Design Principles for Smart Space Management

Sven van der Meer, Brendan Jennings, Keara Barrett, Ray Carroll
Telecommunications Software and Systems Group (TSSG)
Waterford Institute of Technology (WIT)
vdmeer@tssg.org, bjennings@tssg.org, kbarrett@tssg.org, rcarroll@tssg.org

Abstract

Today, previously separated networks converge. Tomorrow, this process is going to challenge the established design principles of service frameworks. The new challenges for management frameworks are approaches that address composition, scalability, reliability and robustness as well as autonomous self-adaptation of services and resources. Research in these areas should result in software systems for management, control and use of fully distributed system.

This paper aims to start the discussion for a roadmap for the development of design principles that allow effective and efficient management of smart spaces. The starting point is the M-Zones model, which describes the relationships between smart spaces, smart services, managed zones and actors within this environment. The vision in short is to develop a management framework that allows for autonomous, self-adapting and self-managed services. Services are autonomous by means of interactions with the user, with resources and with other services. Services are self-adapting realizing the necessary modifications to interact with evolving, heterogeneous environments. And services are self-managed by means of configuration, scalabilit, fault management and security.

1. Introduction

The primary goal is to provide an open and lightweight management framework that supports ubiquitous solutions for smart space management. The overall vision is to develop a management framework that allows for autonomous, self-adapting and self-managed services. Services are autonomous by means of interactions with the user, with resources and with other services.

The established design principles for management frameworks rely on sophisticated long-term experiences in specific application areas. As these application areas converge, the design principles must converge, too. However, the process of a 'simple' convergence is not sufficient to provide future service providers with an environment where services can still be deployed and operated respecting future business needs. And the 'simple' convergence is not sufficient to support a customer-driven market with services tailored to individual user needs.

We think that there is a growing need to investigate the future of management frameworks by starting with a universal vision, instead of starting with specific long-term experiences. This vision is further used as a basis for substantial research in a variety of areas, which results in new approaches and concepts for management frameworks. These approaches and concepts must be exploited in existing or emerging markets to prevent us from predicting a future that will never become reality.

2. Constraints

There are a number of constraints that need to be recognised when defining a management framework for smart spaces. Business models formulate requirements set by service providers and network operators. Distributed Systems and service platforms manage the distribution of service intelligence, applications and appliances represent the available resources and finally, emerging technologies offer new paradigms.

Business Models – There is a need for creating a consensus between different players in shared business opportunities. Formal expressions can be found in business models, which identify the main players (stakeholders), their business role and interactions between them. Those models should be based on the projection that services evolve into an open, multi-provider market and information place.

Distribution of Service Intelligence – Today, the service intelligence is located at the core network or at the edges of the network. Tomorrow, the intelligence will be available at any place, wherever it is needed and wherever it is appropriate. The distribution of this intelligence might follow different constraints, such as the availability of resources, the integration into domains, or the effective cost for

offering a service. This means, services can be activated dynamically everywhere and in every form: centralised, distributed, at the edges, in the core network or on special network nodes.

Applications and Appliances – Applications and appliances need not to be present permanently nor pre-configured for a single use case. They can be available spontaneously, technically backed by mobile or ad-hoc networks. Furthermore, the infrastructure itself might become smart and flexible. Smart devices can become flexible enough to exchange information and to 'combine forces' to offer services by themselves. Wearables are another class of resources that demand for effective strategies to access and manage them.

Service-centric Network – From a service provider perspective, the actual technical environments will become more and more invisible. Instead of using a variety of technologies to access networks, there will be an open interface. This interface is represented by a service-centric network that offers the capabilities of networks towards the service provider and to the end-user. In the current understanding, a service-centric network can provide a unified network model integrating network and smart space technologies and allowing service aggregation towards service providers and end-users.

Emerging Technologies – New technologies reach the market continuously, some of them promise to change the nature of developing and deploying services. Peer-to-peer communication introduces a significant level of autonomy. Intelligent agents combine social and pro-active elements. Web services are the latest hype in the Internet to ease the access to services and to reuse existing components.

