Data Centres Optimization for Energy-Efficient and EnvironmentalLy Friendly INternet

Funding scheme: Specific Targeted Research Projects – STREP

Co-funded by the European Commission within the Seventh Framework Programme

Project no. 609140

Strategic objective: FP7-SMARTCITIES-2013 (ICT-2013.6.2)

Start date of project: October 1th, 2013 (36 months duration)

Deliverable D2.2

DOLFIN requirements and system architecture

Due date:	30/09/2014
Submission date:	30/09/2014
Deliverable leader:	UCL
Editors list:	Alex Galis (UCL), Stuart Clayman (UCL)

Dissemination Level

☒	PU:	Public
☐	PP:	Restricted to other programme participants (including the Commission Services)
☐	RE:	Restricted to a group specified by the consortium (including the Commission Services)
☐	CO:	Confidential, only for members of the consortium (including the Commission Services)

List of Contributors

Participant	Contributor
Interoute	A. Kollaras, M. Biancani, F. Tirabasso
Siemens	S. Cosmin Nechifor
WIND	F. Pandini, A. Reale, M. R. Spada, H. Charlene
Synelixis	K. Pramataris, A. Voulkidis, T. Zahariadis
Nextworks	T. Zini, A. Gronchi
UCL	S. Clayman, L. Mamatas, A. Galis
GRNET	G. Goumas, D. Siakavaras, C. Tziouvaras, A. Sotiropoulos
I2CAT	I. Bueno, E. Escalona, I. Canyameres, J. Pons

Amendment History

Version	Date	Partners	Description/Comments
0.1	12/02/2014	UCL	Initial content list – Barcelona plenary meeting 11-12 February 2014
0.2	19/06/2014	UCL	Revised content list as per Brasov meeting 11-12 June 2014
0.3	24/06/2014	Wind	New contents: 4.3 (4.3.1 DC Monitor description, 4.3.2 DC monitor DB, 4.3.3 DC Monitor Backend) and 7.4 Wind TP
0.4	26/06/2014	Wind	Some editing corrections
0.5	07/07/2014	I2CAT	New content: 7.8. I2CAT Test-bed
0.6	07/07/2014	GRNET	New content: Section 4.4.1 (Performance Supervisor), Section 7.7 (GRNET TestBed)
0.7	23/07/2014	UCL	New context: Section 5.2 (NextWorks), Section 5.3 (NextWorks), Section 4.4 (I2CAT), Section 4.3.1 (I2CAT), Sections 4.2.1 - 4.4 (SYNE), Sections 4.4.1-4.4.3 (UCL), Section 4.4 (WIND)
0.8	27/08/2014	UCL	Next sections: the chilling system modelling section 4.2.1 (SYNE), section 4.3.3 customer profile (GRNET), Section 3 – Requirements Analysis (UCL), Appendix 1(Siemens/UCL)
0.81	28/08/2014	UCL	New sections: contributions to Abbreviations and Definitions.
0.82	2/09/2014	UCL	New sections: update of section 4.3.16 Energy Policy Maker (NWX); Section 6.2 UCL TestBed (UCL); Section 4.3.9 Orchestrator (UCL); Section 6.1 – Introduction to Testbeds (UCL)
0.83	2/09/2014	Wind	First revision of section 6; introduction of Figure/Tables Summary
0.84	3/09/2014	Wind	Final revision of section 6
0.85	3/09/2014	UCL	New sections: Section 4.1 Introduction to the DOLFIN architecture (UCL); Abbreviation and Definitions (UCL); Section 1 – Introduction and Approach (UCL)
0.9	12/09/2014	SYNE	Revision of Section 4
0.91	14/09/2014	UCL	Revision of Section 3 & Appendix I and II.
0.92	14/09/2014	UCL	Revision of Section 4 (add revised /new DC functionality). Consolidate Section 5 – energy system description - content into section 4.
0.93	14/09/2014	UCL	Consolidate introduction & foreword into one section – new section 1; general editorial changes (format & etc)
0.94	14/09/2014	UCL	Revised content list: Ch1 - Introduction, Context and Report Structure; Ch2 - Functional and Non-Functional Requirements; Ch3 - Overall System Architecture Description; Ch4 - Testbeds in the DOLFIN project; Ch5 – Conclusions; Ch6 – References; Ch7 - Abbreviations and Definitions; Ch8 – Appendix I: Requirements Analysis and Ranking; Ch9 – Appendix II: Requirements List
0.10	14/09/2014	UCL	Added conclusions; general editorial changes
0.90	21/09/2014	UCL	Accommodate final reviews & consolidate all chapters: SYN – section 2.2 'Metric and model of Energy consumption', SIE – section 2.1 "DC Energy Functional and Non-Functional Requirements", i2CAT & GRNET – section 4 "Testbeds", NXW – comments on all sections, SYN & UCL – section 3 "Overall System Architecture Description", UCL- comments on all chapters
1.00	24/09/2014	UCL	Consolidate final review comments/contributions from WIND, NWX, IRT, GRNET, SYNE, SIE, i2CAT, UCL. General editorial actions (language & formats).
2.00	27/09/2014	UCL	Added general editorial comments from WIND, NWX, SUNE, UCL
FF	30/09/2014	UCL	Final Version

Table of Contents

Figures Summary

Tables Summary

Executive Summary

This document is the deliverable D2.2 “DOLFIN requirements and system architecture” of Work Package 2” within the FP7-SMARTCITIES-2013 DOLFIN project.

DOLFIN aims to significantly contribute towards improving the energy efficiency of Data Centres and the stabilization of Smart Grids, across networks of Data Centres and Smart Grids, following a holistic approach. DOLFIN will model, monitor, and measure energy consumption and enable seamless, autonomic migration of Virtual Machines (VMs) between servers of the same DC or across a group of energy-conscious, synergetic DCs. The goal is to:

- optimize the overall energy consumption by dynamically changing the percentage of active versus stand-by servers and the load per active server in a DC, and
- stabilize the Smart Grid energy distribution, under peak load and increased demand, by dynamically changing the energy consumption/ production requirements of the local DCs.

The main purpose of this deliverable can be summarized as the presentation of the detailed description of a system architecture and of the main requirements for scenarios and use-cases related to the DOLFIN project. Such a description has been achieved via a workflow that is outlined in the following:

- Identification of the main system energy efficiency requirements based on scenarios, use-cases, and QFD method elaborated in the deliverable D2.1.
- Analysis of requirements in terms of motivation and impact on the DOLFIN architecture, as well as in terms of priority for importance and impact on the realisation of DC business goals and project outcomes.
- Identification of the main Data Centre Energy Metrics and KPIs relevant for the DOLFIN project.
- Identification and description of the DOLFIN architecture and its main functional blocks targeted at future DCs with new key characteristics:
    - SLAs, Orchestration, and Energy control loops per DC and/or group of DCs.
    - Optimisation of DC energy consumption, by dynamically changing the percentage of active/standby servers and load per server.
    - Optimisation of the cumulative energy consumption in a group of DCs, using policy-based VMs allocation.
    - Optimisation of the energy consumption at the DS / Smart City level and provide energy stabilization, by distributing VMs across a group of DCs, following the electricity demand-response approach.
    - Smart grid energy stabilisation, by dynamically changing the energy consumption/production requirements of DCs
- Testbeds description related to DOLFIN project.

1. Introduction & Context - Report Structure

DOLFIN's primary objective is to design, develop, and validate a Data Centre platform capable of monitoring the energy usage of the Data Centre and react accordingly for its efficient management. In this manner, DOLFIN will be able to address the criticalities relating to the growing energy needs and the consequent waste of resources, so that the Data Centre can still guarantee the required level of performance.

The DC optimization process requires addressing different elements jointly, very often connected to each other in a transversal way. This intrinsic characteristic of the "DC system" determines precise architectural choices in the design of the DOLFIN system, which is going to outline two main macro levels.

The DOLFIN approach is based on the following:

- A definition of the **Energy Consumption Optimisation Platform (eCOP)**. The eCOP is the DOLFIN core platform for energy consumption optimization at Data Centre level. It allows for continuous monitoring, energy benchmarking, dynamic control, and adaptive optimisation of the Data Centre infrastructure, including ICT and HVAC equipment.
- A definition of the **Energy-conscious Synergetic Data Centres module**. This provides a dynamic, service-effective, and energy-efficient allocation of demands, across a distributed network of co-operating Data Centres.
- A definition of the **Smart-Grid Energy stabilisation module**, which controls the inter-connection with the smart grid network, providing responses on the changing demands for energy.
- The modelling and negotiation of SLAs as a function of the energy requirement. Online SLA negotiation, taking into account the dimension of the quality of offered services and the energy expected to be consumed, while providing a certain level of service quality, together with the respective cost.

In addition to the characteristics listed above, it is appropriate to reiterate the concept on which the "formula" of optimization proposed by DOLFIN is based on. The main key concept is the workload distribution in a network of Data Centres, managed by rules consistent with the objectives proposed by the project. This aspect imposes the need of a synergistic dialogue between the Data Center entities in this network. DOLFIN acts not only at the individual DC level, but it also defines an aggregation platform between Synergetic Data Centres and the required tools that will permit the integration between the involved actors.

As a consequence of the above, DOLFIN defines an eco-system which has the following high level requirements:

- Integration with elements which are external to DOLFIN, but closely involved in the optimization process (e.g. Smart Grids system and Customers).
- Integration of the Data Centre within the network, as defined in DOLFIN Synergetic Data Centres.
- Integration with entities and devices within a single DC.

Further key components that will be analysed and developed by the project include:

- The **ICT Performance and Energy Supervisor**. This component is responsible for maintaining an easily searchable and well-organised database with information regarding the performance and energy-consumption metrics of the underlying ICT infrastructure. It will actively collect such information at regular time intervals, or, alternatively, utilise any push-technology capabilities supported by the ICT infrastructure. In order to increase inter-operability and usability, the interface for retrieving data should be a standardised one, like the Data Centre Management System interface. In the case of reliance on some proprietary interfaces, this can be addressed through the use of an adaptation block, which will convert any proprietary interfaces via one common interface, based on emerging standards, such as Cisco's EnergyWise or IETF's Eman.
- The **Energy Efficiency Policy Maker and Actuator**. This is the central point of making decisions and enforcing them within DOLFIN. This component allows the definition of energy policies, either manually, through the intervention of an authorised operator, or automatically, based on the information and criteria that are available. This component also maintains a set of energy-related actions that can be applied on ICT devices, while enforcing the energy policies. It is responsible for activating or deactivating energy policies, based on the criteria defined or the input received from other DOLFIN components.
- The **Workload and VM Manager**. This is a specialised component that controls the process of workload and VM migration from one Data Centre to another. It is responsible for maintaining all the information related to migrations done and their status.
- The **Smart Grid Controller**. This component is responsible for interacting with the Smart Grid network. It will maintain information regarding the time-varying energy consumption of the Data Centre and convey that to the Smart Grid. It will also process and respond to requests coming from the smart grid regarding the energy-consumption targets of the Data Centre. It will actually define and utilise the energy demand/response interface.
- The **SLA Renegotiation Controller**. This component is responsible for maintaining SLA information and for negotiating it with the customers. The objective of this component is to automatically negotiate and agree with a customer to lower the SLA performance parameters in favour of decreased service prices and less energy consumption. It will define, provide, and utilise the SLA renegotiation interface.

DOLFIN identified three primary interworking interfaces that will have to be implemented and exploited by the project and represent the interactions with the external systems that allow for monitoring and enforcement of energy parameters. More specifically, the DOLFIN interfaces are:

- The **Energy demand/response interface**. This is the interface between the Smart Grid network and DOLFIN, and it allows the communication between the two systems. This interface will be utilised to transport electricity requests from the smart grid, provide the necessary responses from the Data Centre, and convey the energy operating status of the Data Centre.
- The **SLA renegotiation interface**. This interface is established between the Data Centre and its customers. It allows for the dynamic renegotiation of the parameters that define the SLA agreement between the two parties. The SLA considers not only the performance indicators of the services provided by the Data Centre, but also correlates them with the expected level of energy consumption that would guarantee the performance indicators. The correlation also considers the cost of service provisioning, which allows the SLA renegotiation to have an effect both on the service cost and the expected power consumption.
- The **Workload and VM migration interface**. This interface is maintained between synergetic Data Centres that co-operate in such a way that allow the transferring of workload and VMs from one

Data Centre to the others. Such a migration process allows for the consolidation of SLA requirements under more energy efficient computing conditions.

The envisaged operational scenarios where the DOLFIN project would add significant value are:

1. **Absolute energy reduction**: this class comprises the entire Business Scenario related to the reduction of the energy consumed by the Data Centres (given that equal services will still be provided). This reduction can be achieved in different ways, but the end result is straightforward: less energy consumed means cost reduction, with the possibility for the DC Operator to benefit from energy-efficiency incentives, and other indirect advantages.
2. **Power consumption manipulation to achieve energy stabilisation**: in this Business Scenario, the goal is not necessarily to reduce the absolute consumption of energy, but to better distribute energy both geographically and temporally, in order to achieve efficiency at a higher level, than that of the entire energy system. The benefits for the DC Operators in this situation usually derive from "demand-response" agreements with the electrical power utility.
3. **Improve the quality of the energy mix towards environmental "friendliness"**: this class includes all the Business Scenarios that focus on choosing a given "energy mix" to power different services. This ability can be exploited to foster the use of renewable energy sources, e.g. by letting the final customer select a "green energy mix" to power his/her services. The benefits for the DC Operators can be: to provide new "green" offers to the market, to benefit from national or European incentives, and to improve their brand reputation.

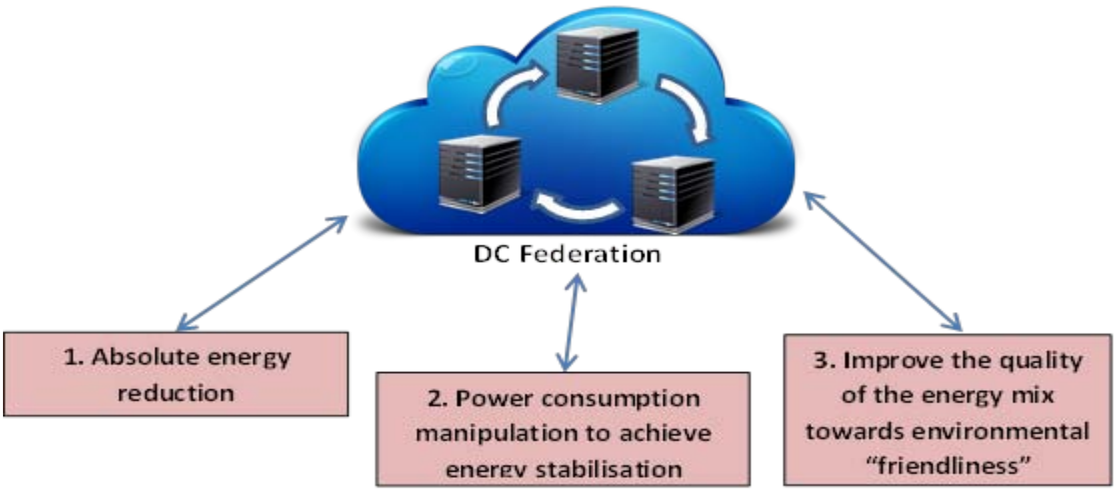

Figure 1-1: Business Scenarios main classes

This document is structure as follows:

- Chapter 2 presents functional and non-functional requirements and their analysis in terms of motivation and impact on DOLFIN architecture, in terms of the main Data Centre Energy Metrics and KPI relevant for DOLFIN project, as well as in terms of priority for importance and impact on the realisation of DC business goals and project outcomes
- Chapter 3 presents the DOLFIN architecture and the detailed description of the main functional DC and cross-DCs components of the DOLFIN system which are targeted at future DCs with new key energy efficiency characteristics.
- Chapter 4 presents the testing environment for DC inter-operability and also the DC testbed infrastructures where the DOLFIN architecture will be integrated and experimented with.
- Chapter 5 presents a summary and conclusions of this report.
- Chapter 6 presents the list of references used in this report.
- Chapter 7 presents the definitions and abbreviations used in this report.
- Chapter 8 – Appendix I presents the aggregated energy efficiency DC requirements identified and detailed Quality Functional Development (QFD) analysis and ranking of requirements
- Chapter 9 – Appendix II presents the list of aggregated requirements and their relationships with row requirements identified using the use cases defined in deliverable D2.1

2. DC Energy Functional and Non-Functional Requirements

A requirement can be defined as a statement that identifies a capability or function that is needed by a system in order to satisfy its customer's needs. In the context of DOLFIN, a functional requirement is meant to describe a function or a feature of a DC system, or its components, capable of solving a certain energy optimisation problem or replying to a certain Energy consumption need/request from key actors. It presents a complete description of how a specific system will function as far as energy management is concerned, capturing every aspect of how it should work before it is built, including information handling, computation handling, storage handling and connectivity handling.

Non-functional requirements are meant as implementation attributes and artefacts that a specific DC system must have. Examples of categories of non-functional requirements are:

- ***Usability:*** it describes the ease with which a system performing certain functions or features can be adopted and used.
- ***Reliability:*** it describes the degree to which a system must work. Specifications for reliability typically refer to stability, availability, accuracy, and maximum acceptable bugs.
- ***Performance:*** it describes the degree of performances of the system (according to certain predefined metrics, e.g. convergence time).
- ***Supportability:*** it refers to a system's ability to be easily modified or maintained to accommodate usage in typical situations and change scenarios. For instance, how easy should it be to add new blocks and/or subsystems to the support framework.
- ***Security:*** it refers to the ability to prevent and/or forbid access to a system by unauthorized parties.
- ***Safety:*** It refers to conditions of being protected against different types and the consequences of failure, error harm or any other event, which could be considered non-desirable.
- ***Resilience:*** it refers to the ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operation
- ***Compliance:*** it refers to the conformance to a rule, such as a specification, policy, standard or regulatory.
- ***Extensibility:*** it refers to the ability to extend a system and the level of effort and complexity required to realize an extension. Extensions can be through the addition of new functionality, new characteristics or through modification of existing functionality/characteristics, while minimizing impact to existing system functions.
- ***Inter-operability:*** it refers to the ability of diverse (sub)systems to work together (inter-operate).
- ***Operability:*** it refers to the ability to keep a system in a safe and reliable functioning condition, according to pre-defined operational requirements.
- ***Privacy:*** it refers the ability of system or actor to seclude itself or information about itself and thereby reveal itself selectively.
- ***Scalability:*** it refers to the ability of a system to handle growing amounts of work or usage in a graceful manner and its ability to be enlarged to accommodate that growth.

In this chapter, we are going to define the DOLFIN System Energy requirements based on the use cases defined in deliverable D2.1, analyse and prioritize them using the Quality Function Deployment (QFD) methodology.

2.1. Synthesis of the main DC Energy Functional and Non-Functional Requirements and Characteristics

2.1.1. DOLFIN DC Energy Requirements

For the priority use cases identified in the deliverable D2.1 (e.g. UC 5.2, UC 5.2.1, UC 5.2.2, UC 5.2.3, UC 5.3, UC 5.3.1, UC 5.3.2, UC 5.3.3, UC 5.4, UC 5.4.1, UC 5.4.2, UC 5.4.3, UC 5.5, UC 5.5.1, UC 5.5.2, UC 5.5.3, UC 5.6, UC 5.6.1, UC 5.6.2, UC 5.6.3, UC 5.7, UC 5.7.1, UC 5.7.2, UC 5.7.3, UC 5.8) 478 rows of functional requirements were identified in terms if interactions between key actors, functional roles, and links with the relevant components. Each row of the functional requirements were also assigned a specific name and numeric identity (e.g. R123). The requirements were compiled in a requirement database and analysis system at http://claydesk2.ee.ucl.ac.uk:8080/DOLFIN/requirement/index.

A similarity and consolidation exercise was performed on the 478 rows of functional requirements with the view of generating a higher-level aggregation of requirements with the following characteristics:

- ***Consistent and atomic:*** a requirement addresses only one function or feature. The requirement is atomic, i.e., it does not contain conjunctions.
- ***Complete:*** a requirement is fully defined in one place with no missing information.
- ***Dependable:*** a requirement does not contradict any other requirements and is fully consistent with all relevant references.
- ***Current:*** a requirement has not been made obsolete.
- ***Feasible:*** a requirement can be implemented and supported by the enabling technology
- ***Verifiable:*** the implementation of a design goal/objective can be determined through one of five possible methods: inspection, demonstration, test, trial or analysis

From this process, 39 aggregated requirements were identified, each requirement identified with a numeric identity (e.g. Q42), plus the name, description, technical actors involved, roles, use cases links and links with the rows of functional requirements. Theses aggregated requirements were also compiled in the requirement database at http://claydesk2.ee.ucl.ac.uk:8080/DOLFIN/requirement/index which could by grouped on different dimensions.

