary API provides the IP addresses of the other APIs the API endpoint should be accessible with domain URL. Even a private DNS name is possible to be used.

5.3.5 Integration Points

[17] http://docs.openstack.org/openstack-ops/content/flavors.html

In order for the STRATEGIC Service Store to be integrated on behalf of a user, all the API endpoints presented above should be accessible.

Also, there are different ports for different OpenStack components. In order support the particular functionalities offered by each component, remote OpenStack site need to open the ports presented in the following table:

OpenStack service	Default ports	Port type
Block Storage(cinder)	8776	
Compute(nova) endpoints	8774	
Compute API(nova-api)	8773,8775	
Compute ports for access to virtual machine consoles	5900-5999	
Compute VNC proxy for browsers(OpenStack-nova-novncproxy)	6080	
Compute VNC proxy for traditional VNC clients (OpenStack-nova-xvpvncproxy)	6081	
Proxy port for HTML5 console used by Compute service	6082	
Identity service (keystone) administrative endpoint	355357	adminurl
identity service public endpoint	5000	publicurl
Image Service (glance) API	9292	publicurl and adminurl
Image Service registry	9191	
Networking(neutron)	9696	publicUrl and adminurl
Object Storage(swift)	6001-6002	
Orchestration(heat)endpoint	8004	publicurl and admin url
Orchestration AWS CloudFormation-compatible API (OpenStack-heat-api-cfn)	8000	

OpenStack service	Default ports	Port type
Orchestration(heat)endpoint	8004	publicurl and admin url
Orchestration(heat)endpoint	8004	publicurl and admin url

Table 4: OpenStack services and default ports

6 Development of cloud services using SEMIRAMIS

STRATEGIC framework has integrated the SEMIRAMIS results with the aim of providing trusted and secure components to cloud services provided by the public administrations. The cross-border attribute exchange service has been integrated in the cloud Certificate of Residence Service provided by both the StariGrad and the Genoa municipalities.

6.1 SEMIRAMIS usage by public bodies for the exchange of cross-border information

European public administrations need to be adapted to a dramatic changing world where the necessity to communicate each other to exchange information, not only at local or regional level but also in a cross-border scenario, are increasingly growing. In this sense the public bodies are covering the increased citizens' demands of online services. In this context ICT helps the administrations to provide more secure and trusted services.

Leveraging SEMIRAMIS outcomes the Certificate of Residence service can improve the secure access to information and the exchange of information in cross–border scenarios.

6.1.1 Current situation

During the project the SEMIRAMIS components were updated for a better integration with the services of the pilots. In order to allow working in a cloud environment provided by the STRATEGIC infrastructure, the SEMIRAMIS components were also configured.

Based on the description of trust and security components provided in D2.3 Framework Architecture and Technical Specifications [2], a prototype comprising two components, the Federation Proxy (FP) and the Identity Aggregator (IA), was implemented [3].

In a second phase these two components were improved [4] including:

- *Policy management*, allowing the use of "attribute release" policies applicable to the FP, which establishes which attributes may not be released by the federation.
- *Integrity of the data*: achieved using a more secure connection through HTTPS connections between the different components.

At this moment, the SEMIRAMIS components, i.e. IA and FP are deployed and running on the cloud environment owned by each municipality (Genoa and StariGrad).

6.1.2 Public administration leveraging SEMIRAMIS

Public administrations can leverage the trust chain and the flexibility SEMIRAMIS infrastructure provides for exchanging information and attributes between each other.

The following figure outlines the different possibilities that SEMIRAMIS components allow. The FP component allows public bodies belonging to a specific

European federation to establish a trust interaction between different countries across Europe. It means, for instance, that any Italian municipality joined to the Italian Municipality Federation can exchange data with other municipality belonging to the Serbian Municipality Federation, as the pilot developed during the STRATEGIC project shows.