3. Smart Space Service Characteristics

Important for future management frameworks are the characteristics all services have in common. In general, all services are developed to act in one or more areas of concern: use, operate, control, administrate, and/or maintain. They are deployed in distributed environments. Services can be distributed in terms of time, location or any type of domain.

The basic characteristics of services are a certain degree of intelligence (goals, reasoning, learning) and a social ability (communication, co-operation). Additionally, a service does not report to other entities nor is it, in general, controlled by other entities. In more detail, we can identify four activities that support these characteristics.

- *Learn from previous events and adapt* behaviour. The ability to learn from previous events and to adapt behaviour covers two characteristics: autonomous behaviour and dynamic configuration. Both are needed for a once deployed service to adapt to changes in its environment.
- *Be aware of other services, applications, and resources.* Available services, applications, and resources must be identifiable and addressable. A service can then support a co-operative environment, passively, by being aware of all other available entities.
- *Advertise functionality to other services.* An active support of co-operative environments is realised by services that advertise their functionality to other services. All services that are aware of each other should share a common understanding on how to advertise functionality.
- *Interact with other services, applications, and resources.* Services must interact with applications and resources to fulfil their defined goals. Furthermore, services should interact with other services in order to organise themselves to work together for a certain task. This self-organisation can be pre-defined or ad-hoc in nature, based on pro-active service behaviour. Services have to develop a strong sense for locality, proximity, and community without losing the characteristic of being independent from their actual location.
- *Allow Service composition to provide value-added services.* Available control and management functionality in the network as well as legacy services can be used as a basis for service composition. Results of this activity are value-added services.

4. Roadmap

A management framework offers concepts and rules to derive concrete management architectures from it. The concepts cover service creation and deployment, while rules apply to run-times environments and service execution. For M-Zones, the starting point is a clarification of service concepts from different communities ('cultures') by means of:

- Business models;
- Terminology;
- Interoperability;
- Product Line Practices; and
- Components, Domain Engineering.

Following the main objectives, the management framework should integrate usage and management from the beginning to cover all areas of concern. Looking at interacting services and integrated management, the management framework must solve questions regarding information mapping and service management.

The three major requirements are interoperability, portability and scalability. Interoperability reflects