The description of the Aggregated Requirements are as follows:

Requirement ID & Name	Requirement Description
Q50 Ask for PostponeExecutionPolicy	The Energy Eff. Policy Maker and Actuator should be able issue a request of the Workload and VM Manager to shift computing load (postpone it for later) when such a need arises, e.g. when this load shifting could result in substantial cost reduction due to electricity price or it is considered to be profitable in the case of an end-user SLA renegotiation process, or such a request arrives from a DCO / Smart City operator.
Q51 Ask for RenegotiatioPolicy	The Energy Eff. Policy Maker and Actuator should be able issue a request of the Renegotiation Policy (e.g. renegotiation of Workload and VM Manager to shift computing load (postpone it for later) when such a need arises, when this load shifting could result in substantial cost reduction due to electricity price or it is considered to be profitable in the case of an end-user SLA renegotiation process, or such a request arrives from a DCO

	/ Smart City operator.
Q52 Ask for VMParamsAdjustment	In case a workload redistribution is necessitated, the Energy Eff. Policy Maker and Actuator should be able to ask from the Workload and VM Manager to adjust the VM characteristics of a certain set of VMs in order to meet the new DC configuration scheme, as generated by the Energy Eff. Policy Maker and Actuator.
Q37 DC2BInterface MgmntSubOptimalDetection	The management subsystem shall detect DC suboptimal states (i.e. high KPI) by utilizing current DC metrics and defined thresholds.
Q15 DC2DCControlInterface	The DC2DC subsystem shall implements CONTROL interfaces to coordinate the VM migration between DCs.
Q16 DC2DCDataInterface	The DC2DC subsystem shall implements DATA interfaces to perform the VM migration between DCs.
Q14 DC2DCInterface	The system shall be able to interface with other federated DCs.
Q53 DC2SmartGridInterface	In order for the Smart Grid Controller to be able to communicate with both the Smart Grid/EPs and the DOLFIN-enabled DCs, a proper interface serving the relevant requests/responses of such communication is needed.
Q54 DCHeatControl	In case the Smart Grid Controller retrieves a demand for a heat exchange state change from the DC part, the Energy Eff. Policy Maker and Actuator should be able to control the heat exchange towards the Smart City (turn it off or on).
Q13 DCManagement	The system shall include a management entity for orchestrating the whole process.
Q20 DCVMHandling	The PolicyEnforcement subsystem shall be able to manipulate VMs (migrate them from server to server, from DC to DC, shutdown, etc.)
Q56 DCVMHandling (VM params)	In accordance with Ask for VMParamsAdjustment" this requirement implies that the Workload and VM Manager should be able to fulfil a request for changing the parameters of a certain set of VMs".
Q62 Energy Provider Cost profile	The Aggregators should be able to request for the tariff schemes of a certain set of EPs, through the Smart Grid Controller.
Q23 Energy Provider Emission rate	The EPs shall be able to produce energy production statistics and forecasts
Q24 Energy Provider Requests Mngmnt	The EPs shall be able to forward requests and to consumers about status changes of the energy provisioning
Q57 HandleExceptions	When a customer denies a SLA renegotiation proposal or a synergetic DC denies cooperation, the DCO should be able to handle the rejection and consider alternative self-optimization actions
Q41 HVACParamsInterface	This interface provides information to the monitor subsystem concerning power dissipation and status of HVAC equipment.
Q19 HVACPowerControl	The PolicyEnforcement subsystem shall be able to manage non ICT devices power dissipation (i.e. control/shutdown HVAC equipment).
Q42 ICTParamsInterface	This interface provides information to the monitor subsystem concerning power dissipation and status of IT equipment.
Q18 ICTPowerControl	The PolicyEnforcement subsystem shall be able to manage ICT devices power dissipation (i.e. DVFS, ACPI, etc.)
Q58 LogRenegotiationSteps	The SLA Renegotiation Controller should be able to persist the SLA renegotiation steps upon happening.
Q69 LogReplanningSteps	This logs should support development for Monitoring, Analytics, Audit and Security. It can be implemented as a utility package for Management of Replanting of energy scheme.

Q63 MgmntCostDetection	The system shall be able to enforce different power dissipation policies (i.e. using DVFS and/or ACPI functions, control HVAC equipment and also through VM migration). The PolicyEnforcement subsystem shall utilize an interface to communicate with the management subsystem that should be common across all platforms.
Q70 MgmntCostsOpportunity	Represents an Analytics functional package used my Management Interfaces to evaluate the impact of potential actions taking in account history logs and current state. Main result is an decision support regarding the opportunity to search and/or accept a different cost scheme and it's time based impact on managed DC.
Q22 MgmntDOLFINResponses	The system shall be able to send feedback and notifications to the Eps.
Q21 MgmntEPsRequests	The system shall handle incoming EP requests and notifications about changes and trends in the energy provision service.
Q5 MgmntPolicyDecision	The management subsystem shall automatically identify the actions that could be taken to optimize the DC energy state.
Q6 MgmntSLARenegotiation	The management subsystem shall be able to (automatically) request the SLA renegotiation to better adapt the SLA KPI to current optimal assets.
Q4 MgmntSubOptimalDetection	[Duplicate of DC2BInterface MgmntSubOptimalDetection] The management subsystem shall detect DC suboptimal states (i.e. high KPI) by utilizing current DC metrics and defined thresholds.
Q9 MonitorInterface	The monitoring subsystem shall be able to present metrics-information to the management subsystem. This interface is considered to expose in a normalized form a number of metrics and their frequency. It could receive the metrics to retrieve, their frequency and response format (e.g. publish to a message broker). The system shall monitor the DC infrastructure: IT devices, power units, other DC facilities that might be involved in the process of energy optimization.
Q12 MonitorOpenDataFmt	The monitoring subsystem should utilize open data formats is available.
Q7 MonitorParams	The monitoring subsystem shall always monitor DC parameters. This requirements identify the necessity to make a formal collection for all parameters that have to be monitored to build KPI and other triggers levels used to evaluate the DC performance and schedule the appropriate policies' actions. The list of parameters to be published should be configurable in terms of frequency of measurements and way to deliver. Also, the access to this functions should be controlled.
Q11 MonitorResilience	The monitoring subsystem should be as much resilient as possible (i.e. implement heartbeat, watchdog functions).
Q66 MonitorSLA Requests	The SLA Renegotiation Controller should keep track of the SLA Renegotiation requests for historical reasons.
Q8 MonitorTranslate	The monitoring subsystem shall always translate parameters to metrics (KPI).
Q10 MonitorVariousData	The monitoring subsystem shall handle many kind of data and with different nature (i.e. performance information, energy consumption, power dissipation, VM status and deployment, semi-static parameters, etc). This function will complement in terms of Data the other peer Monitoring interfaces.
Q67 NegotiateSLAUpdate	When required to do so (e.g. by the MgmntSLARenegotiation requirement), the SLA Renegotiation Controller should be able to actually

	perform the renegotiation
Q73 NegotiateVMTransfer	Since DC to DC transfer may involve a multi-part balancing the way of what and how to transfer should be negotiated at VM level.
Q17 PolicyEnfoercementInterface	The system shall be able to enforce different power dissipation policies (i.e. using DVFS and/or ACPI functions, control HVAC equipment and also through VM migration). The PolicyEnforcement subsystem shall utilize an interface to communicate with the management subsystem that should be common across all platforms.

The relationships between Aggregated Requirements and the row requirements are presented in Appendix II.

2.1.2. Synthesis of the main DC energy requirements and characteristics as the basis of the DOLFIN architecture

Current and future data centres are comprised of diverse cloud management and autonomic functions. This section presents an analysis that motivates the design of the DOLFIN architectural framework which explicitly accommodates the energy management which is aimed at, namely:

- Improve capital and operational efficiencies for DC operators through the use of a common organization, automation, and operations of all energy functions across the different domains
- A migration from an ecosystem of separate energy management functions towards a coordinated arrangement of energy management functions as represented in Figure 2-1.

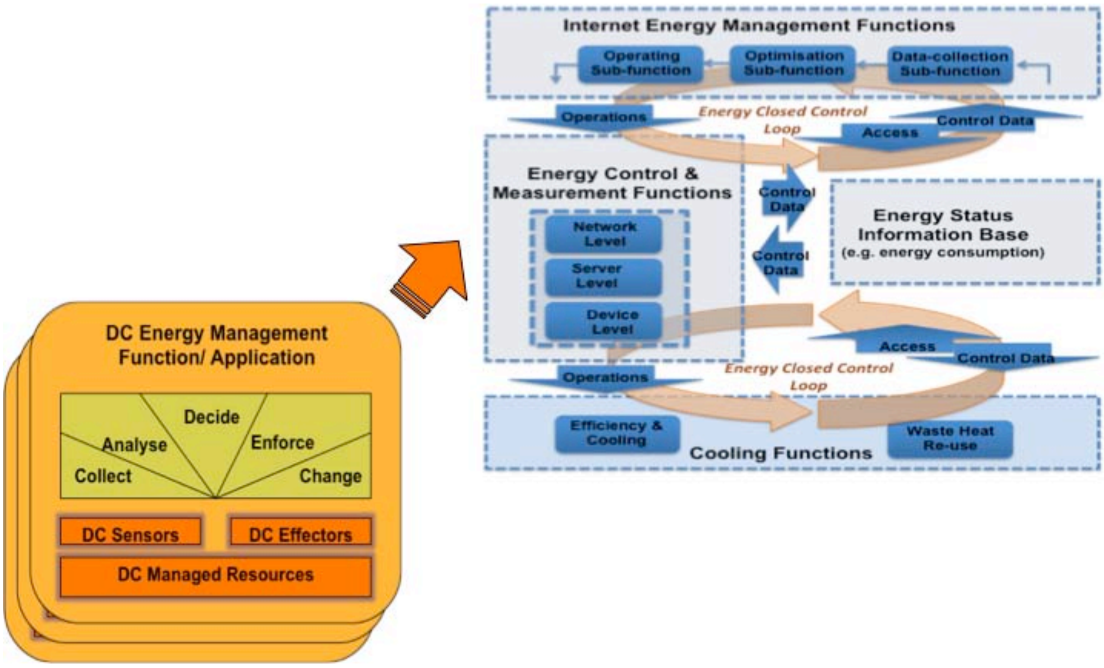

Figure 2-1: Migration from separate energy control loops to a coordinated arrangement of multiple DC energy control loops

The overall DOLFIN requirements list and design goals are derived from individual project partners' expertise, as well as the general vision and research directions for Future DC and Future Internet. The DOLFIN requirements list has two axes (see Figure 2-2):

- a "***bottom-up/row requirements***" expressed through 24 use-case problem specific requirements, addressing DC operators day-to-day problems identified in live DCs and on existing DC architectures. The use-cases as presented in deliverable D2.1 were prioritised via a QFD analysis. 478 row functional requirements were identified in terms of interactions between key actors, functional roles, and links with the relevant components.
- "***vertical requirements/aggregated requirements***" synonymous with high-level functions, new/updated functional blocks and inter-working interfaces in a DC. A similarity and consolidation exercise was performed on the 478 row functional requirements with a view to generating higher level aggregations of requirements with the following characteristics: Consistent and Atomic, Complete, Dependable, Current, and Verifiable. As such 39 aggregated requirements were identified.

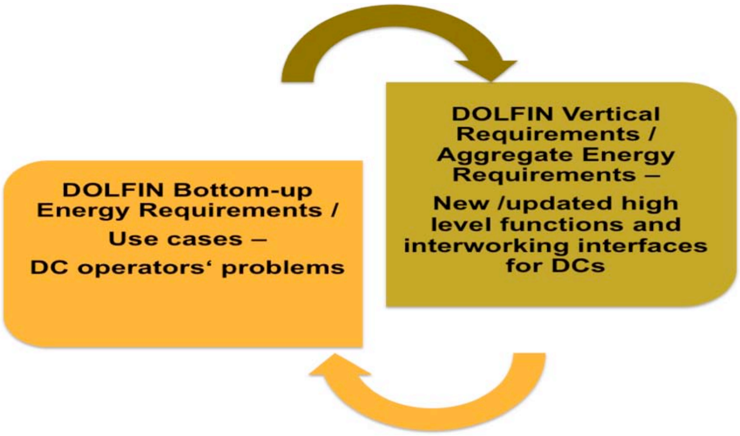

Figure 2-2: DOLFIN Axes of Requirements

The first approach of "bottom-up requirements" aims at addressing the set of requirements elicited for 24 use-cases defined and developed so far within deliverable D2.1 – business scenarios and use-case analysis. The second approach of "vertical requirements" aims at elaborating the expected new management functionality or updated management functionality in DCs. The requirements together as a set, and not necessarily per individual requirement, describe what distinguishes DOLFIN from earlier DC management technologies and what the DOLFIN project intends to design and deliver.

The following is a synthesis (see Figure 2-3) of the main DC energy requirements and characteristics which is the basis of the DOLFIN architecture.

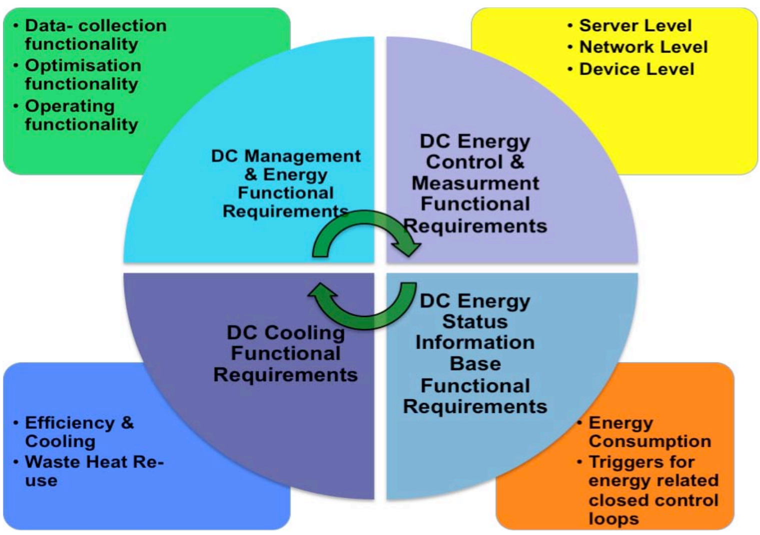

Figure 2-3: DOLFIN Requirements Synthesis

2.1.3. Ranking Requirements for importance and impact on the business goals and DOLFIN project expected outcomes

For each key actor in DOLFIN project (i.e. DC Operator initiator, DC Operator responder, Aggregator, End Customer, Utility Operator, Authority) we evaluated correlations between the aggregated requirements and business objectives/ DOLFIN project expected outcomes. In other words we evaluated how each aggregated requirements is contributing to solve/enable each business objective / expected project outcome as seen from each key actor.

The business objectives and expected outcomes used in the QFD analysis are as follows:

- ***Business Objectives - identified in deliverable D2.1:*** BO1.1- Reduce Server Costs, BO1.2- Reduce Infrastructure Costs, BO1.3- Reduce Power Costs, BO1.4- Reduce Other Costs, BO1.5- Reduce all operational costs, BO2- Increase the performance of offered services, BO3- Increase customer satisfaction, BO4- Improve the overall efficiency of the energy system, BO5.1-Shift processing power, BO5.2- Migrate processing power, BO6- Increase the availability /reliability of DC federation.
- ***DOLFIN Project expected outcomes -- identified in DoW:*** O1- Model, monitor & measure energy consumption; O2 - Energy optimisation & triggers for seamless and autonomic movement of VMs between servers of the same DC or across a group of energy-conscious synergetic DCs respecting/renegotiated SLAs; O3 - Optimize DC energy consumption, by dynamically changing the percentage of active/standby servers and load per server; O4 - Optimise the cumulative energy consumption in a group of DCs (policy-based VMs allocation); O5 - Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs, following the electricity demand-response approach);O6 - Smart grid energy stabilisation, by dynamically changing the energy consumption/production requirements of DCs

The use cases identified in deliverable D2.1 and used in the QFD analysis are as follows:

- **Use Cases - *identified in deliverable D2.1:*** UC1 - Energy efficient workload redistribution, UC2 - Multi tariffs from the Utility companies, UC3 - Multi tariffs from the Utility companies, UC4 - SLA Renegotiation with end customers, UC5 - Optimize benefits/incentives from national/European authorities, UC6 - Smart City.
- **Actors involved in DC operations in each use case- *identified in deliverable D2.1:*** DC Initiator, DC Responder, DC Aggregator, End Customer, Utility Operator, Authority (see section 7.1 – Definitions)

Correlation scores and importance factors for each problem were provided by an expert team from DOLFIN partners. These correlation evaluations were done according to each key actor Quality context for each business objectives and project outcomes and also for all actors (i.e. *Overall value*: correlation values related to the average of each previous value).

The aggregated requirements were ranked for importance and impact on the business goals and DOLFIN project expected outcome using the Quality Function Deployment (QFD) method and full results are presented in Appendix I. The following tables represent the summary of the prioritisation results.

The ordering of the aggregated requirements will be used in the detailed design and implantation of the DOLFIN functional components.

Requirements contributing most (✓) / least (✗) to realising Business Objectives if an Actor is considered in isolation are depicted in the following table.The rest of requirements are emerging as having average importance in realising business objectives as far as each Actor is concerned.

Requirement ID & Name	Actor DCO Initiator	Actor DCO Responder	Actor DCO Aggregator	Actor End Customer	Actor Utility Operator	Actor Authority	Combined All Actors
Q50 Ask for PostponeExecutionPolicy	✓	✓	✓	✓	✓	✓	✓
Q51 Ask for RenegotiatioPolicy	✓	✓	✓	✓	✓	✓	✓
Q52 Ask for VMParamsAdjustment		✓	✓				
Q37 DC2BInterface MgmntSubOptimalDetection		✓	✓				
Q15 DC2DCControlInterface		✓			?	?	
Q16 DC2DCDataInterface		?	?				
Q14 DC2DCInterface							
Q53 DC2SmartGridInterface		✓					
Q54 DCHeatControl			?				
Q13 DCManagement			?				
Q20 DCVMHandling	✓		?				
Q56 DCVMHandling (VM params)		?		?			
Q62 Energy Provider Cost profile							
Q23 Energy Provider Emission rate			?	?	?	?	
Q24 Energy Provider Requests Mngmnt	?						
Q57 HandleExceptions	?			?	?		
Q41 HVACParamsInterface	?	?	?				
Q19 HVACPowerControl	?		?				?
Q42 ICTParamsInterface			?	?	?	?	?
Q18 ICTPowerControl	?		?	?	?	?	?
Q58 LogRenegotiationSteps			?			?	
Q69 LogReplanningSteps			?				
Q63 MgmntCostDetection	?						?
Q70 MgmntCostsOpportunity			✓				
Q22 MgmntDOLFINResponses				✓		✓	✓
Q21 MgmntEPsRequests					✓		
Q5 MgmntPolicyDecision		?			✓		
Q6 MgmntSLARenegotiation			✓				
Q4 MgmntSubOptimalDetection	✓	✓					
Q9 MonitorInterface	✓			✓		✓	✓
Q12 MonitorOpenDataFmt	✓	?					
Q7 MonitorParams							

Q11 MonitorResilience	✓						
Q66 MonitorSLA Requests							
Q8 MonitorTranslate							
Q10 MonitorVariousData		⍰					
Q67 NegotiateSLAUpdate							
Q73 NegotiateVMTransfer		✓					
Q17 PolicyEnfoercementInterface	✓		✓	✓			

Table 1 - Requirements contributing most / least to realising Business Objectives

Requirements contributing most (✓) / least (⍰) to realising Project Outcomes if an Actor is considered in isolation are depicted in the following table. The rest of requirements are emerging as having average importance in realising project outcomes as far as each Actor is concerned.