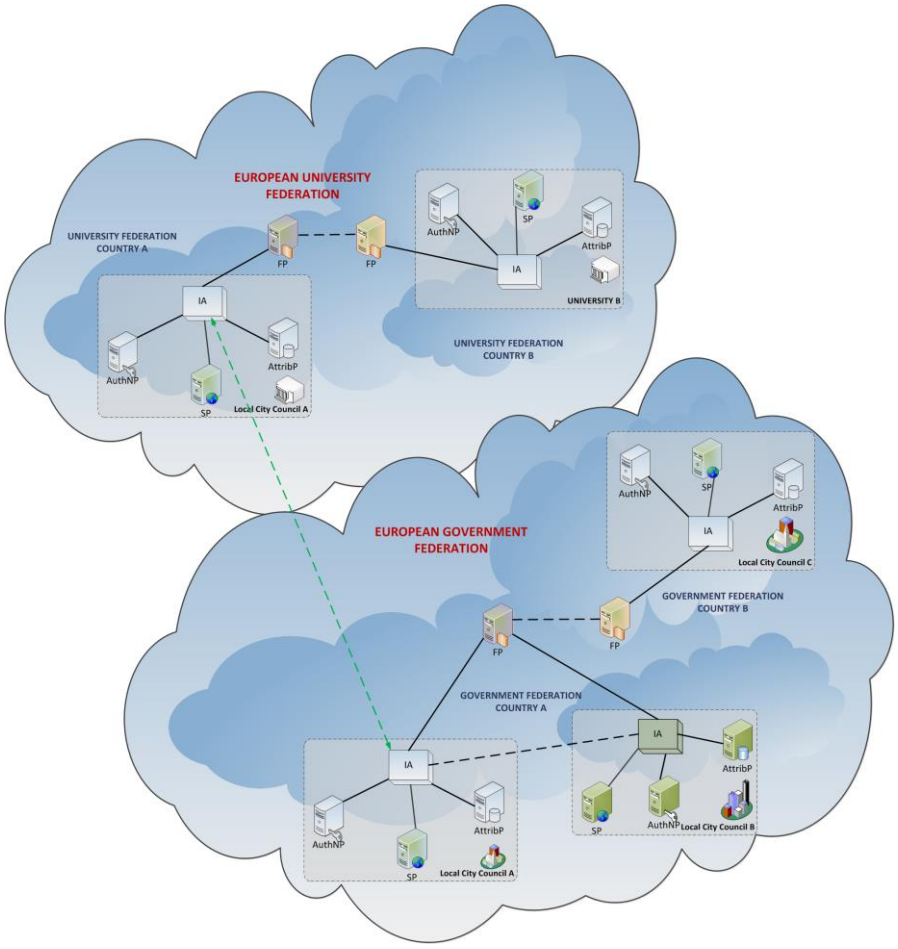

Figure 5:SEMIRAMIS Architecture overview

The IA not only allows the public bodies to interact with the FP for a cross-border interaction, but allows the public body to interact with other public administration belonging to the same federation in the same country. Indeed allows a public administration interacts with another joined public administration belonging to a different federation. For instance, an Italian municipality such as Genoa can exchange data with the municipality of Rome based on the already established trust relationship. Based on the same trust relationship the municipality of Genoa can exchange data with an Italian University. This communication between different federations is not limited to the country boundaries; thanks to the IA can establish a trusted relationship with both the IA and the FP component.

6.1.3 SEMIRAMIS' integration on STRATEGIC

Besides the SEMIRAMIS services components (IA and FF) an additional component called IA client is provided for integrating the Certificate of Residence (COR) service into the SEMIRAMIS service. This client must be embedded into the COR service for accessing the IA component and facilitates the developers' implementation work.

Additionally, mock attribute provider and authentication services are also provided to the pilots for both testing and implementation purposes.

Even though the integration process for developers comprises just a few steps, in order to facilitate the integration of SEMIRAMIS components a training session with developers and technical staff is recommended. In this way a couple of technical session has been developed with the two pilot partners, clarifying and supporting the integration process.

6.1.4 Conclusions and Lessons learnt

After the integration of the SEMIRAMIS components into the Certificate of Residence services, is worth to describe the lessons learnt during the integration process. For this purpose feedback from pilot partners were asked for improving future integrations:

1. The most valuable support for both partners was the training session provided by Atos to the pilot partner's technical staff;

2. The resources and mock applications available for the integration process, also provided by Atos, were highly helpful, avoiding extra implementation effort that could have delayed the piloting;

3. The documentation provided was good enough to start and develop the integration process. The documentation was updated during the integration phase in a regular basis;

4. A fluid communication between partners for solving the arisen issues has been developed; this has facilitated the integration process saving time and effort in both sides;

5. It was difficult to plan, at the beginning of the project the synchronization of Trust and Security solutions implementation and related services developed in pilots. Also, the training session was out of phase with the service implementation.

   Both issues were overcome thanks to the communication established as indicated in the previous point;

6. The configuration of SEMIRAMIS components should be developed by the administrator of the cloud infrastructure. An additional training session to the municipality technical staff should be scheduled for this purpose.