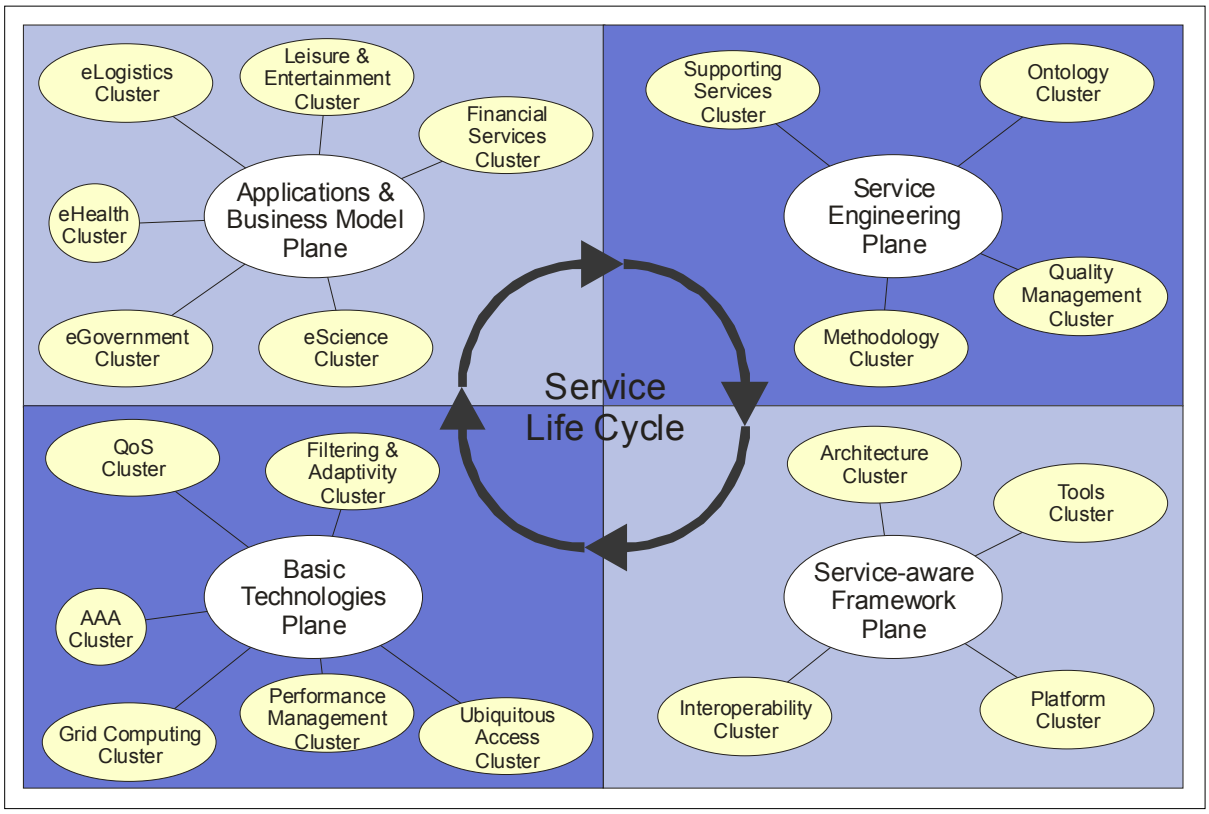

the existence of co-operative structures of services. Portability is introduced because the actual employed technologies are uncertain. Scalability should prepare the framework for its deployment over a large range application domains.

Other requirements, that have to be clearly defined, may be: technology independence, directory-enabled functionality and policy-based functionality. Furthermore, the framework is required to offer basic services like naming, registration, discovery and lookup, messaging and security.

There are two requirements that are crucial for the success of the work. The first is the specification of information that is exchanged among services, applications, and resources. This information must contain more than just signatures or syntactical descriptions. Instead, a common understanding (ontology) of descriptions of capabilities is needed. Such a description can range from technological information (e.g. display capabilities of a device) towards more abstract user requests (like 'I am hungry'). Profiles and databases, probably in form of directories, are needed to store such a description.

The second important requirement is the reliability of the communication between services. Since devices and users are mobile, and therefore unlikely to avail of permanent connections, mechanisms and formats for a reliable exchange of data should be identified.

5. Scientific Planes

5.1. Applications and Business Models

The goal on the application plane is to identify and to exploit new areas of business and to make existing processes more efficient. Research is conducted on business models that are embedded in and specific to application domains. Supporting services are needed for the operation of applications. Major objective of this plane is to phrase requirements for open development platforms to allow the harmonisation of industry needs and research activities.

There are several application domains feasible for the Network. All of them will be found in the core domain of open development platforms that M-Zones can support. Infotainment and gaming applications will drive tomorrow's revenue processes. Logistics and telematics are needed for improving the flow of goods and reducing unnecessary traffic. Banking and insurance services will fulfil citizen's demands of more direct personal money flow control. Health care/mobile health applications are often seen as one of the driving forces towards tomorrow's ubiquitous services. eGovernment is currently pursued all over Europe and will find its continuation in the mobile world as well. However, we concentrate on home, office and education environments.

The business views deal with generating revenues. Business models have to be found to adequately guide the future business processes, still to be defined

and implemented. The business view has to consider the different involved roles. Management support systems will be needed to monitor and control the money flow. The success factors of the services have to be researched, defined, implemented, and checked against.

These success factors particularly depend on the requirements from different areas. Most importantly, the prospective user's needs have to be taken into account. Political, social, ethnical, and market issues as well as telecommunication regulations play their roles in defining the services and applications. Not the least, the technology itself imposes limitations, restrictions and requirements.

To put successful applications into practice, several supporting services have to be provided. Inter-provider communication has to be ensured, dealing with roaming, accounting, charging, billing issues and enabling interoperability and combinability of systems. Personalisation support should be available to yield better results to the user.