Requirement ID & Name	Actor DCO Initiator	Actor DCO Responder	Actor DCO Aggregator	Actor End Customer	Actor Utility Operator	Actor Authority	Combined All Actors
Q50 Ask for PostponeExecutionPolicy	✓	✓		✓	✓	✓	✓
Q51 Ask for RenegotiatioPolicy	✓	✓	✓	✓	✓	✓	✓
Q52 Ask for VMParamsAdjustment		✓					
Q37 DC2BInterface MgmntSubOptimalDetection		✓	✓				
Q15 DC2DCControlInterface	✓	⍰				⍰	
Q16 DC2DCDataInterface							
Q14 DC2DCInterface							
Q53 DC2SmartGridInterface							
Q54 DCHeatControl							
Q13 DCManagement				✓			
Q20 DCVMHandling	⍰	⍰		⍰	⍰		
Q56 DCVMHandling (VM params)		⍰	⍰				
Q62 Energy Provider Cost profile	⍰		⍰		⍰		⍰
Q23 Energy Provider Emission rate	⍰	⍰	⍰	⍰	⍰	⍰	⍰
Q24 Energy Provider Requests Mngmnt	⍰	⍰	⍰				⍰
Q57 HandleExceptions		⍰		⍰			
Q41 HVACParamsInterface		⍰	⍰				
Q19 HVACPowerControl	⍰	⍰	⍰	⍰			⍰
Q42 ICTParamsInterface				⍰	⍰	⍰	⍰
Q18 ICTPowerControl	⍰			⍰		⍰	⍰
Q58 LogRenegotiationSteps				⍰	⍰	⍰	
Q69 LogReplanningSteps							
Q63 MgmntCostDetection		✓	✓				
Q70				✓			

MgmntCostsOpportunity							
Q22 MgmntDOLFINResponses	✓				✓	✓	
Q21 MgmntEPsRequests	✓				✓		✓
Q5 MgmntPolicyDecision	✓	✓		✓			
Q6 MgmntSLARenegotiation		✓	✓				
Q4 MgmntSubOptimalDetection		✓					
Q9 MonitorInterface	✓	✓	✓			✓	
Q12 MonitorOpenDataFmt		✓	✓				
Q7 MonitorParams		✓					
Q11 MonitorResilience	✓	✓					
Q66 MonitorSLA Requests	✓		✓		✓		
Q8 MonitorTranslate	✓		✓		✓		
Q10 MonitorVariousData							
Q67 NegotiateSLAUpdate			✓				
Q73 NegotiateVMTransfer			✓				
Q17 PolicyEnfoercementInterface	✓						✓

Table 2 - Requirements contributing most / least to realising Project Outcomes

The importance of aggregated requirements for realising business objectives based on combined all actors view points is presented in the following Figure 2-4.

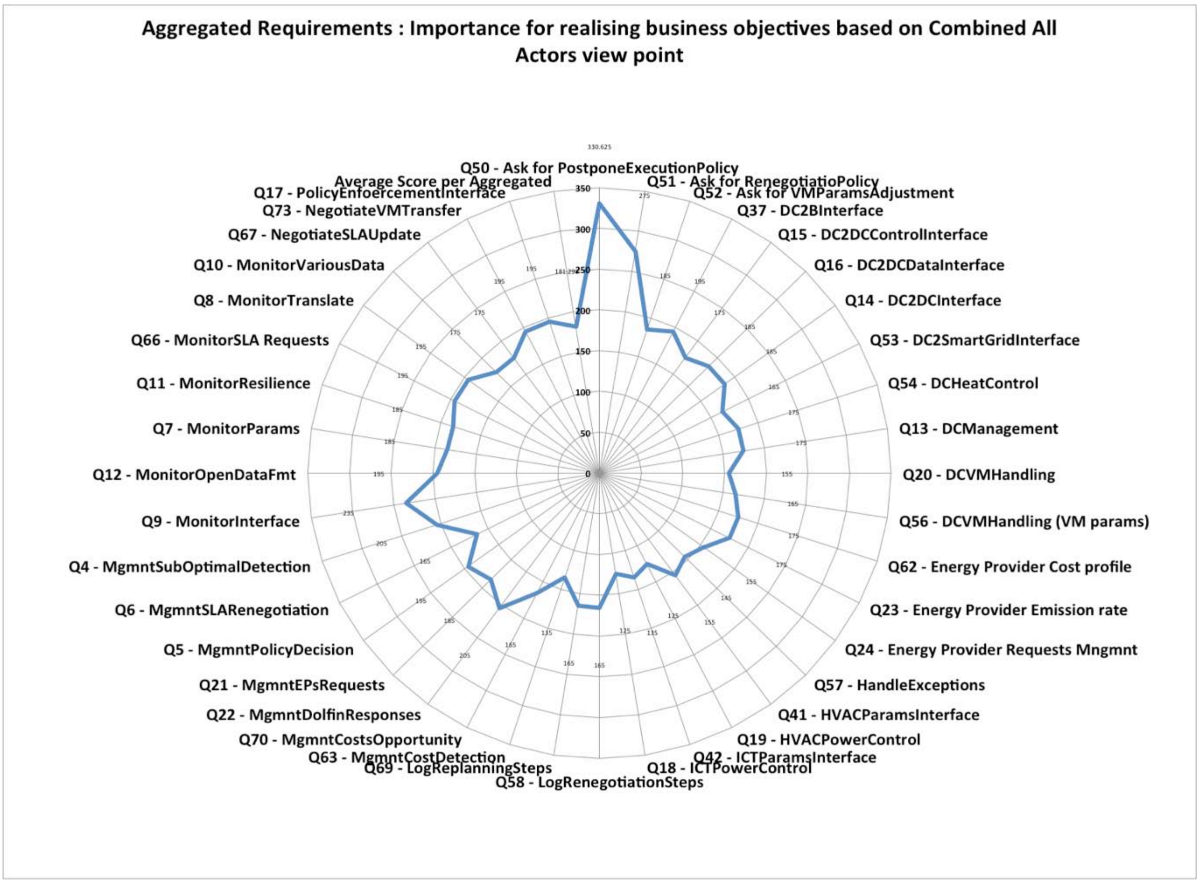

Figure 2-4: Importance for realising business objectives – combined all actors view points

The importance of aggregated requirements for realising project outcomes based on combined all actors view points is presented in the following Figure 2-5.

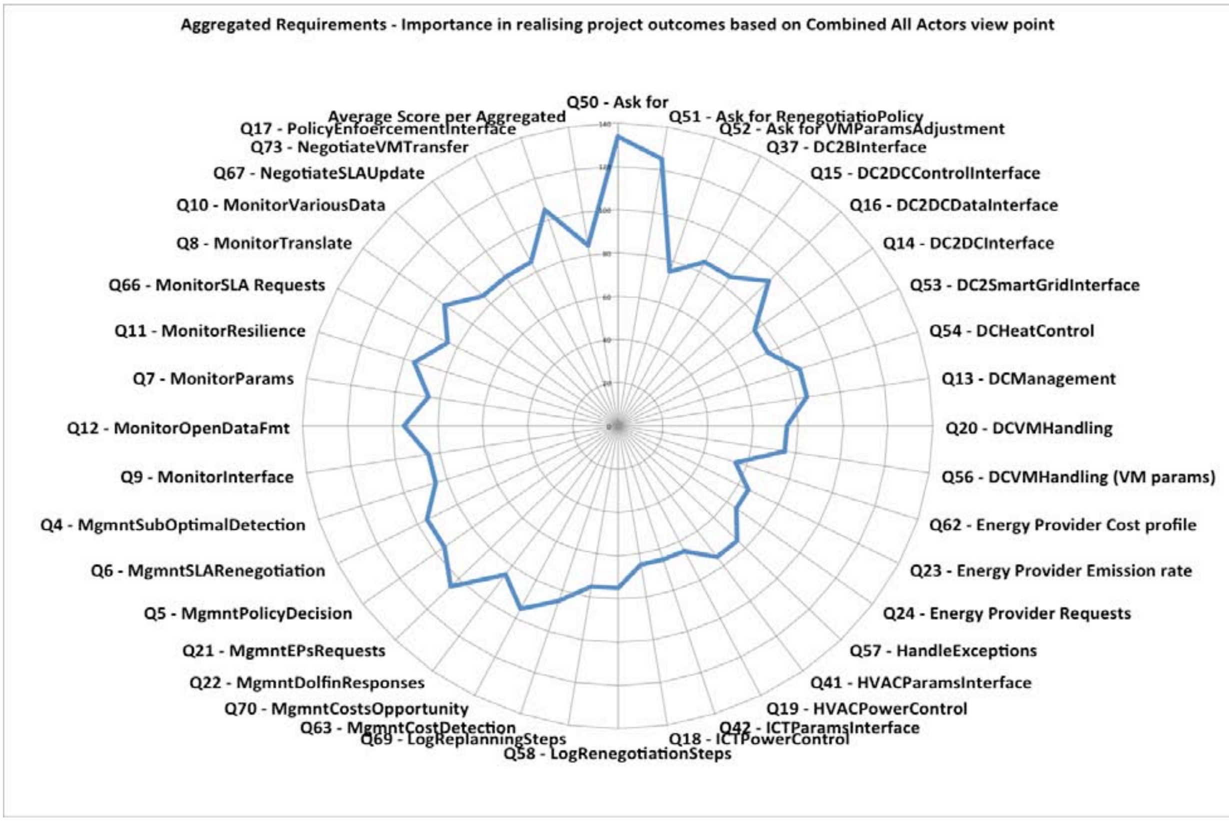

Figure 2-5: Importance for realising project outcomes – combined all actors view points

2.2. Metric and Model of Energy consumption

2.2.1. Review of existing models for Energy Consumption

Within a DC, energy is dissipated by both IT and non-IT components. Focusing on the IT side, 65% of the power is used by servers, 20% by storage, and 15% by network equipment including routers, switches etc Apart from that, running the cooling and power supplies within a data centre is responsible for a significant part of the total DC energy consumption. According to [3], servers are classified as the primary energy consuming elements, accounting for about 56% of the total DC consumption, with cooling consuming about 30%, and network elements being responsible for about 5% of the total DC energy consumption. As reported by Intel Labs [4] and shown in Figure 2-6, the main part of power consumed by a server is accounted for the CPU, followed by the memory and losses due to the power supply inefficiency.

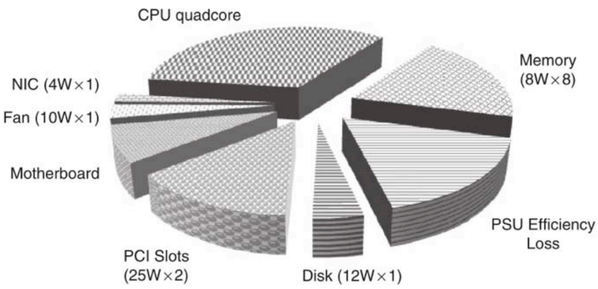

Figure 2-6: Power consumption of server components

It is worth mentioning that it has been shown that effective reduction in the power consumption at the IT side results in augmented energy savings at an overall level, since they cascade across the rest of the DC elements and supporting systems [4].

Power consumption in Complementary Metal-Oxide-Semiconductor (CMOS) circuits is the aggregate of static and dynamic power consumption [6]. The static power consumption, or leakage power, results from leakage currents flowing in any active circuit. This static power is mainly determined by the type of transistors and process technology being, thus, bound to hardware; any improvement on the static power consumption should be considered possible only under hardware change.

On the other hand, dynamic power consumption occurs due to circuit activity and is caused by short-circuit currents and switched capacitance. The former account for about 10%-15% of the total power consumption, while the major load consumer is the switched capacitance, calculated by

$$P_{dynamic} = \alpha C V^2 f$$

where α is a constant dependent on the circuit characteristics, *C* the physical capacitance, *V* the supply voltage and *f* the clock frequency [7].

Focusing on the CPU, which is the major energy-consuming part of the servers' components, *V* and *f* can be appropriately configured in lower values in order to derive reduced power consumption in CPU, employing dynamic voltage and frequency scaling (DVFS) techniques. It should be noted that *V* and *f* share an almost linear dependence and should not be considered as independently variable entities [8].

Identifying power consumption models for DCs refers to modelling the dynamic power consumption, in an attempt to predict the future systemic behaviour in terms of energy consumption, given known runtime system characteristics. Such models can also make use of real time power consumption data, acquired via smart metering and/or sensing/monitoring infrastructure, in order to fine-tune the model parameters and variables.

Server power consumption models can be classified into two categories, depending on the view of the components that are taken into account for the calculation of the total energy consumption [9]. In the first category, several components of the server are considered to contribute to the server power consumption, while in the second category a server is seen as a single power consuming element.

In the first case, the server power consumption is given by:

$$P_{server} = P_{hdd} + P_{mem} + P_{board} + P_{fan} + P_{proc} \quad (1)$$

where P_{server} is the total power consumed by the server, P_{hdd} stands for the power consumed by hard disk drives activities, P_{mem} for the power of the memory chips, P_{board} for the motherboard power consumption, P_{fan} for the power consumed by the cooling fans and P_{proc} for power consumption of the processors.

While this approach appears to be more precise since it considers every single element, it is generally hard/inefficient to measure the consuming behaviour of every single element.

In the second case, the server is seen as a black box, the power consumption of which is easier to be estimated provided tolerance of an acceptable error.

A strong relationship holds between CPU utilization and server power consumption. Specifically, a linear relationship is assumed between power consumption and the respective CPU utilization:

$$P_{server} = P_{idle} + (P_{max} - P_{idle}) \cdot u \tag{2}$$

where P_{server} is the estimated power consumption, P_{idle} is the power consumed by the server in idle state, P_{max} is the power consumed by the server at full load, and u is the CPU utilization at the current load [10]. Equation (2) is arguably the most commonly used formula in the literature modelling the power consumption of servers and will be most probably adopted by DOLFIN as well.

Similarly, a non-linear relationship between the two is presented in [9], as follows:

$$P_{server} = P_{idle} + (P_{max} - P_{idle}) \cdot (2u - u^r) \tag{3}$$

where *r* is an experimentally defined calibration parameter that minimizes the square error.

Although concrete evidence has not been given as to which of the two models is more accurate, linear models are generally preferable compared to the non-linear ones in order to minimize the optimize process time in real testbeds, since they are faster to complete and easier to implement.

The authors in [9] also propose a layered power consumption model, considering layers of power consumption added one on top of the others. In this model, the first layer corresponds to the "standby" power consumption of the server ($P_{standby}$), with the motherboard powered, but with an idle CPU. The second layer represents the power due to the hypervisor in idle mode ($P_{hypervisorIdle}$). In this state the server is started and the hypervisor is running without any VM. The next layer is added when VMs get started in idle mode (without client applications). The last layer represents the power consumption as a result of running client applications on the previously generated VMs. Then, the model is formulated as:

$$P_{server} = P_{standby} + P_{hypervisorIdle} + A \cdot N_{VM} + a \cdot U_{active} \tag{4}$$

where A an experimentally defined coefficient indicating the power consumption induced by a virtual machine when idle, N_{VM} is the number of virtual machines, U_{active} the active CPU usage, a is a coefficient binding power consumption and useful CPU usage. Evidently, this approach can be directly mapped to the one presented by the authors of [10]since both models assume some static and some dynamic energy consumption sources and link them using the relative CPU load of the server in hand.

It should be noted that under the "server as a black-box" assumption, (2)-(4) include the energy consumption due to both the CPU, RAM, Motherboard, PCI and IO devices. Similarly to the linear models of power consumption due to the CPU load of the servers, networking equipment has been also modelled to present a linear dependence on the rate of the ports of the equipment. Specifically, the authors of [10]argue that under a relative networking load (bandwidth used) u, the total power consumption P_{net} of a networking device (e.g. a switch) can be approximated as

$$P_{net} = P_{net}^{fixed} + u \cdot \left(\frac{P_{net}^{max} - P_{net}^{idle}}{B} \right) \tag{5}$$

where P_{net}^{max}, P_{net}^{idle} are the maximum and minimum power consumption of the networking device, respectively, and B is its total, maximum bandwidth. Interestingly, it has been shown that the static part of the networking consumption overpasses the dynamic one by much; changing the rates of all ports from 0% to 100% the power of a switch increases by less than 10% [11]- [12].

Regarding the cooling systems, several approaches exist, each one based on the chiller types in hand and the DC racks topological placement inside the DC rooms. A well-known, commonly adopted chiller energy model is the DOE 2.X [13] which is also used by numerous well known energy consumption simulators including EnergyPro [14]and DOE's eQuest [15]. DOE 2.X suggests that the energy consumption, under constant outside temperature, due to chiller activity is a function of the aggregate DC load, say u, and can be approximated using the following a formula of type

$$P_{chiller} = A \cdot u^2 + B \cdot u + C \quad (6)$$

where A, B and C are constants depending on the DC equipment characteristics. Since these parameters are dependent on the DC setup, proper training of the DOE model is required, in order to come up with a proper constants determination. It should be underlined that there exist tables predefining baseline consumption under usual/representative DC configurations [16]. Moreover, it has been identified that a similar approach exists for determining the coefficient of performance (COP) of a chiller in an average, industrial, production DC [17]:

$$COP(T_s) = 0.0068 \cdot T_s^2 + 0.0008 \cdot T_s + 0.458 \quad (7)$$

where T_s is the temperature of the cold air that that the chiller supplies to the room. Having determined the chillers COP, the expected power consumption of the chilling system of an average DC containing a total of N servers could be calculated as [18]:

$$P_{chiller}(T_s) = \frac{\sum_{i=1}^{N} P_i}{COP(T_s)} \quad (8)$$

Evidently, the aforementioned models refer to instantaneous power consumption of a single server/network/cooling element. At a DC level, all these factors contributing to the total DC energy consumption should be added and integrated over time in order to derive a proper prediction of the DC systemic behaviour under predefined load patters in the near future, based on the current DC configuration, state and dynamics. In short, considering the set of all the energy consuming elements of the DC, the energy consumption of the DC as a whole over time period T, starting at the time instance *t* could be expressed as

$$P_{DC}(t, t+T) = \sum_{servers} \int_{t}^{t+T} P_{server} + \sum_{net} \int_{t}^{t+T} P_{net} + \sum_{chillers} \int_{t}^{t+T} P_{chiller} \quad (9)$$

which cannot be further simplified since the distinct elements contributing to the DC energy consumption are not independent (e.g. the servers load, thus consumption, affects the consumption of the cooling equipment as well).

In order for the prediction of the DC power consumption to be best calculated, load prediction schemes could be also used, in order to come up with the load profile during the time period *T* examined through the model presented via (9). Examples of such models are presented in [18]. Taking for granted that a definite model for the total DC power consumption exists via (9), the power dissipation optimization problem could be modelled as a linear programming minimization problem, as the one presented by in [19]. As an alternative, the authors of [20] suggest that a proper reinforcement learning algorithm could be trained in order to come up with accurate predictions of the DC energy consumption behaviour under the model described by (9).

Similar to the goals of DOLFIN several EU projects have been, recently, investigating the problem of modelling the energy consumption of data centres. One of them, CoolEmAll [21], has proposed an approach to practically model the energy consumption of a DC using a component based-approach called Data Centre Efficiency Building Block (DEBB). In the context of CoolEmAll, DEBB is treated as the smallest element in the thermodynamic modelling process. Having modelled a single element, the whole DC energy consumption could be modelled as the scaled aggregate energy consumption of the sum of the DEBBs. DEBB modelling may be split into two steps, namely load dependent power models, where the distinct load dependent power models which only consider the load of the particular DEBB (DC component) as a variable, and combined load dependent data centre power models where the view of power consumption is aggregated into a more DC-catholic perspective, also considering other parameters affecting energy consumption such as the outside temperature.

2.2.2. Review of existing Energy DC Metrics

The rapid growth of the demand for cloud services has increased the need for high-performance services, indirectly inducing excess energy costs to the DC operators. Numerous ways to reduce energy consumption exist and relevant metrics and KPIs have been defined, in order to quantitatively assess their effect on the efficiency (performance- and energy-efficiency-wise). Several standardisation bodies and independent initiatives have been working on providing concrete sets of metrics that are meaningful and easy to implement in order to render comparison between DCs and DC-related operations comparable. Recently, the International Standards Organisation (ISO) in combination with the International Electrotechnical Commission (IEC) created a Joint Technical Committee Subcommittee (ISO/IEC JTC 1/SC 39) to develop and facilitate standards in the field of sustainable and energy efficient ICT services and infrastructures. The first Work Group (WG1) of ISO/IEC JTC 1/SC 39 is dedicated to resource-efficient DCs [22]. The scope of this WG is to develop a taxonomy targeting at DC resource efficiency, define meaningful and easy-to-implement KPIs for DCs, provide guidance for the creation of resource-efficient DCs and set up a standard for energy management in DCs. Currently, ISO/IEC JTC 1/SC 39 WG1 is working on the standardisation of three KPIs (namely PUE, REF and ITEE – to be defined in the following paragraphs) [23]. Similarly, the European Telecommunications Standards Institute (ETSI) recently published a specification under the name ETSI GS OEU 001 [24] to define global KPIs for DCs, along with measurement procedures and processes. ETSI GS OEU 001 discusses several well-known metrics such as PUE, REF and ERE, also combining them in an attempt to come up with a global KPI that would characterise the resource efficiency of a DC under a multi-criteria perspective. It is worth noting that although PUE and REF are defined both from ISO/IEC and ETSI, their definition is not entirely the same.