7. The cross-border scenarios are examples of applications/services that public administrations can use through the STRATEGIC platform.

8. The main part of discussions between partners involved was related with the list of attributes and data provided by each municipality to each other, and the user definition for piloting and testing (real vs fake users, finally was agreed to use real users but fake data).

6.1.5 Future enhancements

Due to the trust and security components have been developed in a modular way, extensions and enhancement of these features could be implemented in future developments beyond the STRATEGIC project.

7 Development of cloud services using STORK

STRATEGIC framework has integrated the STORK [6] results with the aim of providing trusted and secure component to cloud services provided by the public administrations. The Cross-Border Authentication (CBA) service has been integrated in the cloud Open Market Business Service provided by the municipality of Genoa.

7.1 STORK usage by public bodies for cross-border authentication

The use of the electronic identity (eID) for accessing securely online public administration services is being promoted by the European Commission, not only at country level but extending the use of the online services from other EU Member States in a secure way. Leveraging STORK outcomes these services can provide a secure interoperable authentication in cross-border scenarios, enabling STORK to perform authentication on their behalf.

7.1.1 Current situation

Along the project the STORK component was updated for a better integration with the service of the Genoese pilot. In order to allow working in a cloud environment provided by the STRATEGIC infrastructure, the STORK component was also configured.

Based on the description of trust and security components provided in D2.3 Framework Architecture and Technical Specifications [2] a prototype was implemented [3].

In a second phase this component was improved [4] including:

- *Non-repudiation improvement*: This was assured storing the signed assertions encrypted as a proof of the use of the Business activities engine.
- *Integrity of the data*: using HTTPS connections between the different components.

At this moment the STORK component, i.e. CBA component is deployed and running on the cloud environment owned by Genoa.

7.1.2 Public administration leveraging STORK

STORK infrastructure represents the main identity management initiative in Europe establishing a European eID interoperability platform that will allow citizens to authenticate to across borders, using their national eID.

The use of STORK allows a European public administration provides online services in a cross-border scenario. In this case the Open Market Business service provided by the municipality of Genoa uses CBA component for accessing STORK network, giving European citizens the opportunity to securely use foreign services using their eIDs.

The following figure describes the communication structure and the different components the STORK network is built up.

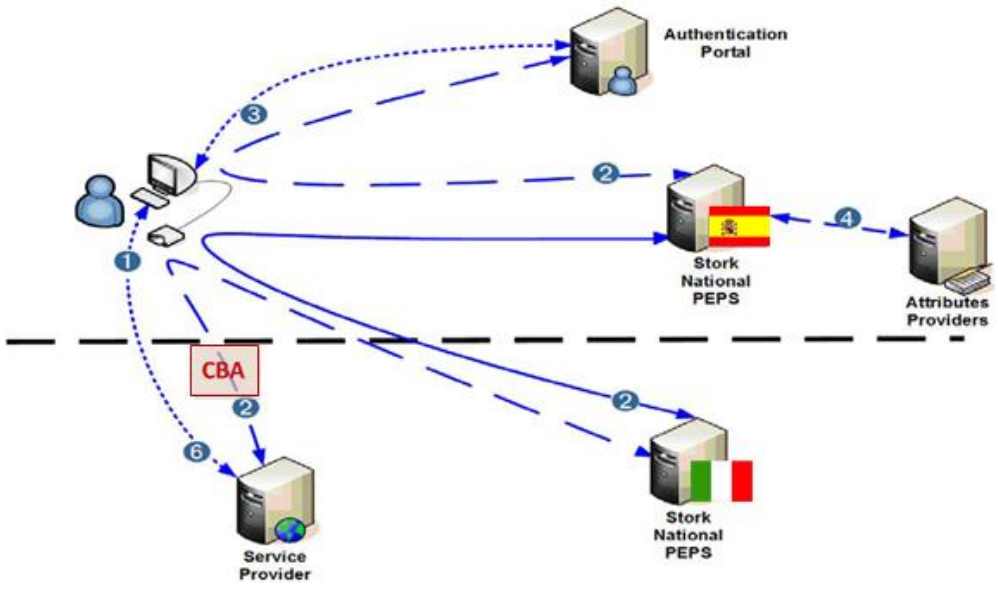

Figure 6: STORK Communication Structure.