5.2. Service Engineering

A sound engineering process in needed for the effective and efficient realisation of software. Based on a methodology that defines the process of service engineering, standard-based ontologies are built. The engineering process contains quality management mechanisms, which ensure a high standard of the realised solutions.

Methodologies, that guide the service creation, deployment, and runtime support process, have to consider many aspects. Applicable concepts must be found, which enable the modelling of new services. Therefore, an adequate notation must exist, that allows to formulate the process and component integration based on the methodology. Reusability of components and services must be taken into account and an engineering of requirements has to take place.

Part of the service engineering process is to identify, define, and share common ontologies. These ontologies define the semantic concept of the respective application domain. Meta-ontologies will be needed to formulate the ontologies themselves. For interoperability, this process has to consider standards, as well as itself must set common standards.

The service engineering process must ensure the quality of the product as well. Standards must exist against which the quality can be measured. Security has to be ensured and measurements have to be taken to validate the output against the requirements. Simulations will help to seamlessly introduce new services.

5.3. Service-aware Framework

The underlying technologies are integrated into a service-aware framework. The framework provides an architecture that is used to design an application. The realisation is based on a specific platform implementation of that architecture. Tools support the complete application life-cycle from planning until termination, including the access to information of the communication infrastructure. Middleware is used to interconnect the different parts of a distributed application.

The architecture provides the concepts, guidelines and mechanisms to smoothly integrate several concepts and approaches. Semantic mechanisms will be provided to deal with ontologies (from the plane above) and the Semantic Web concept. The resulting system must be scalable, which will be supported by infrastructure mechanisms on this plane. Efficient and easy services provisioning must be ensured. Existing services and solutions are to be integrated. Software agents play an important role, in integration, provisioning and generating added value.

The (WIT) platform is an implementation of the architecture. It will specifically include Web Services, recent XML technologies and provide libraries for easy deployment of the platform. Tools are needed to support the service creation, deployment and management process at runtime of the provided solution. Finally, a middleware ensures compatibility of communication in distributed systems. For interoperability, they must comply with standards.

5.4. Basic Technologies and Infrastructure

On the basic technologies plane, a multitude of different technologies that are needed to implement an application can be found. They are provided as libraries and tools for inclusion into the service-aware framework.

Different access technologies are emerging and their respective features must be employed; from traditional cable access, to 2G, 2.5G, and 3G mobile cellular networks, to next generation networks encompassing wireless LAN technologies,to emerging PAN networks, allowing access to a global network.

The network technologies not only include IP and public service telephony networks, but emerging concepts of peer-to-peer and ad-hoc networks have to be utilised, and management concepts must be put into place.

Handling of seamless mobility is still a challenge and an important research topic in today's pan-European community. The challenges include the multitude of devices with differing capabilities and the user's preferences and needs.

Services have to be enriched in order to satisfy the user. Instead of today's 'dumb' information provisioning, for better acceptance and to achieve higher satisfaction, they must be aware of the context of user, especially including the location of use and the preferences and requirements of the user. Further, community building must be supported, and habits of the user must be taken into account for service provisioning.

For managing and accessing the involved huge amounts of data, several aspects must be taken into account: a standard representation has to be used, scalability must be ensured, data mining technologies are involved, data must be aggregated and/or filtered. Grid computing might be employed.

Interface technology has to be researched. Today's representation in textual and graphical forms will be amended in the future by advanced audio and haptic interfaces.

6. Service-centric Environment

Beside the Service-centric Network, a management framework for smart spaces as envisioned by WIT will concentrate on the conception of a Service-centric Environment (SCE). The main objective here is the specification of a framework for integrated smart space services and their management. This framework is the core of the SCE and the foundation of any open development platform.