ITU-T Recommendation Y.3021 [25] describes the framework of energy saving for Future Networks (FNs), including DCs. It presents the need for energy saving of future networks themselves, and it reviews potential technologies. It identifies major functions and their cyclic interactions, analyses possible impacts of introducing the technologies, and itemises the high-level requirements for introducing the technologies. It was developed by the ITU-T Focus Group on Future [26], where UCL acted as vice chair. BSC DC Specialist Group [27] published its report on DC energy efficiency metrics [28] which analyses existing and proposed metrics to provide effective understanding and reporting of DC energy, mostly related to PUE and DCiE.

Following the lead of the main standardisation bodies, and acknowledging the need for faster derivation of concrete results when it comes to DC metrics, several independent initiatives have been initiated aiming at determining more KPIs to cover more DC operation aspects. Green IT [29] is an open initiative targeting at sustainable ICT that has already developed several DC KPIs, including PUE, ITEE, ITEU.

Recently, the European Commission coordinated 8 EC-funded projects to form a knowledge-exchange collaboration known as DC Cluster Collaboration, DOLFIN being part this initiative [30]. In the context of this initiative, several metrics, including the most common ones, are identified. The work programme of the DC Cluster Collaboration was split into distinct phases, subsequently identifying the already existing metrics and KPIs related to DC operation, then determining new metrics that might be of use for some DC operation aspects and, finally, define methodologies to accurately and consistently calculate those metrics. In the following table, the already existing DC-related KPIs are tabulated.

Energy/Power consumption						Energy produced locally	Heat recovered /Energy reused	Power Shifting	CO2 emissions	Performance	
IT	Cooling	UPS	Transformer	Lighting	Building					Economic	Applications
CADE	CADE	CADE	CADE	CADE							
PUE	PUE	PUE	PUE	PUE							
DCIE	DCIE	DCIE	DCIE	DCIE							
CPE	CPE	CPE	CPE	CPE							
DCeP	DCeP	DCeP	DCeP	DCeP							DCeP
THD	THD	THD	THD								
ERF	ERF	ERF	ERF	ERF			ERF				
ERE	ERE	ERE	ERE	ERE			ERE				
CEF	CEF	CEF	CEF						CEF		
GPUE	GPUE	GPUE	GPUE						GPUE		
TUE	TUE	TUE	TUE	TUE						TUE	TUE
DCU, ScE, DCcE, DCPD, DCD, Swap, Useful Work, ProductivityDC, TPS/Watt, DH-UR, SI-POM, H-POM, U_server, U_network, U_storage, ITUE											TPS/Watt
CUE, CEB, WUE									CUE, CEB, WUE		
CIUD, PAR4									PAR4		
DPPE	DPEE	DPEE	DPEE	DPPE	DPPE	DPPE	DPPE		DPPE		
DC FVER										DC FVER	
DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct	DCMM, Code of Conduct		
...	...	...	...	...	...	...	...	...	...	...	...

Table 3 - Existing DC resource efficiency metrics as identified by the DC Cluster initiative.

Next, the main, most well-known, already existing metrics will be analysed, followed by a short overview of the metrics that the EC DC Cluster Collaboration programme has delivered.

2.2.2.1. Existing metrics

Power Usage Effectiveness (PUE)

PUE is defined as the ratio of the total power drawn by a DC facility to the power used by the IT equipment in that facility:

$$PUE = \frac{Total\ DC\ Power}{IT\ Equipment\ Power} = \frac{P_{cooling} + P_{lost} + P_{lighting} + P_{IT}}{P_{IT}}$$

where$P_{cooling}$ is the power used by the entire cooling system of the DC, P_{lost} is the power lost in the power distribution system through line-losses and other infrastructure (e.g. UPS, PDU) inefficiencies, $P_{lighting}$ is the power used to light the DC and P_{IT} is the power used by the IT equipment (server, network, storage) of the DC in hand.

The total DC power is the sum of the power consumed by the IT and the non-IT equipment in a DC. The IT equipment power is the power drawn from the DC UPS by the equipment used to manage, process, store or route data within the DC. Measuring the IT power requires sub-metering of the rack distribution power, which is often incorporated in PDU equipment.

In addition to the default definition of PUE which refers to instantaneous power demand, several additional definitions of PUE have been defined to reflect efficiency based on different data measurement points, frequencies ad averaging periods. As a result, four distinct PUE categories have been identified by the Green Grid [31], namely PUE_0, PUE_1, PUE_2 and PUE_3, the definition of PUE_0 coinciding with the PUE definition already presented. The rest three PUE categories are defined as follows:

$$PUE_1 = \frac{E_{DC}(Wh)}{E_{UPS}(Wh)}$$

$$PUE_2 = \frac{E_{DC}(Wh)}{E_{PDU}(Wh)}$$

$$PUE_3 = \frac{E_{DC}(Wh)}{E_{IT}(Wh)}$$

where E_{DC} is the total energy drawn at the DC level, E_{UPS} is the energy drawn by the DC UPS system, E_{PDU} is the energy drawn by the PDU equipment and E_{IT} is the energy drawn by single IT equipment.

The primal difference between the definition of PUE_0 and the respective ones of PUE_{1-3} lies on the fact that the former refers to the maximum instantaneous power demand over a reference time period (usually a full year) whereas the latter refers to the aggregated power demand over the same time period. The distinction point among PUE_1, PUE_2 and PUE_3 is based on the metric measuring points, the latter being the DC UPSs (PUE_1), the PDU equipment (PUE_2) and single IT equipment (PUE_3).It is worth noting that proper combination of PUE_{1-3} can help DC operators identify the points of the DC where most energy is drawn and need, hence, to be optimized.

PUE has received broad industry adoption as an overall facility efficiency metric (the closer to 1 the better). DC PUE of 2.4 to 3 (and higher) were not uncommon, indicating that as much as twice the power consumed by the IT equipment was required for the supporting facilities.

Data Centre Infrastructure Efficiency (DCiE)

An alternative to the PUE is the Data Centre Infrastructure Efficiency (DCiE), defined as follows:

$$DCiE = \frac{IT\ Equipment\ Power}{Total\ DC\ Power} = \frac{1}{PUE_0}$$

DCiE is a more intuitive measure of the overall efficiency of a DC. This metric is similar to traditional efficiency measures and indicates the percentage of the total energy drawn by a facility that is used by the IT equipment. DCiE has been adopted as the key metric for infrastructure efficiency in the European Code of Conduct on DCs Energy Efficiency and is expected to gain wider adoption in the future.

Renewable Energy Factor (REF)

The Renewable Energy Factor (REF) is a unit-less metric indicating the greenness of a DC and is defined as the ratio of *local* renewable energy consumed in a year over the total DC energy consumption during the same period, namely

$$REF = \frac{E_{REN}(Wh)}{E_{DC}(Wh)}$$

where E_{REN} is the renewable energy drawn by the DC. According to ETSI specification ETSI GS OEU 001, an energy source is considered renewable if originating from renewable non-fossil sources, namely wind, solar, aerothermal, geothermal, hydrothermal and ocean energy, hydropower, biomass, landfill gas, sewage treatment plant gas and biogases. In the same specification, ETSI refers to REF as KPI_{REN}.

IT Equipment Energy Efficiency (ITEE)

IT Equipment Energy Efficiency (ITEE) is a unit-less metric indicating the energy efficiency of IT equipment defined as

$$ITEE = \frac{Rated\ capacity\ of\ IT\ equipment}{Rated\ power\ of\ IT\ equipment} = \frac{\alpha \cdot \sum_{DC} Server\ Capacity\ (GTOPS) + \beta \cdot \sum_{DC} Storage\ Capacity\ (GB) + \gamma \cdot \sum_{DC} Network\ Capacity\ (Gbps)}{Rated\ power\ of\ IT\ Equipment\ (W)}$$

where α, β, γ are coefficients to integrate the capacity of the server, storage and network equipment, respectively [32]. These values of these coefficients are defined as the inverse of the nominal capacity values for average, standard server, storage and network equipment as of 2005 and can be extracted from relative reference tables; an average DC comprising equipment of 2005 is expected to exhibit an ITEE value of 1. The definition of ITEE indicate that more efficient equipment should exhibit $ITEE > 1$, i.e. it should be able to exhibit higher processing capacity per unit of power consumed.

Evidently, ITEE only considers IT equipment classified under three categories (server, storage and network). Hence, DC equipment that cannot be categorized under one of these three categories should be excluded from the calculation of ITEE.

IT Equipment Utilisation (ITEU)

The IT Equipment Utilisation (ITEU) is a metric indicating the operation efficiency of IT equipment and is defined as

$$ITEU = \frac{Actual\ energy\ consumption\ of\ the\ IT\ equipment\ (Wh)}{Rated\ energy\ consumption\ of\ the\ IT\ equipment\ (Wh)}$$

ITEU is used to estimate the utilisation of the DC equipment; in a DC with no operating IT equipment the ITEU value is 0, whereas a DC with IT equipment operating at full load, the ITEU value is 1.0. The calculation scope for ITEU includes all the IT equipment (server, storage, network) and other IT equipment. On the other hand, equipment powered-off are excluded from the calculation as they do not incur explicit rated energy consumption. The measurement period and measurement points of the measured power level shall be based on the published measurement method of PUE.

Energy Reuse Effectiveness (ERE)

Similarly to the PUE definition, the Energy Reuse Effectiveness (ERE) metric is defined as

$$ERE = \frac{E_{cooling} + E_{lost} + E_{lighting} + E_{IT} - E_{reuse}}{E_{IT}}$$

where $E_{cooling}$ is the energy used by the entire cooling system of the DC, E_{lost} is the energy lost in the power distribution system through line-losses and other infrastructure (e.g. UPS, PDU) inefficiencies, $E_{lighting}$ is the energy used to light the DC, E_{IT} is the energy used by the IT equipment (server, network, storage) of the DC in hand and E_{reuse} is the total energy (e.g. thermal) that the DC reused during its operation under the same time period.

As per mathematical definition, the nominal range of ERE is 0 to infinity, values less that 1.0 indicate that part of the energy brought into the DC is reused elsewhere, outside of the DC control volume. Interestingly, if an ERE value of 1.0 does not imply that the DC in hand is necessarily efficient; it could represent an efficient DC design combined with a small amount of energy reuse, or an inefficient DC design combined with a lot of energy reuse. Hence, to properly translate the effect of the ERE value of a DC, a combined analysis of both ERE and PUE should be performed [33].

Global Synthetic KPI (KPI_{GP})

Proposed by ETSI in its GS OEU 001 specification, the Global Synthetic KPI (KPI_{GP}) is a global metric combing information from various other metrics to assess the energy efficiency of a DC. Since the definitions of ETSI, ISO/IEC and the rest of the initiatives presented in the current document are slightly deviating, the direct translation of the DC parameters involved in the definition of the metric are used.

KPI_{GP} is calculated based on two values, that is DC_G and DC_P where DC_G defines the energy consumption *gauge* (level) of the DC and DC_P defines the performance of the DC based on DC_G. The default values of DC_G are tabulated in Table 4.

DC_G	E_{DC} value
S	$E_{DC} \leq 1GWh$
M	$1GWh \leq E_{DC} \leq 4GWh$
L	$4GWh \leq E_{DC} \leq 20GWh$
XL	$E_{DC} \geq 4GWh$

Table 4 - Default DC_G values according to ETSI GS OEU 001

Keeping in mind this DC categorisation, the default classes of DC_P are tabulated in Table 5.

DC commissioning date	since 2005		since 2005	
	DC_P			
Class	$\geq$	$<$	$\geq$	$<$
A		0,70		1,00
B	0,70	1,00	1,00	1,40
C	1,00	1,30	1,40	1,70
D	1,30	1,50	1,70	1,90
E	1,50	1,70	1,90	2,10
F	1,70	1,90	2,10	2,30
G	1,90	2,10	2,30	2,50
H	2,10	2,40	2,50	2,70
I	2,40		2,70	

Table 5 - Default DC_P values according to ETSI GS OEU 001

Note that 2005 is the year when the Kyoto protocol was put into force. Formally, following the notation already presented, DC_P is defined as

$$DC_P = \frac{E_{DC}}{E_{IT}} \cdot \left(1 - W_{REUSE} \cdot \frac{E_{REUSE}}{E_{DC}}\right) \cdot \left(1 - W_{REN} \cdot \frac{E_{REN}}{E_{DC}}\right)$$

where W_{REUSE} and W_{REN} are coefficient factors for E_{REUSE}/E_{DC} (ETSI interpretation of ERE) and E_{REN}/E_{DC} (ETSI interpretation of REF), respectively. Interestingly, the factor E_{DC}/E_{IT} could be also directly related to the PUE metric.

Having defined the two values DC_G and DC_P, KPI_{DG} is defined as the tuple (Gauge, Class), e.g. (S,E).

Cooling Effectiveness Ratio (CER)

Cooling Effectiveness Ratio (CER) should quantify the energy efficiency of the cooling system of a DC, namely specify the percentage of the total energy used to power the DC cooling system that is translated into actual cooling services. The metric is under discussion by various standardization bodies but a definite interpretation has not yet been given.

Summing up, the following table summarises the standardised/close to be standardised metrics presented so far and how each standardisation body/initiative considers each of these metrics. In green colour we denote the metrics that have been taken into consideration by the standardisation bodies/initiatives, whereas red colour denotes the opposite.

		Standardisation Bodies/Initiatives			
		ISO/IEC	ETSI	GREEN IT	EC DC Cluster Collaboration
Metric	PUE	✓	✓	✓	✓
	DCiE	×	×	×	×
	REF	✓	✓	✓	✓
	ITEE	✓	×	✓	✓
	ITEU	×	×	✓	×
	ERE	×	✓	×	✓
	KPI_{GP}	×	✓	×	×
	CER	×	×	×	×

Table 6 - Existing metrics and their relation to the already mentioned standardisation bodies/initiatives.

2.2.2.2. Energy Metrics defined by EC Cluster Collaboration initiative

As aforementioned, the EC Cluster Collaboration initiative aims at determining meaningful KPIs, able to assess numerous aspects of DCs that are handled by current metrics and KPIs, like the ones already mentioned. The work already performed has resulted in the definition of 25 new metrics, classified into 8 distinct categories, depending on the target aspect of the DC of interest. The 8 categories are identified are (1) Energy/Power consumption, (2) Flexibility mechanisms in DCs: Energy Shifting, (3) Flexibility mechanisms in DCs: Energy being federated, (4) Renewables Integration: Energy produced locally and Renewables usage, (5) Energy Recovered: Heat Recovered, (6) Primary energy savings and CO_2 avoided emissions, (7) Economic savings in energy expenses and (8) Capacity planning and management. As exhaustively describing all the new metrics identified by the DC Cluster Collaboration partners is outside of the scope of this document, only the ones that are of interest to DOLFIN will be quickly presented; the interested reader is requested to refer to [34] for further clarifications

2.2.2.2.1 Energy/Power consumption

This category comprises metrics that help quantify the energy efficiency of the various DCs. Specifically, three metrics have been classified into this category, namely PUE, CER (already analysed) and Energy Effectiveness for HVAC cooling mode in a session (EE), analysed in the following.

Energy Effectiveness for HVAC cooling mode in a season (EE)

Energy Effectiveness for HVAC cooling mode in a season (EE) is a metric quantifying the energy efficiency of the HVAC system in a DC. Formally, it is defined as

$$EE = \frac{\sum_{i=1}^{N} m_i \cdot C_p \cdot \left(T_i^{inlet} - T_i^{outlet}\right)}{\sum_{i=1}^{N} P_{HVAC,i}}$$

where i represents the time of the measurement, m_i is the air flow rate (kg/s), C_p is a system-specific heat coefficient $(J/kg\ ^oC)$, $T_i^{inlet}$and T_i^{outlet} are the inlet and outlet air temperature $(\ ^oC)$ and $P_{HVAC,i}$ is the electricity power demanded by the HVAC equipment at time instant i.

The nominator of the EE metric definition refers to the theoretical consumption of the DC HVAC system under consideration, whereas the denominator stands for the actual consumption due to HVAC operation. An EE value of 1.0 suggests that the actual consumption matches the theoretical one (there are no power losses in the DC HVAC system) whereas lower values of EE indicates some power loss; an EE value of 0 indicates that theoretically the HVAC system should be either idle (the inlet and outlet temperature are identical) or acting as plain ventilator.

2.2.2.2.2 Flexibility mechanisms in DCs: Energy Shifting

This category comprises metrics that characterise the ability of the DCs to react to explicit changes of their operation or environment. The identified metrics are Adaptability Power Curve (APC), Adaptability Power Curve at Renewable Energy (APC_{REN}), DC Adapt (DCA), Flexible Energy Rate (FER), Managed Energy Rate (MER) and, last, Managed Flexible Energy Rate ($MFER$).

Adaptability Power Curve (APC) and Adaptability Power Curve at Renewable Energy (APC_{REN})

APC determines the capability of a DC to adapt its operation to a planned optimal energy consumption pattern. It is formally defined as

$$APC = 1 - \frac{\sum_{i=1}^{n} \left|E_{DC,i} - K_{APC} E_{P,i}\right|}{\sum_{i=1}^{n} E_{DC,i}}, K_{APC} = \frac{\sum_{i=1}^{n} E_{DC,i}}{\sum_{i=1}^{n} E_{P,i}}$$

where $E_{DC,i}$ stands for the energy consumption of the DC at time instance i, $E_{P,i}$ is the planned energy consumption at time instance i and K_{APC} is a mitigation factor to render the two energy consumption patterns comparable. An APC value of 1.0 indicates that the energy consumption of the DC is completely adapted to the optimal, planned energy use, whereas values close to 0 indicate inability of the DC to adapt to the desired state.

Similarly, APC_{REN} considers DC adaptability on the basis of the power production pattern of the (locally produced or not) renewable energy consumed by the DC, namely

$$APC_{REN} = 1 - \frac{\sum_{i=1}^{n} \left|E_{DC,i} - K_{APC_{ren}} E_{ren,i}\right|}{\sum_{i=1}^{n} E_{DC,i}}, K_{APC_{ren}} = \frac{\sum_{i=1}^{n} E_{DC,i}}{\sum_{i=1}^{n} E_{ren,i}}$$

where $K_{APC,ren}$ is a coefficient similar to K_{APC}, but referring to the renewable energy consumption.

DC Adapt (DCA)

This metric indicate the ability of a DC to deviate from a baseline energy consumption profile and is defined as

$$DCA = 1 - \frac{\sum_{i=1}^{n} \left|K_{DCA} \cdot E_{DC,real,i} - E_{DC,baseline,i}\right|}{\sum_{i=1}^{n} E_{DC,baseline,i}}, K_{DCA} = \frac{\sum_{i=1}^{n} E_{DC,baseline,i}}{\sum_{i=1}^{n} E_{DC,real,i}}$$

following a notation similar to the one used in APC and APC_{REN}. DCA values close to 1.0 indicate that the DC is able to considerably shift its energy consumption baseline whereas values close to 0 suggests that the adaptability mechanisms applied to the DC are less efficient than expected.

Flexible Energy Rate (FER)

This metric evaluates the percentage of the DC energy consumption that can be adapted. Formally it is defined as follows:

$$FER = \frac{FE}{E_{DC}}$$

where FE is the energy consumption of the DC that is flexible/adaptable. High FER values (close to 1.0) indicate higher ability of the DC to adapt itself. It is worth mentioning that although an energy consumption source may be considered as flexible, it may not be manageable. The latter consideration is handled by MER and $MFER$ metrics, whose analysis follows.

Managed Energy Rate (MER) and Managed Flexible energy Rate ($MFER$)

These two indices evaluate the ability of the DC equipment to be actually managed, can be, in fact, adapted. The definition of MER is similar to the one of FER:

$$MER = \frac{ME}{D_{DC}}$$

where ME is the total energy consumption of the DC that can be managed. MFER combines FER and MER, being defined as the percentage of the flexible energy consumption that can be managed, namely

$$MFER = \frac{MER}{FER} = \frac{ME}{FE}$$

2.2.2.2.3 Flexibility mechanisms in DCs: Energy being federated

This KPIs category contains metrics that consider energy federation among DCs, namely evaluate the DCs ability to share workload (hence, energy) with other DCs of the same or different administrative domain in order to achieve better energy and cost efficiency. The metrics that were identified by the DC Cluster Collaboration projects were Federated Energy Weight (FEW), Federated COP ($COP_{federated}$) and Federated RES (RES_{FED}).