The figure shows two PEPS, one in the country of the user (called citizen), one in the country of the service provider. In STORK, the former is called C-PEPS (Citizen-PEPS), the latter S-PEPS (Service-Provider-PEPS). The authentication process is as follows:

1. The process starts with the user accessing a STORK service provider.
2. The STORK service provider makes issues a call to the CBA component which issues an authentication request to the S-PEPS. This request declares the required attributes applying the national attribute domain of the service provider.
3. The S-PEPS translates the required attributes to the attribute domain of the citizen's country (a step called "mapping" in STORK), locates the responsible PEPS for the citizen (i.e., C-PEPS), and forwards the authentication request.
4. In response, after authenticating the user via the Authentication Portal, C-PEPS sends a SAML assertion back to S-PEPS who then applies the inverse attribute translation and issues a new assertion containing these translated attributes.
5. The service provider, on receiving this assertion, grants access to the resource.

This kind of STORK services is not restricted for cross-border authentication; Italian citizens also can take advantage of the Italian STORK network.

In summary, the use of the CBA component gives online public bodies services the access to the STORK network taking advantage of its trusted and security features.

7.1.3 STORK integration on STRATEGIC

The integration of STORK network into the cloud STRATEGIC platform is an easy process and involves the deployment of the CBA on the managed cloud, and the configuration of this STORK service.

Besides the CBA service component an additional component called CBA client is provided for integrating the Open Market Business Service into the CBA service.

This client can be embedded into the Business service for accessing the CBA component and facilitates the developers' implementation work.

The CBA client can also be used for both testing and implementation purposes.

7.1.4 Conclusions and Lessons learnt

After the integration of the STORK components into the Open Market Business service, is worth to describe the lessons learnt during the integration process. For this purpose feedback from pilot partner was asked for improving future integrations. The conclusions achieved for the SEMIRAMIS components (1 to 7 in section 6.5) also apply to the STORK component, including an additional one:

- The main part of discussions was related with the user definition for piloting and testing (real vs fake users, finally was agreed to use real users but fake data) and the data needed by the Open Market Business service that must be retrieve from the STORK network (name, surname, eIdentifier and fiscal code were agreed).

7.1.5 Future Enhancements

The STORK network has evolved from version 1.0 to STORK 2.0. In September 2014 entry into force the eIDAS Regulation (910/2014) [6] and will be compulsory to be adapted in September 2018 for public administrations as the following figure shows.

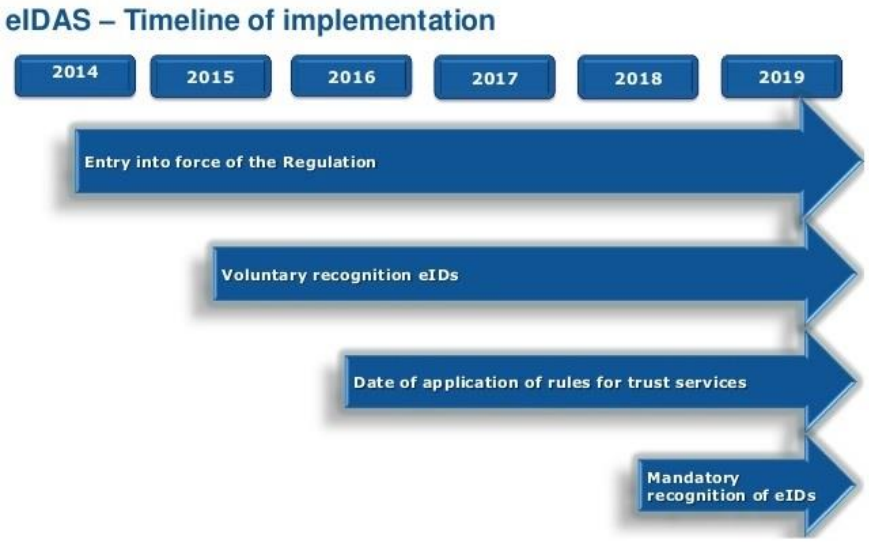

Figure 7: eIDAS Timeline of implementation

eIDAS Regulation, which includes the development of an interoperability framework and encourages governments to make their eID schemes more cross-border friendly, can be positive for STRATEGIC. The influence of the eIDAS Regulation might help to lower the legal and technical hurdles and increase the amount of possible eID means that Brokers service can support.

The eIDAS Regulation enables the use of electronic identification means and trust services (i.e. electronic signatures, time stamping, registered electronic delivery, etc) by citizens, businesses and public administrations to access on-line services or manage electronic transactions.

Currently the eSENS European project [7], based on STORK project, is trying to facilitate the deployment of cross-border digital public services through generic and re-usable technical components. There are 20 European countries involved, which means that most of the European public bodies are developing the new nodes for the future network. The evolution of the CBA service will be upgrading to this new network, in the near future.