6.1. Modelling the SCE

This activity will be carried out from two main perspectives:

- *The integration of the Business view*. The objective is to integrate the business view into a functional view of the Service-centric Environment, referenced to the SCE layers as defined in a reference architecture. This functional view will support service life cycle implementation focussing on the delivery and operation of smart spaces. The common model will identify a set of smart services by integrating and abstracting network services already present in different access technology domains with different characteristics. This set of smart services will define a set of functional blocks to be designed and implemented during the project.
- *Service modelling*. The modelling approach is an important element of the M-Zones programme, when defining a Service-centric Environment that enables the creation and management of services from an abstract view. Service modelling activities are defined by the WIT M-Zones team as one of the major outputs of the SCE work as modelling will bring the necessary flexibility when creating, deploying and managing services.

Service modelling is based on the definition of meta-models and will benefit from the wide adoption of object modelling techniques and methodologies as well as a number of available tools. This approach fulfils the requirements defined by us for the SCE in a number of ways. At first, the SCE is based on some abstraction layers, first one is defined as Service-centric Networks layer. This approach also fully supports reusability of some of the SCE components for different services. More and more services should be based on well-defined components (e.g., COTS) as well as composed from other services. Composition and aggregation of services is highly dependent on service definition language, as well as service discovery processes. Service modelling also allows the possibility to integrate the mapping from the end-user perspective towards network and other resource level, e.g., configuration of services can start from the Service Level Agreement (SLA) defined between the customer and the service provider and then be mapped into specific routers' and smart space elements' configuration. This brings possibilities for self-adapted services.

One main challenge for M-Zones will be to model the network services based on different network technologies into a common model, enabling upper layers to be built independently of specific network technologies. This will provide the foundation of the SCE and will allow it to build and manage services with maximum flexibility, with respect to underlying network and device technologies. These open access methods allow retrieval of resources information as well as control resources, and they can not only be used by service execution processes but also by service management ones. This unique abstract layer is quite essential for the support of the evolution of services and to add value to services in the sense of seamless mobility, service adaptability, etc.

Let us consider the authentication and authorisation network services that are managed differently and implemented by different access domains (UMTS-FDD, TDD, 802.11a, 802.11b, Bluetooth). An M-Zones management framework will allow for an enhanced authentication service that will be able to integrate all authentication methods available in the underlying access networks. Specific APIs will publish this service to the application domain and the communication service layer will integrate authentication and authorisation information coming from different domains. Authentication and authorisation based on subscriber modules (i.e., UMTS / TDD) could be used also to integrate RADIUS / DIAMETER based authentication/authorisation (i.e., 802.11a; 802.11b), depending on the security requirements coming from the application layer. This simple example highlights the necessity of abstracting a unique user identity, an

integrated user profile and flexible / scalable access to network services from applications.

In the service modelling work, we include the modelling of resources used to build the service, such as the user profile or the terminal capabilities. Some work, such as CC / PP or 3GPP GUP, are ongoing and will be taken into consideration by M-Zones. The work done by the TMForum under the umbrella of NGOSS concerning the definition of Shared Information Model (SID) will also be taken into consideration. SID group provided first results whereby they provide definition of service and SLA. In today's situation, a number of ongoing initiatives exist, looking at various parts of the problem, with different perspectives. The aim of the WIT M-Zones team is to incorporate these different pieces into consolidated models enabling user requirements support. Our vision regarding modelling and description of services is a step-by-step view with a long term perspective. Indeed, we think some exploitable results are already there, however, as far as the description of services can include the semantic aspect of the service, or components of the service, we can envisage the link with the semantic web as defined by Tim Berners-Lee. This provides the long-term vision but also introduces some very interesting perspectives. Service modelling also concerns the decomposition of services into service components. These ones could be defined at different levels, such as business processes, or system components. It therefore brings lot of possibilities, such as service aggregation, based on services provided by a third party service provider.

6.2. Specification of an Open Framework

This activity will be focused on the definition of the architecture principles and concepts capable of supporting the functional and business requirements and how this can be applied to SCE reference architecture. This work will integrate ongoing work around enhanced distributed architectures and will evaluate state of the art technologies capable of fulfilling specific SCE components requirements, such as negotiation processes in distributed environments. This activity will also be a starting point for other activities, as it will provide the overall view of the Service Centric Environment, based on overall definition of the reference architecture. This will provide a clear separation between the different layers of the Service-centric Environment, the definition of the main components to be implemented in each layer and a definition of the interfaces between layers as well as with external systems. The specification of an Open Framework aims to build an architecture framework from state of the art technologies and architectural concepts. A number of initiatives need to be considered. An important output from the network will be to define an integrated view of these different technologies or standards, as well as results from other projects, based on a user driven approach. A number of areas are still not matured, so the project will contribute in the standardisation work through this specification work.