Federated Energy Weight (FEW)

This metric measures the percentage of energy that has been federated during a specific time period, namely the energy that was required for a computing load that was re-located to another DC, over the DC energy consumption of the DC during the same period:

$$FEW = \frac{\sum_{i=1}^{N} FedE_{DC,i}}{\sum_{i=1}^{N} E_{DC,i}}$$

where $FedE_{DC,i}$ is the total energy (kWh) that was federated over the reference time period, say $T = N \cdot t$, where t is the measurements period. Evidently, high FEW values indicate that the DC relocated a significant amount of its workload (thus energy consumption) to another DC, whereas low FEW values suggest that the DC handled most of its assigned workload on its own.

Federated COP ($COP_{federated}$)

This metric evaluates the efficiency of federating a number of services to other DCs. Specifically, it identifies if it is more efficient to share computing load, if possible. Assuming that there is a pool of M cooperating DCs, the $COP_{federated}$ of DC $j \in [1, M]$ is defined as

$$COP_{federated,j} = \frac{\eta_j \cdot \sum_{i=1}^{M} WD_{DC,i}}{\sum_{i=1}^{M} \eta_{DC,i} \cdot WD_{DC,i}}$$

where $WD_{DC,i}$ is the work (number of processes) performed by DC_i, $\eta_{DC,i}$ is the energy needed to satisfy a given working load $WD_{DC,i}$. $COP_{federated,j}$ values smaller than 1.0 mean that it is less efficient to relocate workload from DC_j to the rest of the DCs, values larger than 1.0 indicate that it is more efficient and $COP_{federated,j} = 1.0$ indicates that a federation option is indifferent to DC_j.

Federated RES ($Federated_{RES}$)

This metric quantifies how much the use of renewable energy was increased due to energy federation during a predefined period of time $T = N \cdot t$, where t is the measurement period:

$$Federated_{RES,j,k} = \frac{\sum_{i=1}^{N} FedE_{DC,i} \cdot RenPercent_{i,j} - FedE_{DC,i} \cdot RenPercent_{i,k}}{\sum_{i=1}^{N} FedE_{DC,i}}$$

where $FedE_{DC,i}$ is the energy federated at the time instance $i \cdot t$ and $RenPercent_{i,j}$ represents the percentage of RES usage of DC_j at the time instance $i \cdot t$. Values of $Federated_{RES,j,k}$ higher than 0 indicate that the utilization of RES was decreased due to the workload federation from DC_j to DC_k and vice versa.

2.2.2.2.4 Renewables Integration: Energy produced locally and Renewables usage

This metrics category deals with the evaluation of the DC capabilities to properly operate using the highest possible percentage of energy coming from RES. The identified KPIs where $RenPercent$, $RenEPPercent$, $RenThermPercent$, $RenEPThermPercent$, $TotalEPPercent$, REF (already defined) and a small set of KPIs related to the DC-Power Grid interaction. $RenPercent$, $RenEPPercent$, $RenThermPercent$, $RenEPThermPercent$ and $TotalEPPercent$ indicate the share of RES in the total electricity consumption of a DC, either aggregated or segregated depending on the type of power each metric refers to (primary energy, thermal energy, or total). For more information, the interested reader is requested to refer to [34].

Grid Interaction Indicators

The Grid Interaction Indicators proposed by the EC Cluster Collaboration initiative indicate the level of interaction between a DC and a power distributing source, typically the power grid. Assuming that a DC has access to owned energy sources, these indicators can be used to quantify the dependence of the DC in hand on the power grid it is attached to. In other words, these indicators are not related to performance or energy efficiency but, rather, to DC energy sufficiency.

Load cover factor (γ_{load})

Load cover factor (γ_{load}) reflects the percentage of the electrical energy demand covered by local energy generation over a predefined time period T. Formally, it is defined as

$$\gamma_{load} = \frac{\int_T \; min[P_{REN}(t), P_{DC}(t)] \; dt}{\int_T \; P_{DC}(t) \; dt}$$

where t indicates a time instance in T. By definition, the values of γ_{load} vary between 0 and 1.0, the former value indicating low utilisation of the local energy resources and the latter indicating abundance in the local energy generation (fully satisfying the total DC energy demand).

Supply cover factor (γ_{supply})

The supply cover factor (γ_{supply}) is a dual indicator of γ_{load}; instead of using the DC energy demand in total as a reference, it uses the local energy production, namely

$$\gamma_{supply} = \frac{\int_T \ min[P_{REN}(t), P_{DC}(t)] \ dt}{\int_T \ P_{REN}(t) \ dt} = \gamma_{load} \cdot \frac{\int_T \ P_{DC}(t) \ dt}{\int_T \ P_{REN}(t) \ dt}$$

where the notation is similar to γ_{load}. In contrast to γ_{load}, low γ_{supply} values indicate that the local energy sources satisfy a large portion of the DC energy consumption and vice versa.

Loss of load probability ($LOLP_b$)

The loss of load probability ($LOLP_b$) is defined as the probability that the local generation does not satisfy the energy demand of the DC in hand. The value of $LOLP_b$ can be approximated making use of historical DC power measurements over a (large-enough) period of time T, as follows

$$LOLP_b = \frac{\int_T \ f(t) \ dt}{T}, f(t) = \begin{cases} 1, P_{REN}(t) - P_{DC}(t) < 0 \\ 0, P_{REN}(t) - P_{DC}(t) \geq 0 \end{cases}$$

Evidently, $LOLP_b$ values close to 0 indicate that the DC usually covers its electricity demand through local power generation, whereas values close to 1.0 may be translated into abundance of local electricity generation.

2.2.2.2.5 Energy Recovered: Heat Recovered

This category deals with the evaluation of the DCs based on the energy that they can reuse, e.g. by reusing thermal energy to produce electricity. Two distinct metrics have been identified, the one being ERE (already defined) and the other being $ReusePercent$, defined as the share of waste energy reused by a DC (either the one who produces the wasted energy or another). $ReusePercent$ is not considered of interest to DOLFIN; thus further analysis is omitted.

2.2.2.2.6 Primary energy savings and CO_2 avoided emissions

This category considers the energy saving were achieved by adjusting DC operation as well as the CO_2 emissions that were avoided as a result of the application of a proper resource management plan.

Primary energy savings: PE savings

This metric calculates the effect that the application of a resource efficiency plan to the DC operation has to the primary energy consumed by the DC. Since this metric is out of DOLFIN scope, its analysis is omitted and the interested reader should refer to [34] for further details.

CO_2 Savings/Avoided emissions

This metric compares the CO_2 emissions reduction that was achieved as a result of a proper resources management plan at DC level. Granted that proper baseline measurements are available, the definition of the metric follows:

$$CO_2S_{DC} = \frac{\sum_{i=1}^{N} \left[\left(E_i^{DC} + E_i^{other} \right)_{baseline} - \left(E_i^{DC} + E_i^{other} \right)_{current} \right]}{\sum_{i=1}^{N} \left(E_i^{DC} + E_i^{other} \right)_{baseline}}$$

where E_i^{DC} is the CO_2 emissions of the DC at time instant i and E_i^{other} represents the CO_2 emissions from other sources at the same time instance, t. In the definition of CO_2S_{DC} it is supposed that measurements are taking place periodically every t time instances, during time period T, namely $T = N \cdot t$.

CO_2S_{DC} is not a metric characterising the computational performance or energy efficiency of a DC but, rather quantifies the change in the environmental friendliness of a DC after a relevant management plan has been applied. It should be underlined that the baseline CO_2 emissions are subject to proper adjustment in order to acquire a subjective CO_2S_{DC} assessment. Such an adjustment is imperative since the baseline reference for the CO2 emissions are taken prior to measuring the current DC CO2 emissions. The adjustment process depends vastly on the measurement process but, in any case, should consider changes in the DC equipment and DC operation in order to make the comparison meaningful.

2.2.2.2.7 Economic savings in energy expenses

This category comprises a single metric, namely Energy Expenses (EES), which determines how did the energy expenses change after the application of a certain resource management scheme to upgrade the resource (in terms of energy, cost and environmental friendliness) usage of the DC. Since this metric is out of DOLFIN scope, its analysis is omitted and the interested reader should refer to [34] for further details.

2.2.2.2.8 Capacity planning and management

In this category, the DC Cluster collaboration included the already presented metric $ITEE$.

3. Overall System Architecture Description

3.1. Synthesis of the DC energy efficiency characteristics

The DOLFIN system is an ecosystem of collaborative Data Centres. DOLFIN considers a number of Data Centres, each one having its own DC customers and links with the energy network, as shown in Figure 3-1. From the energy network perspective, DOLFIN considers both traditional Energy Providers and Smart Grid Networks in a Smart City scenario. Moreover, DOLFIN considers that each DC may achieve further internal energy efficiency, by recycling and reusing the warm water that is used for cooling the ICT equipment for warming the DC offices (e.g. with an under floor warming system). We consider the DOLFIN System as a complete ecosystem that has a group of customers and an interaction with the energy network. These are not directly connected with a specific DC but use the DOLFIN system as a Virtual Data Centre. The main objective of the DOLFIN solution is to reduce the energy consumption of the elements within the DOLFIN ecosystem as a whole (especially the consumed brown energy) and stabilize the smart grid electricity network, wherever needed, without breaking the agreed SLAs with the end users (or renegotiating the SLAs based on signed contracts).

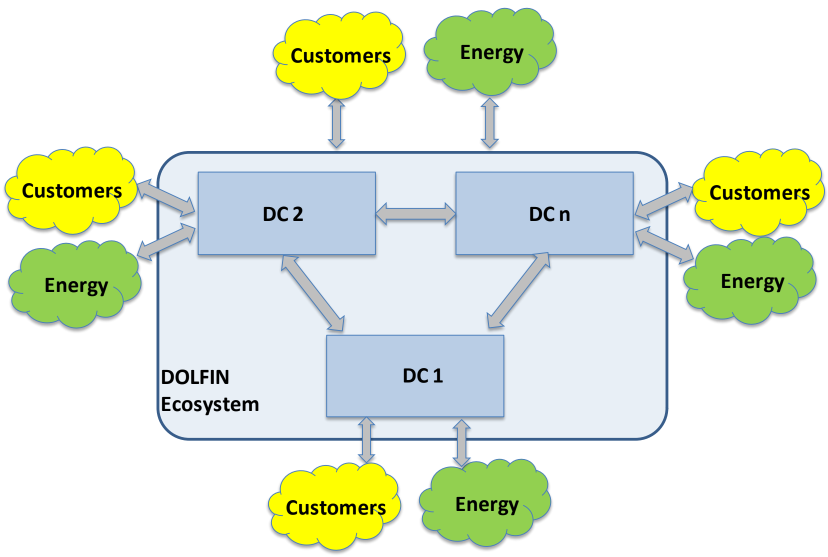

Figure 3-1: DOLFIN Interworking DCs – High Level Architecture

Current and future Data Centres are complex systems, comprised of diverse cloud management and autonomic functions, Virtual Machines (VM) hypervisors, legacy Energy Management and Uninterrupted Power Supply (UPS) systems, Building Management Systems (BMS) and HVAC (heating, ventilation, and air conditioning) systems. In most cases, DCs have legacy Data Centre Infrastructure Management (DCIM) systems which are quite sophisticated. Therefore, totally replacing this existing infrastructure would be unrealistic. Only new DOLFIN-aware DCs would significantly lower the deployment potential of DOLFIN.

Consequently, the DOLFIN approach is to design an architectural framework, which explicitly factors in and accommodates the existing DC ICT and energy management systems, and in parallel:

- Improves both capital and operational efficiencies for DC operators through the use of a common organization, automation, and operations of all energy functions across the different domains, and
- Migrates from separate DCs energy management functions towards a coordinated ecosystem of energy management functions in a group of DCs

To achieve these objectives and meet the DOLFIN requirements as expressed in this chapter, we follow a distributed spiral approach that aims to optimize the energy (as shown in Figure 3-2). During each optimization cycle of the Figure 3-2 spiral, the energy is optimized by an internal control loop, meeting the system requirements (SLAs and Smart Grid stabilization) and then moves stepwise to the outer loop. Each optimization process, may recursively initiate optimization at the hierarchically lower optimization level in a closed loop approach. Initially DOLFIN optimizes the energy consumption at server rack level, then continues to energy consumption at DC segment level, and then energy consumption at DC level. Finally the outer process, performs optimization at DOLFIN level, enabling movement between energy conscious DCs.

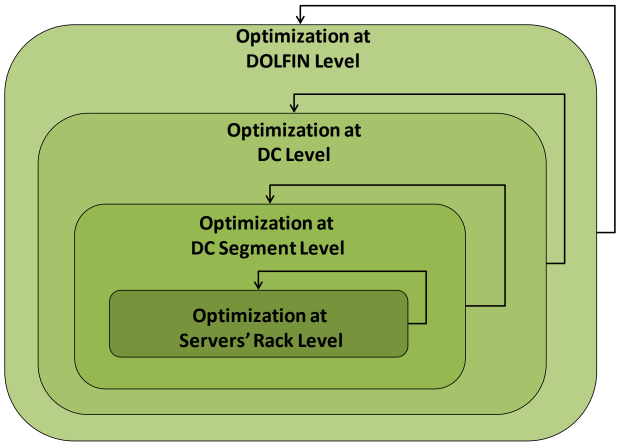

Figure 3-2: DOLFIN Optimization Process

The justification for this approach is due to the fact that moving VMs between DCs requires migration not only of the processes, but also the swap space and disk space. This requires the movement of significant amount of data and consequently stresses the end-users SLAs. The above process is further analyzed in the section DC Functional Architecture.

3.2. DC Functional Architecture

The DOLFIN functional architecture of each individual DC is defined in terms of new DC functional blocks, revised DC functional blocks, and the interworking interfaces needed to realise different energy closed control loops and their interactions with the normal operation of the DC. The functional architecture is depicted in the Figure 3-3.

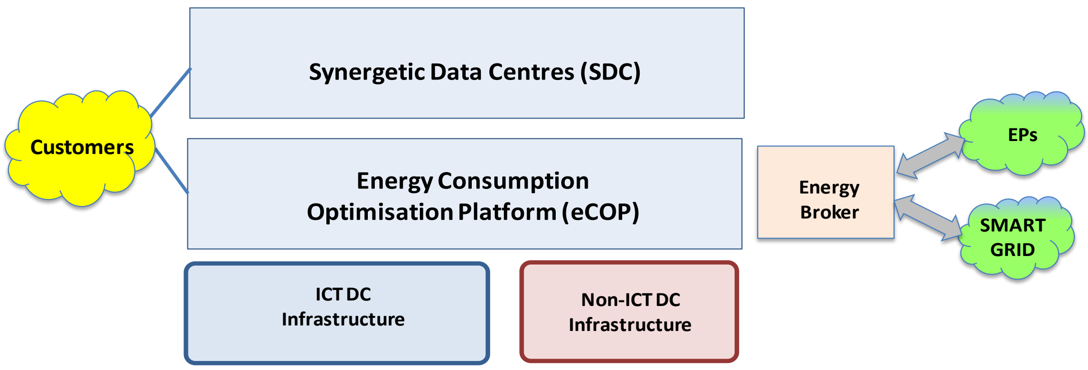

Figure 3-3: DOLFIN Functional Architecture

The main DOLFIN components are:

- The **Energy Consumption Optimisation Platform (eCOP)**, which is the DOLFIN core platform for energy consumption optimization at Data Centre level. It allows for continuous monitoring, energy benchmarking, dynamic control, and adaptive optimisation of the Data Centre infrastructure, including ICT and HVAC equipment.
- The **energy-conscious Synergetic Data Centres (SDC) module**, which provides a dynamic, service-effective and energy-efficient allocation of demands, across a distributed network of co-operating DCs.
- The **Energy Broker** that extends the **Smart-Grid Energy stabilisation module**, which controls the interconnection with the smart grid network, providing responses on the changing demands for energy, with a module that controls the legacy (brown) Energy Providers.

The above architecture is analysed in further detail whereby the SDC and the eCOP are presented and the sub-components of each are described. This functional architecture refers only to the S/W platform in a DC and it is depicted in Figure 3-4.

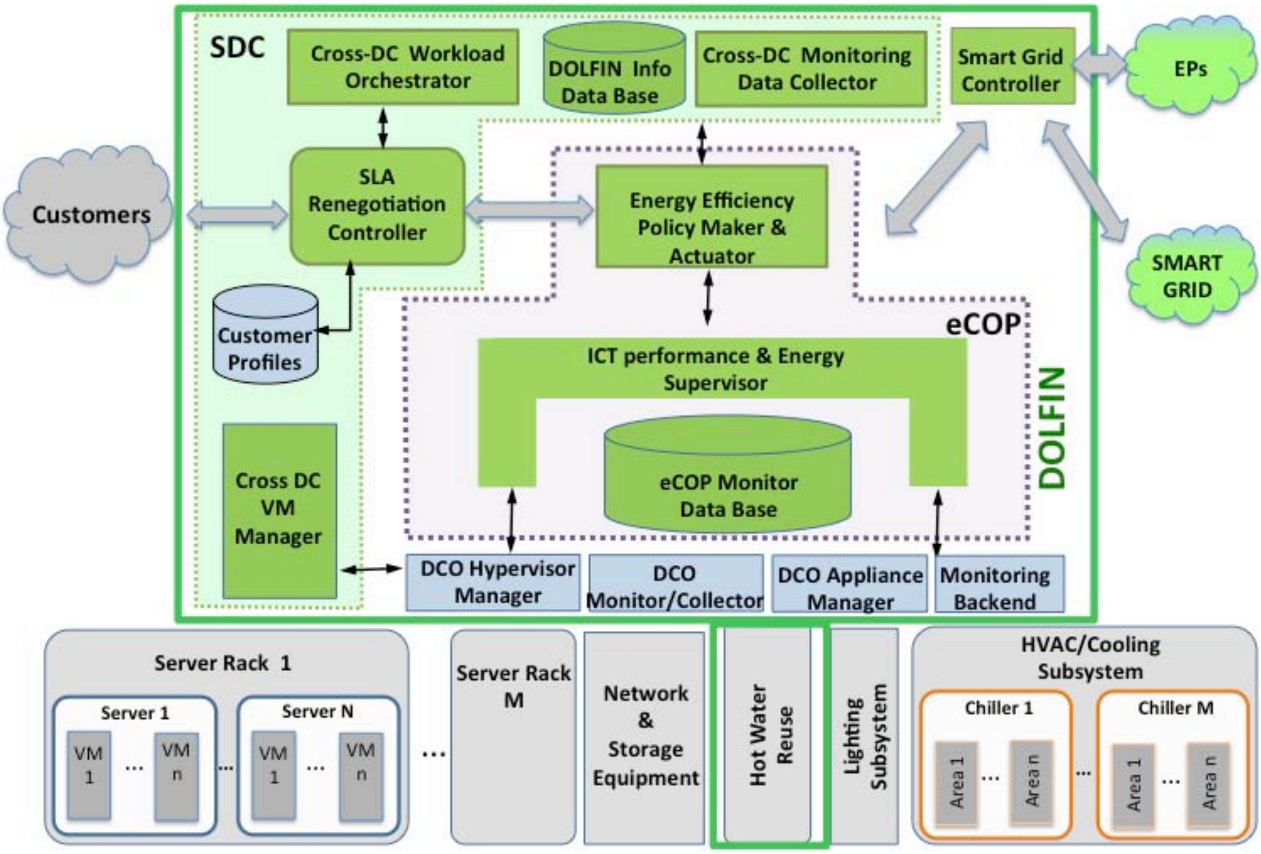

Figure 3-4: Detailed DOLFIN Functional Architecture

Using a bottom up approach, the architecture has at the lower level two types of infrastructure:

- **ICT DC Infrastructure.** This group includes the equipment of the DC servers, per Server Rack, including all VMs and software components, the network and storage equipment, together with all of the swapping and storage functionality of all applications.
  In this infrastructure group we also include any existing or legacy DC management system that pre-exists. There are already a number of DC Management systems, of both hardware and software, to support the ICT DC infrastructure (e.g. ABB Decathlon). Interfacing and inter-operating with these,

instead of replacing them with the DOLFIN system, would ensure faster deployment and easier acceptance of the DOLFIN system.

- **Non-ICT DC infrastructure.** This group of equipment includes the HVAC/Cooling Subsystem, which may be organized in chillers that heat/cool specific areas of the DC, the lighting subsystem (which may also be organized in areas), and the hot water subsystem.