8 Dissemination of the best practices

The experiences of the consortium with the public sector and the adoption on cloud is an important outcome of the project and shall be shared with the public audience. The produced best practises are therefore publicly available as part of the project website[18].

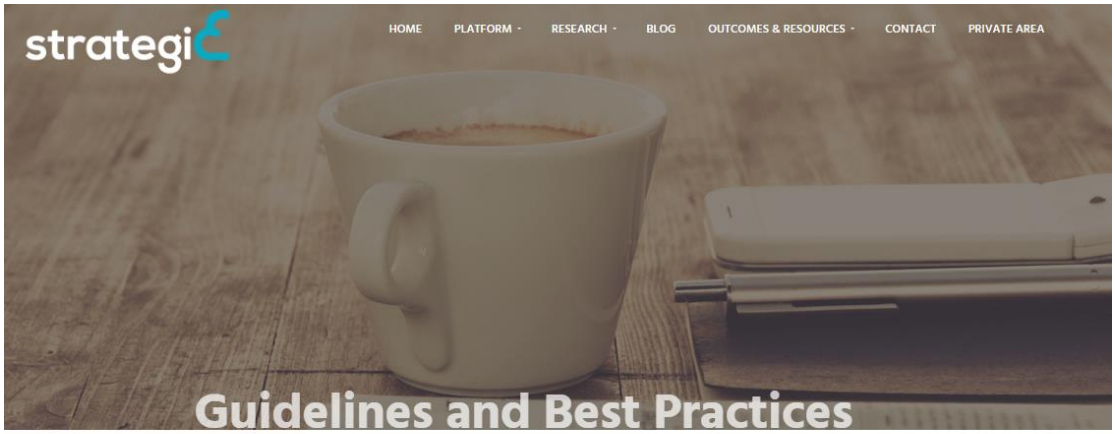

In this page you can find parts of our experience with the cloud adoption by public bodies.

#	Title
1	Cloud Migration for Public Bodies: Benefits and Obstacles
2	Selection of appropriate cloud computing service model
3	Selection of Appropriate Cloud Computing Deployment Model
4	Selection of Cloud Provider and the Benefits of Cloud Brokerage
5	Preparation and Design of a Private OpenStack Installation
6a	Installation of OpenStack IaaS Using MaaS and Juju
6b	Mirantis based installation of OpenStack IaaS
7	Protection of Applications and Data on Cloud Environments Using STRATEGIC

Figure 8: Best Practices in the project website

We also consider the role of the best practises sharing important in order to attract more audience to the project website through the SEO and the inbound marketing approach that has been described in deliverable D8.3[9]. Posts in social media based on the best practices are used as well in order to promote STRATEGIC to a wider audience.

[18] http://strategic-project.eu/guidelines-and-best-practices/

9 Conclusions

Based on the research and development that was done on STRATEGIC, and also based on the piloting and evaluation of the created platform and tools, the collected experience has been used for the elicitation of best practices and guidelines. In the chapters 2-7 the actual text of the best practices can be found organized under similar topics.

This material can help on promoting and boosting cloud adoption in public sector and it mainly focuses on stakeholders on governmental and public sector organizations. As the topics covered include the whole frame of cloud adoption, from selection of proper cloud model and deployment type, to installation experiences and usage of external tools, we consider that the best practices and guidelines are also useful for cloud providers, solution providers, ISVs, and actually anyone interested to adopt cloud.

Overall fourteen best practices and guidelines have been created, exceeding the initial goal imposed. The produced material covers all aspects that a public body moving to the cloud will have to be aware, and the same time illustrated the added value of STRATEGIC, either by the Stork and SEMIRAMIS integrations or with the STRATEGIC Service Store usage. These best practices are public available, through the project website and we believe that helpful to mainly to public organizations (i.e. public bodies), but also to other related stakeholders (cloud application developers, cloud providers).

10 References

[1] STRATEGIC Annex I - "Description of Work", 2014

[2] D2.3-STRATEGIC Framework Architecture and Technical Specifications

[3] D4.3a- Trust and SecurityComponents

[4] D4.3b- Trust and SecurityComponents

[5] SEMIRAMIS project: http://semiramis-cip.atosresearch.eu/

[6] STORK project: https://www.eid-stork.eu/

[7] EU eIDAS Regulation: http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.257.01.0073.01.ENG

[8] eSENS project: https://www.esens.eu/

[9] D8.3 Sustainability, Business Marketing and Financial Plans