System architecture based on modelling is already addressed in a number of ways. From a methodology viewpoint, OMG provides the Model Driven Architecture (MDA) allowing to define the different aspects of component-based systems. This approach is promoted by a number of bodies, such as TMForum for NGOSS. Particularly, the definition of technology neutral models (PIM – Platform Independent Model) is a point of interest not only for system developers but also for business-oriented people. This should increase the possibility of the system architecture being re-used as the technology evolves; business individuals can easily comprehend the beneficial consequences. This is a corner stone in our approach.

Some new approaches are stimulating wide interest such as the Service Oriented Approach. This can be related to Web Services applications development. M-Zones needs to consider interest groups from the IT world as well as the telecommunications word. It is essential for a project such as M-Zones to validate such a technological shift and give the necessary feedback to the industry, therefore evaluating the real progresses that can be expected from any new technology.

Today, there is a high importance level surrounding web services technologies and other works aiming to support mobile Internet services. This brings an already high level of expectation from the industry as the shift from component-based to service-based approaches maps very well with the shift from a system oriented approach towards a business oriented one. This fits with the approach promoted by us, i.e., business driven. At the same time, the OMG will support the development of a software defined access network. This approach does not require additional efforts from classical access providers deploying new physical infrastructure. It simultaneously increases the usage of existing networks and optimises network resources.

Integrated Communication Services include solving end-user-programmable aspects of multi-network operation whose abstraction requires an integrated approach, such as criteria (programs) to react to network changes, single sign-on facilities, programming of asynchronous events, positioning, profiling, unified and secure context provisioning across multiple domains, feature discovery, etc. For the sake of simplicity, it is essential to have a single access point for the users to program adaptability for their services to react to context changes following a simple event-call-back model. The main purpose of this is to facilitate the interaction between a smart space, network and services so as to make it extremely simple for the user to create fully

adaptable services, despite the many resources accessed simultaneously and the inherent complexity of providing seamless mobility. Hence, it is the responsibility of this work package to provide applications with a consistent view of the status of available communication capabilities and with methods for automatic optimisation of resources, to ensure that the end user experience is consistent with the user's environment and profile. The main research challenge in this area is providing a really integrated and scalable solution for event-based service-network interaction in 4G communications.

6.3 Covering the complete Life-Cycle

Business Models. Otologies such as MDA or Service-oriented Approach start from an integration of business requirements. This activity will map defined business models and business requirements, which define the business view, into the system view. This mapping is quite strategic and could be defined as an essential element allowing the implementation of business requirements into systems. Business models bring the definition of the stakeholders as well as the contractual relationships between them. On the other side, the business modelling activity brings the business requirements, which can be extracted from business cases and scenarios. Based on these inputs, one important step consists in the definition of business processes. The correlation between business processes and the system components implementing these business processes is a major challenge, and can be viewed as the enabler for the definition of business driven systems. Here, we have lots of different working groups looking at this issue. There are new standards providing the global and long term vision, such as ebXML for B2B. There are some groups looking at more specific issues such as choreography or orchestration of web services (WSCI, BPEL4WS, BPML). Web services orchestration is about: 'providing an open, standards-based approach for connecting web services together to create higher-level business processes' (Chris Peltz, Hewlet Packard Co., Web Services orchestration). An interesting aspect of web services is interoperability through the service definition language (WSDL), which is technology neutral. Based on this, you can easily plug the different pieces together, based on the business logic. However, the vision is there, a lot of effort and interest from the industry, but real evaluation is still necessary. Moreover, the mapping between web services and business processes still needs to be refined. Another quite interesting approach of the problem comes from the telecommunications world. TMForum defines DENng as well as an integration architecture enabling the easy plug-in of new components. The latter is based on the definition of a system integration map, which defines logical business components. Components are defined through technological-neutral contract interface specification. NGOSS Technology Specific Architecture team have already defined a mapping of NGOSS Technology Neutral Architecture framework onto XML (and related technologies SOAP, UDDI).