  We further highlight, the DOLFIN **Water Cooling Systems**. Traditionally, in most DCs, cooling at the server level is air based. Air cooling is simple and safe. On the other hand, air is not thermally efficient and thus impractical for cooling high-density implementations. The most attractive alternative to air cooling is liquid based cooling. One of the main attractions of liquid cooling is the greater ability of liquids to capture and hold heat, relative to air. Therefore, a smaller amount of water can accomplish the same heat capture and removal as a relatively large amount of air, enabling a targeted cooling strategy. For high-density implementations, air cooling is often simply insufficient. In such cases, water (or, more generally, liquid) cooling offers an alternative in which not only can greater cooling efficiency be applied, it can be applied only where it is needed.

  In this group of infrastructure we also include any existing or legacy **Building Management System (BMS)** that pre-exists and has to be interfaced and inter-operated by the DOLFIN system.

Moving up the architectural stack, we have various elements plus the eCOP module. We highlight and summarise these components which have revised or new DC functionality. A more detailed description is provided in Section 4.3.

- **DCO Hypervisor Manager (revised DC functionality) –** The DCO Hypervisor Manager is an adaptation layer that maps the high-level decisions taken by the DOLFIN system into low-level technology-dependent commands that control the DC ICT infrastructure. This process may be executed in two steps: first, the high-level decisions are translated into low-level technology-agnostic actions; second, for each action to be executed, the action must be translated into real operations, specific to the DC control framework, as supported by the actual DC control system in use. The DCO Hypervisor Manager analyses the new DCs ICT hardware including DCs configuration and status. It is responsible for managing new ICT hardware and VMs configuration. In addition, it is also responsible for the translation of new hardware configurations to dedicated commands to manage the specific virtualization hypervisor and the execution of commands to provide a new DCs hardware configuration.

- **DCO Appliance Manager (revised DC functionality)** – The DCO Appliance Manager is the counterpart of the DCO Hypervisor Manager, on the non-ICT DC infrastructure. The DCO Appliance Manager adapts and guarantees the execution of strategies elaborated by DOLFIN expert system on the non-ICT infrastructure by translating high level commands into low-level technology-agnostic actions, and then realises the real operations, specific to the DC control framework, either directly or via the relevant BMS system. Moreover, the DCO Appliance Manager analyses the new DCs non-ICT hardware including the configuration and status.

- **DCO Monitor/Collector (revised DC functionality) –** The role of the DCO Monitor/Collector is to interface with both the ICT and the non-ICT DC infrastructure and collect all operational and energy related information to be stored in the eCOP Data Base. Information that is collected includes the DC server performance data (usage of both CPU and RAM for the server level and the VM level), electricity consumption data, DCs temperature data, mapping of DCs infrastructures (servers and VMs topology).

- **eCOP Monitor Data Base (new DC functionality)** - The eCOP Monitor Data Base has the sole purpose of storing all real-time and historical energy related data collected from DCs. This information is used for energy efficiency decisions and VM load predictions, along with data from all the components within a DC architecture. The eCOP Monitor database has interfaces to both the DCO Hypervisor and the DCO Appliance Manager, which store the relevant information into the data base. The eCOP Monitor DB may be also accessed by a monitoring backend component, which can be connected with a ***dashboard*** to offer a Graphical User Interface (GUI) monitor tool on top of the ICT Performance and Energy Supervisor system. The main objective of this interface is to provide an easy way to consult all the relevant information related to ICT Performance and Energy Supervisor.

- **ICT Performance & Energy Supervisor (new DC functionality) -** The ICT Performance and Energy Supervisor focuses on the analysis of performance monitoring data and energy data. The objective of this component is to provide information on the actual performance of the applications (typically VMs) utilizing the resources of the DC devices and the energy consumed by the non-ICT components. Primarily, this module will receive as input performance utilisation data collected at the DOLFIN eCOP Monitor DB, including CPU utilisation, memory consumption, link utilisation, etc that are typically available for ICT devices, namely servers, storage equipment and networking devices. It will then perform an analysis to describe the utilisation levels of each device.
- **Energy Efficiency Policy Maker & Actuator (new DC functionality).** This is the most intelligent part of the eCOP that makes the decisions realizing the requested policy of the DC. Following the spiral approach of Figure 3-2, this component realizes the 3 lower parts of the energy optimization, achieving optimization at DC level. Moreover, it aims at initiating a stream of control commands, which can be translated into actual actions by other DOLFIN subsystems within a DC.

To achieve energy efficiency at cloud level, the DOLFIN eCOP collaborates with the ***Synergetic Data Centres (SDC)*** module which consists of the following components:

- **Cross-DC Monitoring Data Collector (new DC functionality)**, which collects knowledge not only from the Synergetic Data Center resources, but also from the network routers and in general from the network resources point of view. This is an important prerequisite to achieve energy efficiency across DCs.
- **Cross-DC Workload Orchestrator (new DC functionality)**, which is a distributed software element that gets the SDC decisions and does most of the different types of resource optimisation including energy optimisation and management of their trade-offs in a cross DC optimization scenario. It is in charge of managing the full lifecycle of the virtual routers in the network and the allocation of the applications running on the virtual nodes
- **Cross-DC VM Manager (new DC functionality)**, which realises a DC Interconnect (DCI) interface and performs the actual migration of VMs cross DCs. It will typically apply a set of standard alternatives for coping with high/peak workloads, more precisely with allocation of VMs, data, services and tasks.

The Cross DC Monitoring Collector, Workload Orchestrator and VM Manager (also referred together as ***Workload and VM Manager***) collaborate closely in order to offer a flexible, reliable, and fast communications solution to handle the increased network traffic arising from the DCs operation.

- **DOLFIN Information Data Base (new DC functionality)**, offers an abstracted and logically-centralised information manipulation (including information collection, aggregation/processing, storage / indexing and distribution) across all DOLFIN architectural components. It includes information on the local energy requirements e.g. Demand/Response requests from the local Smart Grid operator. The Information Base uses two separate interfaces for communication with the DOLFIN software entities: (1) the Information Management Interface is used for information manipulation configuration, including the information sources/sinks registration to the Information Base, the management of internal Information Base, information manipulation functions and the establishment, operation and optimisation of information flows; and (2) the Information Exchange Interface that offers the actual information exchange capability to the DOLFIN components. These two interfaces allow a unified yet abstracted information handling for all types of participating entities.

The DOLFIN expert system needs to have the topology data of infrastructures in DCs, the electricity consumption, and infrastructure performance data of hardware infrastructures of the DCs constantly available in order to manage systems in the DCs in a smart optimized, and automated way and as much as possible in "real time". This is necessary to achieve energy savings and the resulting C02 emission reductions and to meet the different demands from EPs in terms of optimization or different distribution of the power consumption of the smart grids,. This info is also stored in the DOLFIN Energy Information Base.

- **SLA Renegotiation Controller** **(new DC functionality)** - The SLA Renegotiation Controller guarantees to steer the data centre operation towards environmentally sound behaviour. This module will take into account the existing approaches in modelling SLA criteria (e.g. the overall cost of an offered service) and augment these approaches with the use of energy related criteria. This module will also handle the negotiation process between the DC and the end-users, allowing the DC to request from an end-user to decrease the performance metrics of their SLA, in exchange of lower service costs. The component will allow the negotiation to be dynamic, in the sense that a new SLA offer by the Data Centre provider may be followed by modified suggestions by the customer. This component will also consider the SLA negotiation process between co-operating Data Centres for the SLA that govern the services that are being transferred between them.
- **Customer profile management** **(revised DC functionality)** - DC and cloud infrastructure users exhibit a wide variety of profiles and requirements as they originate from different sectors such as industry, academia, education, research and are widely geographically and culturally distributed. To that extent, the customer profile management is greatly influenced by the technical requirements of the services requested by the DC of cloud infrastructure, but also by the financial, societal and even moral status of the customer.

The DOLFIN System includes the **Smart Grid Controller (SGC)** which is a Smart City Component. The SGC interconnects the DC with the Smart city electricity network. This most critical part of the system acts as a broker between the core DOLFIN infrastructure and subsystems, represented by the Energy Efficiency Policy Maker & Actuator, and the outside world, namely the Smart Grid (SG) and the Energy Providers (EPs).

3.3. Detailed Description - Revised and New DC functionality & Data and Control I/Fs

In order to increase deployment potential of the DOLFIN system, DOLFIN will be designed to incorporate as many existing and legacy systems as possible. In this section, we present the detailed description and WPs responsibility of the needed revised and new DC functionality and interfaces for data and control in order to fully realised the energy multi objectives optimisation and management for future data centres.

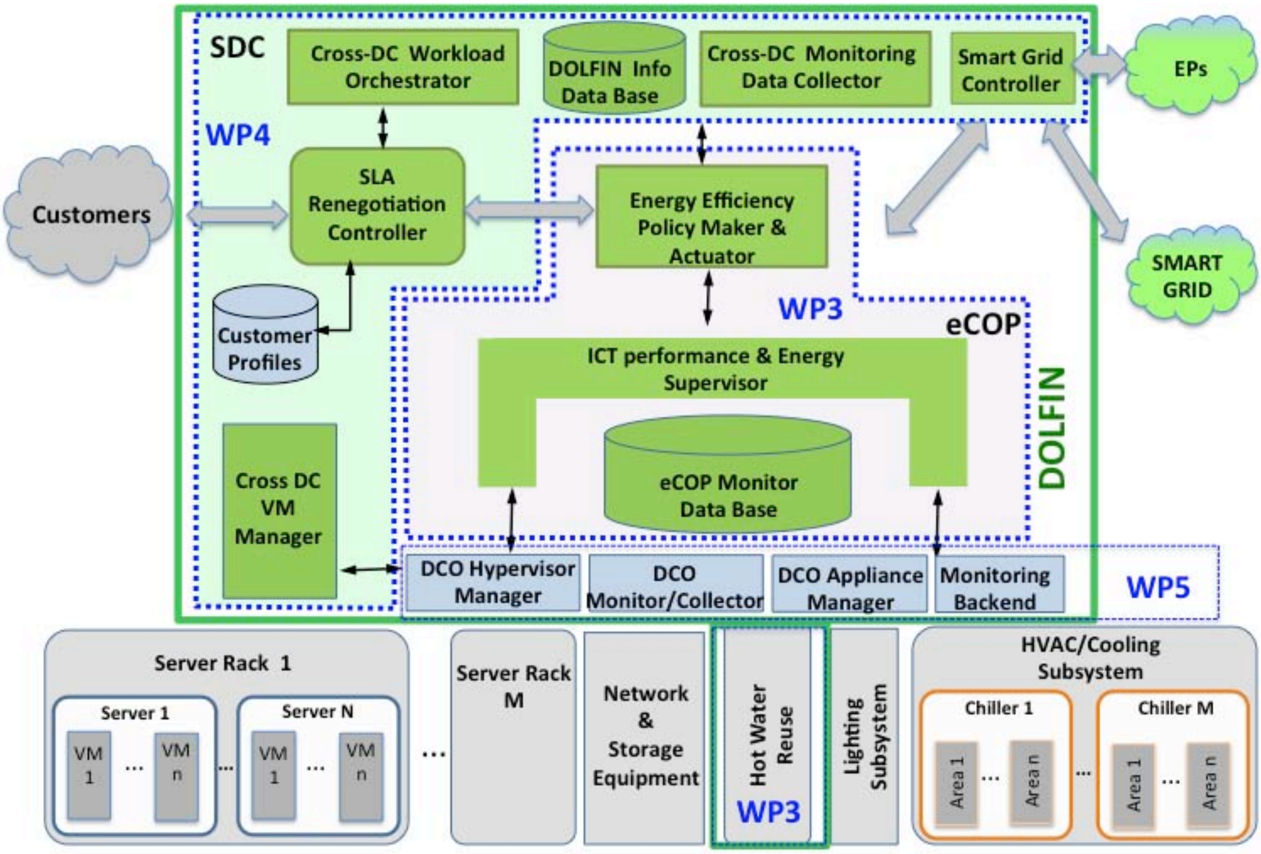

Figure 3-5: DOLFIN S/W Components with WPs responsibility

3.3.1. DCO Hypervisor Manager (revised DC functionality)

As already presented, the DCO Hypervisor Manager is a very important component in the DOLFIN architecture as it is responsible for the ICT DC infrastructure management, including the VMs configuration. In addition, it translates new hardware configuration in dedicated commands to manage the specific DC virtualization hypervisor and executes commands to provide a new DCs hardware configuration.

Legacy DCs based on virtualized architectures have a Hypervisor running on every server, which allows several virtual machines (VMs) to coexist and run concurrently on the same physical server. This practice is now industry-standard and has a great number of advantages. The most noticeable consequence of this is that virtualization technology is a key factor in achieving high energy efficiency: contrary to traditional system images, VMs can be stopped and frozen when not needed, leaving room for other running instances and decreasing the server load. They can also be migrated from server to server while retaining all their features, allowing for a physical server to always allocate as much physical resources as possible. Overall, this is very beneficial, since an idle server will consistently make sub-optimal use of its energy when compared with a correctly loaded and balanced one.

The DCO Hypervisor Manager is the component in DOLFIN which address the coordination and feeds orders to the hypervisor instances inside a DC. It turns high level messages produced by the Workload and VM Manager into a consistent sequence of lower level actions compatible with the specific brand of hypervisor which is to be dealt with. DOLFIN should be able to interface the ICT infrastructure of the DC to perform actions or get information. Thus, the DCO Hypervisor Manager should also operate as an adaptation interface, able to retrieve new DCs' ICT infrastructure configurations and send the new DCs' ICT hardware configuration via the injection of setup commands. Moreover, via the DC to DC component, it should control VMs migration between the DCs in a DOLFIN ecosystem.

The virtualization models commonly used in DCs, contain their own method for power management through their hypervisor. However, there is no unified model that enables the monitoring and management, in a coordinated manner, of the resources used by each VM. Some existing DC Hypervisors that we are going to analyse are:

- **XEN Hypervisor** [44] is the only type-1 (or bare metal) hypervisor that is available as open source. It makes possible to run many instances of an operating system or indeed different operating systems in parallel on a single machine. It is used as the basis for a number of different commercial and open source applications, such as: server virtualization, Infrastructure as a Service (IaaS), desktop virtualization, security applications, embedded and hardware appliances. The Xen Project hypervisor is powering the largest clouds in production today. It offers offline and live migration of VMs, used to transfer a VM between physical hosts. Offline migration moves a VM from one host to another by suspending, copying the VM's memory contents and then resuming the VM on the destination host. Live migration allows transferring a VM without a noticeable suspension. A live migration requires that both hosts are running Xen and the destination has enough spare resources. Xen's live migration technology can facilitate the development of various energy management techniques, utilising dynamic VM consolidation. Some of the key features are: Small footprint and interface (around 1MB in size), operating system agnostic and Driver Isolation.
- **VMware Hypervisor**. The VMware ESX Server and VMware ESXi, which are enterprise-level virtualization solutions [45] support host-level power management via Dynamic Voltage and Frequency Scaling (DVFS). The system monitors the CPU utilization and continuously applies appropriate ACPI's P-states. VMware supports live migration of VMs, between physical nodes, which can be initiated manually or programmatically. VMware monitors the resource usage in a pool of servers and continuously rebalances VMs, according to the current workload and load-balancing policy. VMware employs Distributed Power Management subsystem to reduce power consumption by a pool of servers by dynamically switching off spare servers. VMware utilizes live migration to reallocate VMs, keeping the minimal number of servers powered on.
- **Kernel-based VM Hypervisor (KVM)** [46] is implemented as a module of the Linux kernel. Under this model, Linux works as a hypervisor and all VMs are regular processes, scheduled by the Linux scheduler. KVM supports S4 (hibernate) and S3 (sleep/stand by) power states. KVM's parts are licensed under various GNU licenses. By itself, KVM does not perform any emulation. Instead, it simply exposes the /dev/kvm interface, with which a user space host can then: a) Set up the guest VM's address space, b) Feed the guest simulated I/O or c) Map the guest's video display back onto the host

DOLFIN will analyse the available solutions installed in each DOLFIN trial, including features of the OpenStack and realise the DOLFIN DCO Hypervisor Manager, measuring the energy consumption during VMs migration.

3.3.2. DCO Appliance Manager **(revised DC functionality)**

Similarly to the DCO Hypervisor Manager, the DCO Appliance Manager is responsible for the non-ICT DC infrastructure management, and the translation of new hardware configuration in dedicated commands to manage the specific DCs hardware configuration and BMS system. Some well known Appliance Managers are:

- **Oracle Appliance Manager.** Among the DCO Appliance Managers, one of the most often utilized is Oracle Appliance Manager [47]. It has unique knowledge of the system and is aware of the operating environment. It is able to: a) Make installation, configuration and tuning decisions as well as automate the setup, b) Enable customer support to diagnose and resolve issues very quickly, c) Issue "system" patch bundles, combining OS, database, clusterware, and storage updates, d) Proactively provide patches once known issues, raised by other customers, are resolved and fixed, e) Employ updates/patches at any element immediately when available — no cross-certification of multi-vendor technology to wait on, some of which can take more than a year to get tested and certified.

- **IBM *WebSphere* Appliance Manager** simplifies the management and monitoring of environments that consist of multiple WebSphere DataPower Appliances [48]. This web-based application provides centralized multi-appliance administration and key services, such as: centralized firmware management, disaster recovery, domain and service configuration, configuration life cycle deployment and monitoring multiple appliances, collecting key metrics, and presenting them in a central location. The WebSphere Appliance Management Center consists of the following two independent components: a) ***Management component***: is used to manage multiple IBM DataPower appliances by deploying configuration changes and firmware upgrades to multiple appliances, and securely backing up and restoring appliances and b) **Monitoring component**: is used to monitor the behaviour and status of appliances. It gathers information to detect problems, notifies you of common issues with the appliances that it monitors, and provides a central point of in the organization.
- **ABB Decathlon** [49] provides tools to manage a flexible network of power, cooling and IT for Data Centres, aiming to optimize: cost, capacity and control. More than a traditional building management system (BMS) or a building automation system (BAS), Decathlon provides the visibility, decision support and centralized control technologies across a DC. Decathlon enables monitoring of system performance and environmental factors that would typically go unnoticed until it is too late. Decision support is provided in a context-sensitive manner, so that historical trends, forecasts, workflow processes, intelligent alarms and incident reporting provide information quickly for rapid response and resolution.

3.3.3. DCO Monitor/Collector **(revised DC functionality)**

In order to manage systems in the DCs in a smart optimized, automated and "near-real time" way, the DOLFIN expert system needs to have clear information of the topology data of infrastructures in DCs, the electricity consumption, infrastructure performance data of hardware infrastructures of the DCs constantly available, along with energy consumption of non-ICT systems such as the lightning subsystem, the HVAC cooling subsystem etc.

According to the DCO Monitor description introduced, the main purposes of this component are the:

- Collection of DCs performance data (usage of CPU, RAM at servers and VMs level)
- Collection of DCs electricity consumption data (at servers level and DC segment level)
- Collection of DCs temperature data (at servers level and DC segment level)
- Mapping of DCs infrastructures (servers and VMs topology)
- Provides information to DOLFIN modules for the necessary analysis

The monitoring system collects ICT infrastructure energy consumption in near real-time (e.g. every 5 minutes): CPU and RAM usage of the servers and VMs, the disk usage etc, electricity consumption and temperature of the servers and the DC segments. The DCO Manager/Collector does not provide data elaboration or data analysis but only the real-time and the historical “row” data injected by DCs. It also keeps track of the topology of the servers and the location of the VMs across servers.

The DCO Monitor/Collector needs to be integrated with the existing ICT hardware and software infrastructure by exploiting already available management systems/interfaces. In order to increase interoperability and usability, standard interfaces should be provided, like the Data Centre Manageability Interface (DCMI)[50] which derived from the Intelligent Platform Management Interface (IPMI) 2.0 [51]. In case of some proprietary interfaces, these can be addressed through the use of an adaptation block, which will convert any proprietary interfaces in just one common interface, based on emerging standards, such as Cisco’s EnergyWise Suite [52] or the IETF Energy Management (Eman) WG Interfaces [53].

3.3.4. Energy Information Base (new DC functionality)

The Energy Information Base (EIB) has the sole purpose of storing all data collected from DCs. As stated previously, the information coming from the ICT part of the DCs can be classified into the following main classes (from this moment for convenience we define them as entities):

- Server electricity consumption (TAB_POWER_USAGE)
- Server operating temperature (TAB_TEMPERATURE_USAGE)
- Server CPU usage (TAB_CPU_USAGE)
- Server RAM usage (TAB_RAM_USAGE)
- VMs CPU usage (TAB_VM_CPU_USAGE)
- VMS RAM usage (TAB_VM_RAM_USAGE)
- DCs topology (TAB_DC)
- Servers topology (TAB_SERVER)
- VMs topology (TAB_VM_RAM_USAGE)
- Services topology (TAB_SERVICE_DC)
- Services/servers topology (TAB_SERVICE_HW)

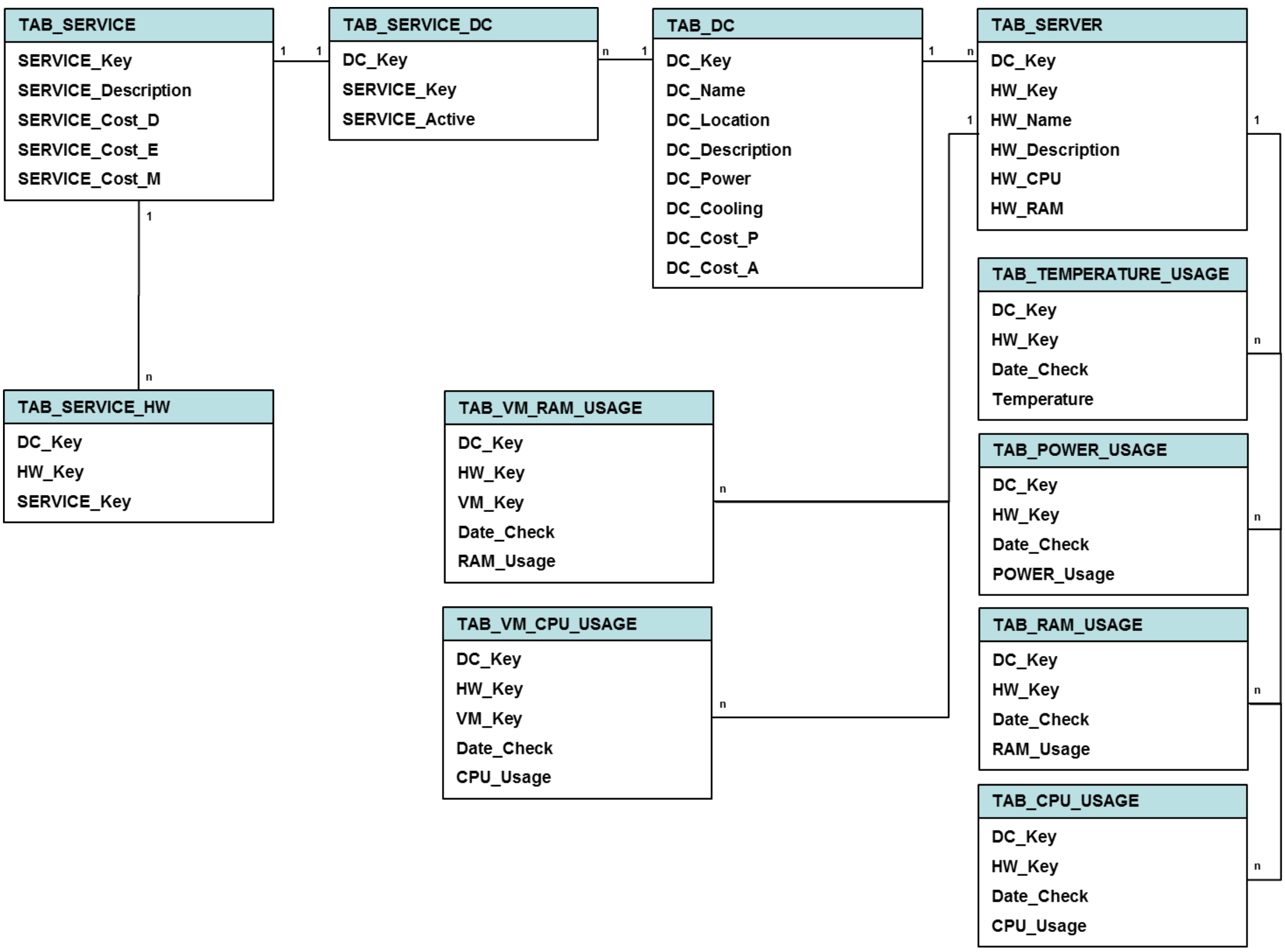

Figure 3-6: ICT only energy consumption ERwin Model

A similar list will be defined also for the energy consumption of non-ICT DC infrastructure. Each entity is stored in a specific table within the database and are related to each other following the ERwin data model shown Figure 3-6. All the information is collected by DCs every 5 minutes. Data retention policies provide a data historical depth of 15 calendar days.

3.3.5. Monitor Backend (revised DC functionality)

The Monitor backend can be considered as a set of dedicated interfaces available through SQL-Net protocol. These stored procedures will be available to the DOLFIN applications in order to get the basic information of usage of the hardware infrastructures deployed in the DCs and the infrastructure topology.

Store Procedures are developed to retrieve "near real-time" information or "historical data" and moreover guarantee all levels of aggregation and detail. As mentioned before, the Monitor Backend does not provide functionality for data analysis and/or data aggregation, but only for data retrieval and delivery.

The monitor backend may provide an interface to a ***Dashboard***, which will be a Graphical User Interface (GUI) acting as a monitor tool on top of the ICT Performance and Energy Supervisor system and providing an easy way to consult all the relevant information related to ICT Performance and Energy Supervisor.

3.3.6. ICT Performance and Energy Supervisor (new DC functionality)

The objective of this component is to provide information on the actual performance of the applications (typically VMs) utilizing the resources of the DC devices along with information on the energy consumption. Primarily, the module will receive as input performance utilisation data and energy consumption data collected by the DOLFIN DCO monitor/collector module (via the eCOP Monitoring DB) including CPU utilisation, memory consumption, link utilisation, energy consumed by the lightning of each DC segment etc. that are typically available for ICT devices, namely servers, storage equipment and networking devices and from the BMS systems and will perform an analysis to describe the utilisation levels of each device.

The key challenge of this module will be to devise an effective model capable of predicting the actual performance levels of applications as perceived by the user (e.g. throughput in operations/sec for compute intensive applications, response delay for interactive applications, bytes or frames/sec for streaming or media applications) based on raw performance data such as CPU utilization, memory consumption, network utilization, etc. Under certain circumstances, lower-level metrics may also need to be utilized relevant to the memory hierarchy of the servers (memory bandwidth utilization, cache misses etc.). Based on the module's analysis, the DOLFIN system will be able to detect suboptimal states of operation where excessive compute, network, storage and consequently energy resources are consumed for execution scenarios that do not deliver the desirable performance to the user. For DCs, where energy and power are not explicitly measured by the relevant monitoring devices, the Performance Supervisor will also provide information on these metrics as well based on models that correlate ICT resource utilization metrics with power consumption.

Finally, apart from studying, analysing and building models utilizing the available performance metrics, the module will also include a component that will be used for retrieving and efficiently storing such information. Storing of the utilisation and energy consumption information is necessary for two primary reasons. The first one is for maintaining the historic course of the data, so as to identify trends over different time and date periods. The second one is to utilise that data, at a later stage, for processing and further transformations. The structures that will be used for storing that data will be efficient enough to support fast search, insertion and retrieval operations. With this purpose, an information model will be defined to support the abstraction of heterogeneous resource information and characteristics including energy consumption profiles.

ICT Performance and Energy Supervisor utilises the eCOP Monitor Data Base and correlates it with real-time information of the DC energy consumption. Additional interface is needed with the Energy Policy Maker to provide real-time metrics about DC's energy exploitation. The ICT Performance and Energy Supervisor component computes and makes available this information that can be periodically collected.

3.3.7. Energy efficiency Policy Maker & Actuator (new DC functionality)

The central element of the DOLFIN infrastructure which makes decisions which are taken with the aim of improving the energy efficiency of the system is the Energy Policy Maker & Actuator. This component applies a set of well known criteria and evaluation patterns which are able to highlight conditions of sub-optimal energy exploitation and, consequently, can produce a set of operations meant to improve the DC's energy efficiency. Following the spiral approach of Figure 3-2, this component realizes the 3 lower parts of the energy optimization, achieving optimization at DC level.

The Energy Policy Maker does not directly enforce actions; instead it produces a stream of requests which can be translated into actual actions by other DOLFIN servant subsystems within a DC. Particular care must be taken in the design of the Energy Policy Maker to avoid it being over-optimizing as this could probe to too disruptive to the DC functionality. The Energy Policy Maker should then sanction corrective actions taking their tradeoffs and costs into account.

Examples of actions which could be ordered as part of an optimization policy are:

- ***Multiple-VM migration from sparse hosts across a DC toward a small number of tenant servers***, a strategy which aims at obtaining an higher density of sustained workload for a smaller number of physical server, powering down the remaining, now-idle ones;
- ***Offloading of a sizable portion of DC workload*** by the means of negotiation and transfer of VM groups toward a cooperating DC which needs an higher workload in order to improve its own energy efficiency;
- ***Negotiate a power supply strategy with the local smart grid***, in order to ensure access to lower power fees, or to achieve better balancing of the power grid;

The way the Energy Policy Maker works is by evaluating real-time metrics about DC's energy exploitation which are computed and made available by the ICT Performance and Energy Supervisor component. Knowledge of the structure DC and layout of the established services is also essential in this evaluation phase and it is provided by the ICT performance and Energy Supervisor. Another component whose status is also factored in by the Energy Policy Maker is the Smart Grid Controller, which is dealing with the local Smart Grid and/or Energy Providers. Its task is to sort and maintain information about the underlying energy provision, including energy fees and availability of local energy production, as well as providing feedback to the Smart Grid about the status and planed energy absorption for the DC.

In the case of single site DC, the algorithms to consolidate/deconsolidate workload locally in order to save energy need to take into account the trade-off between performance and the impact on QoS due to the energy saving actions. Energy efficiency can be increased by dynamically consolidating or spreading the use of the software on a minimal and efficient set of virtualized hardware resources – servers, storage and network components, which are still able to meet the SLA requirements. Unused servers can be turned off (or hibernated) to save energy. Some hardware elements such as CPU, disk, network devices, need to temporarily support a higher load, in order to reduce the number of physical servers needed. To cope with reduction of heat and Greenhouse Gas (GHG) emissions, VMs and services can be moved from DC segments with high load or high temperature to segments with fewer loads and lower temperature. For this, the thermal, workload and energetic characteristics of each DC site are important, in order to determine the time slot, regular peaks, and hotspot computing locations. An energy saving strategy that does not rely on workload relocation is the compression of input and output data used by data processing services.

Data relocation across different storage units is another class of policies potentially useful to improve energy consumption, by achieving a better data storage consolidation and allowing one DC to free up some of the units. One promising option to investigate is relocation from local data storage attached to a server (elastic service typology) to SAN (Storage Area Network) connected disk units. This alternative can be employed to turn off a server node at times when they don't contain used images or instances.

The measures, which can be employed by Energy Policy Maker & Actuator to reduce workload, energy consumption and emissions are:

- **Workload Optimization**: when a server is not servicing requests on its virtual machine instances, the server is either turned off or hibernated, within the limits allowed by the green criteria applicable for the DC.
- **Workload Consolidation**: rear time migration of all related tasks on the machines which are most frequently and most heavily used. This can be done in combination with Server Consolidation and with Virtual Machine Migration and/or Replication.
- **Workload Control:** rescheduling and redistribution of related tasks based on negotiations across the DOLFIN the Customer and the Energy Providers **based on Customers agreed SLAs**. Usually DCs offer services in full respect of acting SLAs (centred on KPIs such as performance and availability). It is possible to introduce the **Customer SLAs** concept, which allows Data Centres to cope with certain constraints in a regulated way. By this high energy savings can be achieved with a controlled impact on the performance and availability of the system, acceptable and sustainable from the ICT User's standpoint. In general it should be possible to reward the "flexibility" of the system using this new concept

Combining these energy optimization measures in an integrated and collaborative way within a DC leads to a situation in which not only the EP is responsible for detecting and reacting to peak energy demands or a sudden supply of renewable energy, but the DC and their collaboration play an increasingly important role in controlling the energy consumption and emissions together with the QoS.

When the Energy Policy Maker highlights a specific strategy which is potentially able to improve DC's energy efficiency, a list of the SLAs potentially affected by the proposed change is produced and the SLA Renegotiation Controller is interrogated in order to re-negotiate them. In case of success, and only at this point, the new policy becomes eligible for actual enforcement.

To summarize, interactions which the Energy Policy Maker may have with other components of the DOLFIN infrastructure are:

- **SLA renegotiation query**, a request which is sent to the SLA Renegotiation Controller in order to evaluate the feasibility of a potential optimization policy and green light it before actually enforcing it;
- Various status reports and inspections are requested at the **ICT Performance and Energy Supervisor**, which is queried to obtain the overall functioning of a DC, as well as retrieving its current or historical energy metrics data; Via the ICT Performance and Energy Supervisor, the energy efficiency Policy Maker & Actuator components forwards a set of action to the DCO Hypervisor and Appliance Managers, which typically results in server consolidation or VM migrations across servers or different DCs and changing of the non-ICT DC infrastructure status;
- The **Smart Grid Controller** is typically queried by the Energy Policy Maker in order to sense the cost-effectiveness of the current DC working conditions, or in order to spot possible forms of optimization which may lead to the exploitation of local energy production made available by the energy grid (with lower fees) or energy quotas which, if met, grant access to lower fees.

3.3.8. Cross-DC Monitoring Data Collector (new DC functionality)

In order to support the VM migration between federated DCs, while keeping the SLAs, the DOLFIN ecosystem needs to have excellent knowledge not only from the federated SDCs resources, but also from the network resources.

Existing monitoring systems such as Ganglia [41], [42], [43] and GridICE [37] have addressed monitoring of large distributed systems, but they do not address the rapidly changing and dynamic infrastructure seen in service clouds. A monitoring system for clouds needs to be for the whole of infrastructure and service management, and so it should cover SLA compliance, elasticity, QoS, etc. It is important to recognise that it is the monitoring mechanism that closes the loop from the initial deployment, through execution, and back to a Service Manager. The monitoring system is there to gather data from all the components within cloud architecture, and so monitoring is a fundamental aspect of a service cloud that is used by the infrastructure and for service management. The following section presents the UCL Lattice Monitoring Framework which we have designed and built for the purpose of monitoring dynamic environments such as service clouds [38], [39].

In many systems, probes are used to collect data for system management [40], [41]. In this regard, Lattice will follow suit. However, to increase the power and flexibility of the monitoring we introduce the concept of a data source. A data source represents an interaction and control point within the system that encapsulates one or more probes. A probe sends a well defined set of attributes and values to a data consumer at a predefined interval. The goal for the monitoring system is to have fully dynamic data sources, in which each one can have multiple probes, with each probe returning its own data. The data sources will be able to turn on and turn off probes, or change their sending rate dynamically at run time. A further useful facility for a data source is the ability to add new probes to a data source at run-time. By using this approach we will be able to instrument components of the system without having to restart them in order to get new information. It is also beneficial to interface with existing frameworks in order to collect data for Lattice. We would need these frameworks to fit in with the concept of data source and probe, and so they can be encapsulated with the relevant adapter in the implementation.

To meet all the criteria outlined requires careful architecture and design. Many monitoring systems rely on simple data transmission. From a design point of view, this approach is successful, although, we have found it is better if the monitoring framework design encapsulates separate planes for data, for meta-data, and for control. This allows us to build a system that has the desired behaviour and meets the requirements.

In Lattice, the separate planes for connecting the monitoring framework are:

- the ***data plane***, for distributing measurements from the Data Sources to the Data consumers.
- the ***control plane***, for distributing control messages to the Data Sources and the Probes.
- the ***information plane***, which holds all the meta-data relating to measurements sent by Data Sources and Probes.

These three planes are shown in Figure 3-7, together with the Data Sources and Probes, Data Consumers, and a Regulator. In a running system there will be multiple Data Sources, multiple Data Consumers, but only one Regulator. The regulator’s job is to ensure that the monitoring system does not flood the network or overload any applications. In many systems, the Probes collect data at a given data rate, and transmit measurements immediately, at exactly the same data rate. In Lattice we can decouple the collection rate and the transmission rate in order to implement strategies which aid in efficiency. For information in UCL Lattice please refer to [38], [39].

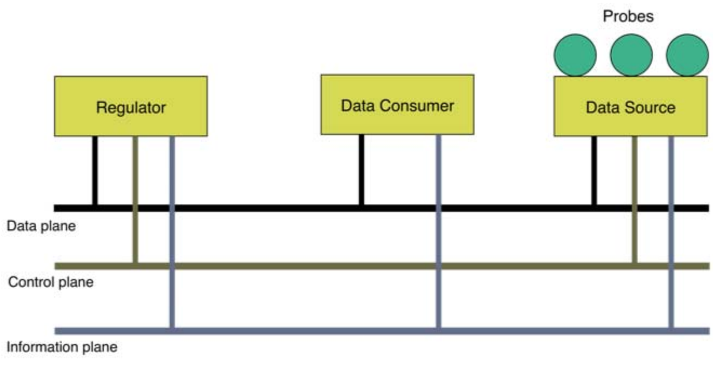

Figure 3-7: Lattice Multiple Planes

3.3.9. DOLFIN Information Data Base (new DC functionality)

The DOLFIN Information Data Base (IDB) offers abstracted and logically-centralised information manipulation (including information collection, aggregation / processing, storage / indexing and distribution) across all DOLFIN architectural components. Moreover, it includes information of the local energy requirements e.g. Demand/Response requests from the local Smart Grid operator.

It uses two separate interfaces for communication with the DOLFIN software entities: (1) the Information Management Interface is used for information manipulation configuration, including the information sources/sinks registration to the Information Base, the management of internal Information Base, information manipulation functions and the establishment, operation and optimisation of information flows; and (2) the Information Exchange Interface that offers the actual information exchange capability to the DOLFIN components. These two interfaces allow a unified yet abstracted information handling in all types of participating entities.

The DOLFIN IDB provides storage and indexing, along with processing functionalities. It maintains a registry, storing specifications for the available information to be collected, retrieved, or disseminated. Moreover, it may apply aggregation functions (e.g. MAX, MIN, AVERAGE) to the collected data before they are stored or disseminated or filter data at the aggregation level for optimisation purposes.

Last but not least, in order to manage systems in the DCs in a smart optimized, and automated way and as much as possible in "real time" and in order to achieve energy savings and the resulting C02 emission reductions or to meet the different demands from EPs in terms of optimization or different distribution of the power consumption of the smart grids, the DOLFIN expert system needs to have the topology data of infrastructures in DCs, the electricity consumption and infrastructure performance data of hardware infrastructures of the DCs constantly available. This info is also stored in the DOLFIN Energy Information Base.

3.3.9.1. Energy Information Base

The Energy Information Base is a logical database of the status of energy consumption for all elements in the DC at network, server and device levels. DOLFIN efficiently stores the performance metrics that describe the utilisation and energy consumption information of ICT devices. This way DOLFIN maintains the historic course of the data. It makes easy to identify trends over different time and date periods. In addition, it allows processing and data transformations as well information collection, aggregation indexing and distribution.

The structure used for storing that data will be efficient enough to support fast search, insertion and retrieval operations. Well-designed databases are easy to enhance, simpler to work with, and has fewer problems when you need to extract information. For this reason the database will meet the following characteristics:

- Self-describing nature of database system: the data base has to contain a complete definition of the structure and constraints, not only the database itself. The information is saved in the DBMS. It contains information about the structure of each file, the type and storage format of each data item, constrains on the data and the relationship between the data.
- Concurrency and consistency: The control of data concurrency and data consistency is essential. Data concurrency involves that many users can access data at the same time and data consistency implies that each user sees a consistent view of the data. In a multiuser database, the statements within multiple simultaneous transactions can update the same data and need to produce consistent results. Concurrent and consistent use of data increases the economy of a system. Data capturing and data storage is not redundant, the system can be operated from a central control and the data can be updated more efficiently.

- Data Integrity: The concept of data integrity guarantees that all data in a database can be traced and connected to other data. Enforcing data integrity ensures the quality and reliability of the data in the database. It includes the protection of the database from unauthorized access and unauthorized changes. Having a single, well defined and well controlled data integrity system increases stability, performance, reusability and maintainability
- Data Persistence: Data persistence means that the data is maintained as long as it is not deleted explicitly. The life span of data needs to be determined directly or indirectly be the user and must not be dependent on system features. Additionally data once stored in a database must not be lost. Changes of a database, which are done by a transaction, are persistent. When a transaction is finished even a system crash cannot put the data in danger.
- Data abstraction: It will offer an abstracted and logically-centralised information manipulation across all DOLFIN architectural components.