Terminology. M-Zones will evaluate these different works and level of pertinence for integration into the SCE. Most of these works are still in progress; therefore M-Zones will be in position to contribute actively to some of them. The SCE will be defined from architecture principles and a set of building blocks. However, to ensure wide adoption of this framework, M-Zones will also provide methodology guidelines enabling developers and integrators to be compliant with SCE architecture principles.

6.3.1. Interoperability

Secure service discovery – will focus on protocols for service discovery both for the service instance and for the end-user. Options to be explored include: directory/registry based, multicast/paging-based, overlay/peer-to-peer-based, etc., and the utilisation of widely accepted standards as building blocks, like SLP, LDAP, DHCP, and W3C's Web Services (e.g. SOAP, WSDL, UDDI, WSFL). The activity also focuses on access control mechanisms for using and accessing public interfaces after discovery.

Reconfigurability – explores new forms of programming the behaviour of the communication services directly by the end-users (rather than by the service provider) under different circumstances (type of access, terminal, QoS, cost, etc), by enhancing today's very limited expressivity of web-based personal configuration interfaces with visual programming techniques including flow control operators, and special extensions in the service environment.

End-to-end service delivery – concentrates on the assured and secure delivery of end-user services to the user. When service centric networks offer network capabilities to users and services, these capabilities must be provided end-to-end. Capabilities such as a choice of network QoS must reflect the end-environment-to-end-environment network QoS, taking into account the different network domains and types that may exist between the service and the end user. Techniques to tackle these problems should be explored for service centric networks.

The definition of a consistent and efficient session model that incorporates various aspects related to application-level connections: host mobility, AAA credentials, media components, QoS and other low-level network parameters, billing information and location. Session control for multiple administrative domains and should address the issues that are particularly relevant to mobile applications.

The proposed solutions need to deal with discontinuity of SCN platform functionality offered through the session interface. For example the billing capability may be linked to having connectivity through a certain mobile network (billing through the operator of that network). This capability may be lost when roaming to another network. The session model and associating session APIs should support the service/application developer in a constructive way. Note that at the same time this model must provide an abstraction from functionality that varies widely over the different network technologies. The session model will interface with other service components in the federated platform. Session control will have to be incorporated in the overall federated platform architecture.

6.3.2. Product Line Management

Product line management defines a unified software development process, including a set of concepts, processes, methods and tools, which is important to enable open development platforms for the Service-centric Environment

6.3.3. Components and Domain Engineering

Service Adaptability to User Context – One of the objectives of M-Zones is to enable full exploitation of Communication Service Adaptability in all its possible meanings. This is the objective of the activity: providing a set of tools that enable and facilitate the adaptation of the service to: the current location of the user, to his/her profile, to the environment, to the capabilities of the terminal used, to the network

Session mobility – research on procedures to optimise the way communications services are used by the mobile user deciding upon choices. In the cases where, conditioned by the service, a handover to a different technology is needed, the component developed in this activity will provide seamless session mobility, as complementary to the solutions defined for terminal mobility. The outputs will include mechanisms to support the adaptation of the service to available network resources, the capability to change terminal characteristics during a session, roaming (accessing services from different domains), and methods to suspend/resume/reconfigure a session.

Integrated security –will feature single sign-on for users and users' applications to access network resources (gateways, bandwidth brokers, etc) from the service environment, by the definition of a unified architecture for federated security and authentication. This activity will also address user privacy and security management of context information and produce a complete specification and implementation of a scalable security architecture compatible with the requirements induced by the business model, the actors in the value chain, local privacy policies, interoperability and roaming of M-Zones. Special emphasis is placed on management of end-user privacy and profiles (which service provider party can access what and when of a particular end-user), to allow for end-user control on what information – such as location or status information – should be provided to which parties.