The Energy Information Base uses two separate interfaces for communication with the DOLFIN software entities:

- The Information Management Interface: is used for information manipulation configuration, including the information sources/sinks registration to the Information Base, the management of internal Information Base, information manipulation functions and the establishment, operation and optimisation of information flows
- The Information Exchange Interface that offers the actual information exchange capability to the DOLFIN components.

3.3.9.2. Energy Information Systems

DOLFIN keeps the DC performance monitored at all times. Hardware and software probes are deployed across the DC infrastructure and provide a continuous flow of measurement data which tracks the functioning of the services and ICT hardware at the highest possible level of detail.

DOLFIN uses this flow of data to produce and update various measurement metrics which can be better used to perform energy efficiency evaluations and to screen the DC in search of sub-optimal energy exploitation conditions. In order to produce the needed metrics, the ICT Performance and Energy Supervisor component perform what effectively is an operation of normalization and aggregation of the data stored in the DCO Monitor database. The task of the ICT Performance and Energy Supervisor is to extract all the relevant data from the DCO Monitor database, process it with various aggregation patterns and produce aggregated readings which briefly describe the overall DC functionality in simple form, at low detail level.

The extracted readings have multiple purposes:

- they are used as inputs for metric computation and updating. The various metrics produced by the ICT Performance and Energy Supervisor are the indexes which express in very clear, numeric form, various aspects of the current energy efficiency and effectiveness of the DC, and are the primary input factored in by the Energy Policy Maker component
- they can be stored for later reuse. These are not only useful to produce statistics and for historical reporting of the DC performance, as the Energy Policy Maker may also access this historic data to extract trend features, or evaluate potential plan-ahead policies.

The ICT Performance and Energy Supervisor exposes a query interface which allows easy and searchable access to various quantities. In particular, it exposes:

- access to the current computed metrics, and their historical data;
- access to the aggregated values which are inputs for the metrics evaluation, and their historical data;

- access to the un-aggregated data being reported by the monitored appliances, in a clear and normalized format.

3.3.9.3. Implementation of Energy Model

The DOLFIN system relies on a tree-like structure to organize the inbound measured data across a whole DC. This structure also plays the main role in aggregation and consolidation of the dataset, providing an unified view of all the measured attributes of the DC functions at different levels of detail; exploring the aggregated data neat the tree root (or the root) would provide a more high-level, generalized report of the status of a DC branch, or the DC as a whole, while exploring the aggregated data in a note proximal to the leaf level, would provide a detailed view of the working conditions and sustained load of an aisle, rack, single server or a specific VM. This of course also applies to non-ICT appliance such as HVAC, which can also be represented with their own class in the topology tree of the DOLFIN expert system.

The choice of the tree structure for organization and representation of the stored and aggregated data is functional to some objectives of the energy systems. First, it must be quick in propagating updates and new aggregated data upon its own structure. A tree structure representing the effective electrical placement and power distribution across a real DC provides the smallest computational challenge which completely solve this problem, with few exceptions. Moreover, it ensures which the aggregated data is, and stays, consistent at all times, while keeping the computational load low: it is in fact extremely easy to validate an aggregated data in in a node of the topology tree just by considering its descendent leafs. For this same reason, an updated measurement produced from a probe, can find its way in the aggregated measurements from its entry leaf (representing the reality at the highest level of detail) up to the topology root (representing the overall working conditions of the DOLFIN system) in very few steps, since only the upstream nodes form the leaf to the root, potentially, require an update.

This tree-structured information is not only retained in an organized form and stored, it must also be presented to other DOLFIN subsystem with an interface which is both secure, and easy to use. This query interface will have to provide a set of important basic functions:

Tree traversal. A set of functions which allow an external entity to dynamically explore the structure and the contents of the topology tree. This interface introduce overhead small enough navigation of the topology tree at considerable speed, avoiding the need for other consumers to rely on local caches and allowing them to always have the fresh view of the system;

- Stable watch functions. Other consumer components of the DOLFIN expert system must be able to instruct the manager of the energy model in order to have fresh data delivered to them directly, with zero delay lag, and with no overhead introduced due to active polling. The watch interface lets an observer to dynamically declare which are the data series of interest among those available in the topology tree. An active watch causes, at the update of the value of any topology data - resulting from a direct probed measurement or from a subsequent aggregation (existing in any level of the topology tree) to be notified to the consumer interested in it. This mechanism of course shall be robust enough for other consumers to rely on it, and must be implemented in a way which allows the manager of the energy model to scale indefinitely. This feature allows to completely avoid the active polling techniques which would not scale at this level of complexity, and allow the implementation of the consumers of the data to access fresh data in an extremely simple way;
- Node metadata. Not only probed measurements and aggregated values, but every node of the topology tree must support a hierarchical set of metadata which describe what the node refers to, how the measurements (or aggregated values) must be implemented and where they come from, as well as providing all the necessary information to reference the handled entity across other

components of the DOLFIN expert systems. For example, a VM can have its own node in the topology tree which stores the measurements relative to the VM work (cpu load, storage, network usage), whilst the node metadata describes the VM class, attributes, and references the VM record in the Info Base;

- Arbitrary, hierarchical node categorization and selection. While remaining strictly a tree structure which allows for strict validation and organization of the data, each topology node can be categorized arbitrarily, making it part of other data sets, by the controlled addition of named hierarchical tags. There is no limit to the number of tags which a node can carry, and any combination of these can be used in conjunction with the stable watch functions to obtain views which would normally not be provided by the structure of the topology tree. For example, a DOLFIN component may be interested in finding out the current power expenditure of all the PSU units across a specific datacentre. Being, typically, the PSU units layered horizontally in a DC power distribution network, there would be no easy way to simply declare a stable watch against a specific topology node to get this information. Instead, a consumer system could rely on the tagging provided to the topology nodes referring to PSU units to locate them and put them under watch;
- Historical data reporting. Likewise as the stable watch functions, a similar set of functions must also be provided in order to allow other DOLFIN system components to retrieve and explore the historical data of specific measurements, or aggregated values. These functions must allow the consumer to specify a time window and an optimal sampling resolution, in order to retrieve a complete, or incomplete (in case of missing data) time series containing the requested data. It is necessary to note that not all measurements and aggregated values are subject to long-term historical archival, as this could have very little benefits compared to the problem this constitutes in terms of storage utilization and I/O. Instead, a specific set of meaningful quantities and aggregated values is elected for long-term storage archival, which must be configured upfront.

3.3.10. Cross-DC Workload Orchestrator (new DC functionality)

The Cross DC Workload Orchestration is a software element which does most of the different types of resource optimisation including energy optimisation and management of their trade-offs at cross DC and network level. It is in charge of managing the full lifecycle of the virtual routers in the network and the allocation of the applications running on the virtual nodes. Moreover, it closely collaborates with the Monitoring Data Collector, which collects and manages all of the monitoring data received from the DC to DC network, including virtual and physical resources.

The Cross DC Workload Orchestration component is responsible for the placement of the virtual routers, the allocation of service components and functions running on the routers as applications. In particular, it is responsible for instantiating the virtual routers as it is the software element which does most of the management and orchestration and is in charge of managing the full lifecycle of the virtual routers. Mainly, the Orchestration component has the following functions:

- It starts and stops the network controllers on each physical machine;
- It acts as a control point for the platform by sending out the command and,
- It acts as a management element for the platform by collecting monitoring data and enabling reactive behaviour.

The Placement Engine of the Orchestration component is the module in charge of performing the actual placement of the virtual routers according to the initial topology and the usage of the virtual network elements. This is an important feature because, when we configure a network, considering some initial information, some of these parameters may change during the course of the system's operation and a re-configuration may be required to maintain optimized collection of information. For this reason, our approach considers a mechanism to achieve adaptation in a flexible manner. The decision on the Placement Engine is encoded in an algorithm which can be rather simple, such as counting the number of virtual routers on a host, or it can be based on a set of constraints and policies that represent the network properties. In this work we have considered i) infrastructure based measures for the placement of the virtual nodes and ii) as constraints, the usage of the virtual network entities (in this case, the usage of the virtual links) for placement. The results of using different Placement Engines will be highlighted later in this paper.

Each virtual router has a probe to monitor the usage of the network resources (e.g. the state congestion of the links). The data provided by the probes is collected by the Monitoring Data Collector and stored in the DOLFIN IDB to create or remove the virtual routers according to the current state of the network.

3.3.11. Cross-DC VM Manager (new DC functionality)

The Cross-DC VM Manager, extends the role of Energy Policy Maker & Actuator in multi-site federated DCs. Efficient energy management is considered in the DOLFIN Ecosystem of Synergetic DCs as a distributed constraint satisfaction problem: "how to distribute the temporarily higher workload within the DC Federated nodes given the current SLA constraints of each node, while also trying to shape the energy consumption of each node and the energetic and computational stability of the DOLFIN ecosystem as a whole". For this, the energetic and quality impact of a service is measured and taken into account, and explicit relations are maintained between control actions and their energetic costs with respect to supporting resources. This helps to keep the services operational and to meet their QoS target parameters.

To cope with reduction of heat and Green House Gas (GHG) emissions, VMs and services can be migrated from DC with high load or high temperature to DCs with fewer loads and lower temperature. For this, the thermal, workload and energetic characteristics of each DC site are important along with the network speed and load, in order to determine the time slot, regular peaks, and hotspot computing locations. An energy saving strategy that does not rely on workload relocation is the compression of input and output data used by data processing services.

The major differentiating factors that influence the level of collaboration between the member nodes of a SDC are:

- **Type of Managed Services**: when the DCs are compliant to the same technological level, the type of managed services provided (e.g. IaaS, SaaS, PaaS services) is a strong influencing factor in the cooperation schema. It is more likely that providers of the same types of services are to cooperate, as their organization and the types of services provided are more compatible with each other and they could trust each other
- **Administrative and Organizational Topology**: The administrative boundaries and organizational topology of the DC Federation (DCF) is important in determining the coordination mechanisms that can be employed to shape and optimize workload. By organizational topology we understand the level of organization provided by the DCF and the distribution of nodes across administrative boundaries.
- **Openness and Persistency**: The openness and persistence of relations in the DC Federated is another determining factor in establishing the available actions at disposal of a DC. DCF workload management based on dynamic coalition formation could be more dynamic and exploit the heterogeneity of the DC, but might incur more coordination overhead, while the more static, closed-garden type of DCF can exchange more information and could cooperate more intensively in order to maintain their QoS. According to our experience this model requires a stronger cooperation; in the DOLFIN project, we don't explore this kind of theoretical possibility to use the DC federated where

the persistency of relations influences the trust that emerges based on repeated experience, and typically leads to a situation in which controls and formal agreements between parties can be relaxed.

- **DC Workload Management Approaches**: Companies use virtualization to solve infrastructure consolidation, to reduce downtime of servers and to improve business continuity and disaster recovery contingency, and to create more efficient development and test platforms. As workloads become more and more virtualized, and use newer and more capable hardware, resources can be dynamically allocated based on real-time and power-aware analysis, so that "green criteria can be met. Virtualization technologies (e.g., VMware vSphere, Citrix XenServer, Microsoft Hyper-V, etc) enable a dynamic environment where VMs can be moved and load balanced across the physical infrastructure. This provides the basis for a solid and redundant shared storage environment and a predictable high performance network. Building an IT Infrastructure geared to delivering the SLAs to the business requires a tightly coupled set of infrastructures, management tools, and operating environments. These must be coupled with policies, data management, storage, security and network control required to dynamically allocate resources based on application needs.

The potential for energy consumption reduction through improvement of IT control, particularly through more power-aware applications, can be realized through several alternatives, as indicated in the following sub-sections.

Existing DC actions for running workload externally, in another DC, are:

- **Consolidation/deconsolidation of workload capacity.** The main workload consolidation/ deconsolidation approaches which are involved in the context of resolving resource allocation requests, i.e. VM data and task allocation on remote nodes are Server Load Consolidation (SLC) and Network Traffic Consolidation (NTC). SLC is an essential mechanism for energy efficient workload management. To facilitate optimized workload management and to minimize energy consumption, one needs to take into account the existing network traffic and server node topology of a given DC. After performing server load consolidation, one can further reduce energy consumption by resorting to NTC as well, but this is outside of the main DOLFIN project scope. An indirect way to consolidate network traffic onto fewer links and to allow the controller to turn off non-utilized ports (and switches) is to migrate jobs, such that a fewer number of servers are being used with less network traffic. To satisfy such constraints, one needs to ensure that other server resources, such as CPU and memory, are adequate to handle the assigned jobs. Consolidation of server load or network traffic has benefits such as energy savings, but also disadvantages, such as a temporary reduction of availability and reliability

- **Delaying, shifting and prioritization of tasks.** Another approach to run workload externally is re-scheduling of workload across DCs, by using a common (grid-like) workload allocation/reservation mechanism through which planned computing tasks are prioritized and scheduled based on their impact on the overall DC Federated energy consumption. We consider several options for shaping workload, energy consumption and emissions through collaboration between the individual DC of a DCF taking into account the Customer SLAs: in this way the possible solutions are based on prioritization and shifting of tasks to a later point in time. This so-called task shifting is one of the dominant approaches for cases when the DCF governance applied in the nodes is too different to allow interoperation, or when the VMs dependencies on other VMs, i.e. data and/or custom services, are too complex.

- **Task relocation for not started tasks.** Relocation of tasks is aimed at maintaining the QoS and / or reducing energy consumption and / or emissions. In case of DC server loads above the maximum load considered safe for the current DCF node, all tasks that are scheduled with high priority, but have not started yet can be relocated to other available nodes of the DCF. Their relocation needs to be preceded by a negotiation and scheduling step agreed with a remote DCF node.

- **Relocating to 1 other DCF node:** This is the situation when one task is completely migrated to another node. The remote node provides sufficient capacity for the task to be executed. It still has to be taken into account whether the task is relocated before or after it has been started. In case the task already started, the possibility to suspend the task, to create a restore point and to make a snapshot of the current state, which can be sent to the remote node, needs to be taken into account.
- **Relocating to n other DCF nodes:** This is possible through negotiation, when a task is larger than the available workload capacity that can be outsourced to any given node of DCF. In that case, the task needs to be split into multiple smaller chunks, and a more complex orchestration between these chunks needs to be planned. The task itself, or each different chunk of it, is distributed to the different DCF nodes found available and with sufficient capacity, and it is run remotely within the resource limits provided by each node. There are issues associated with passing of the intermediate results between the nodes or with passing of the overall task result back to the original node on which it would have started, should enough workload capacity have existed. This can impact the service levels of the SLA, due to the uncertainty about the available bandwidth to transfer data, and about the data replication and synchronization issues for the task dependencies and data to be processed by the task.
- **Task migration for started tasks.** This is aimed at minimizing the impact of task migration on the current QoS of each DC. Migration of an already started task requires several conditions: (1) the possibility for the task and for the execution framework to have control points in which the task can be suspended or stopped, its current state saved, and its resources freed; (2) the possibility to negotiate, within a reasonable amount of time, of a new location of computation, i.e. a new VM image on which the suspended task can run; (3) the possibility to transfer the task to the remote node, together with its task dependencies, and to resume execution within the limits allowed by the SLA. Note that in most cases the task can be suspended only after a suitable candidate solution has been found for it, either shifting/re-scheduling or task migration.
- **Workload outsourcing of non-critical tasks.** Workload outsourcing of non-critical tasks refers to migration of non-critical tasks (e.g., periodic security checks, monitoring reports, backups and other tasks scheduled to run periodically) that should be run in a specific DC server in a critical load period, to other servers of the same DC or (for a small class of tasks) even to other federated DCs that can accommodate these tasks. Under some conditions, this can impact the SLA due to time to migrate/transfer or due to other end user tasks that depend on the administrative services migrated. This flexibility depends on the granularity and flexibility provided by the signed criteria which formally define how Data Centres can federate and collaborate to exchange load among them.
- **Workload outsourcing in Traditional DC Federations.** The workload outsourcing in this case follows the steps encountered in workload balancing for application servers. This implies that a load balancer exists, which monitors the load status of two or more nodes, and if the load on the primary server exceeds a certain threshold value, then it assigns the new workload to the other nodes which have more available capacity. The assignment follows the policy specified in the criteria concerning the conditions and rules for migration of workload. The traditional aspects will be taken into account but always considering the possibility to deal with all these elements in the other cases.
- **Workload outsourcing in Cloud DC Federations.** Several mechanisms exist for service provisioning in a CDCF of which the main ones investigated are: ***a) Dynamic selection of virtual availability zones by DC***, a simple mechanism by which not the end customer, but an automated DC determines the location of computation for any given task. This mechanism requires rather homogeneous entities, managed by the same organization. It also assumes that data stored is stored in volumes, racked up, and replicated across the federated DC nodes and **b) Cloud bursting**: an on-demand reservation of supplementary workload, which is done by computation offload; a DC creates local resources as required by its ITC needs. When the allocated resources prove insufficient for the local demand, the environment supporting cloud bursting ensures that additional resources are allocated. During peak

energy and workload, insufficient resources will need additional coordination, in the limits allowed by the SLA. On the one hand, the cloud bursting mechanism is more permissive, as it allows more heterogeneity of DC: different cloud, different owners, heterogeneous configurations of the VMs are possible. On the other hand, the workload transfer from one DC to another requires a more complex "workload transaction": Configure/Convert VM in cloud A; Transfer the VM from cloud A to cloud B; Import the VM into cloud B. These operations have costs in terms of duration, energy consumption, financial cost and impact on service availability.

The cross DC VM Manager will be based to perform the actual VM migration on the DC Interconnect (DCI) framework. CISCO was among the first to present a complete DCI platform to support the contemporary networking needs of the DCs to properly support their cloud capacity in a multi-cloud environment, followed by many other vendors such as HP, BTI, Brocade, TelX [55]. Among the most commonly used protocols as Layer 2 extenders are (Advanced) Virtual Private LAN Services (VPLS) [62][63], Overlay Transport Virtualization (OTV) [64] Ethernet over MPLS (EoMPLS) [59]. Recently, IETF proposed alternatives named Ethernet VPN (E-VPN) [60] and Transparent Interconnection of Lots of Links (EPLS/TRILL) [61]. HP has also proposed an alternative Layer 2 DCI interface to CISCO's OTV, named Ethernet Virtual Interconnect (EVI). OTV provides both Layer 2 and Layer 3 connectivity using the same connections, also applying services for native Spanning Tree Protocol (STP) isolation. Moreover, it is optimized in terms of ARP caching, making it an overall faster and more reliable DCI solution than VPLS. IETF's E-VPN offers Ethernet connectivity among data centres spanning metropolitan area networks (MANs) and WANs assuring that one VLAN will be available for each MAC VPN, in combination with automatic route distinguishers. As E-VPN bases its operation on the hierarchical utilization of a number of MPLS edge routers, E-VPN could be described as implementing OTV over an MPLS network. TRILL, another IETF standard, creates a cloud with a flat Ethernet address, so that nodes can move around within the cloud and not need to change their IP addresses. Although nodes attached to the cloud perceive the cloud as an Ethernet while the packet is traversing the cloud, it is encapsulated with a TRILL header, which like a Layer 3 technology, contains a source (ingress RBridge), destination (egress RBridge), and hop count. HP EVI is a MAC-over-GRE-over-IP solution. Ethernet frames are encapsulated into GRE/IP at ingress to the switch. The GRE/IP packets are then routed over the WAN connection between the data centres. EVI's operation is very similar to the one of OTV supporting, however, larger DC and VLAN federations and offering support for multi-pathing and load balancing. However, in contrast to OTV, it does not explicitly support multicasting.

DOLFIN will analyse the available solutions, including features of the OpenStack and realise the DOLFIN DC to DC Manager Interconnect, measuring the energy consumption during VMs migration. Yet:

- The ***migration of VMs from one DC to the other is problematic***, as it requires high bandwidth, dedicated lines, compatible service catalogues (i.e., shared images) or shared data storage, all of which are uncommon in DCFs managed by different organizations.
- ***Live migration of tasks is ever more demanding***, as it requires shared storage source and destination. The migration does not cover the VM image, but only the content of memory is stored by the source node and transferred to the destination node, after an exact VM image replica has been started on the destination node. To cope with SLA requirements, very high speed (1 GB) network connections are required, and the same hypervisor program, running the same version on the source and the destination node. This is more likely to happen in single-owner DCF than with multi-owners, no shared storage. Therefore we can conclude that live migrations are very unlikely in heterogeneous federations, because the required conditions - same version of virtualization software, stopping of existing VM, transfer of memory contents and re-start of the task on the destination VM, which is an exact replica of the source VM -, are very difficult to meet in practice. It will be very interesting to test these features using the Wind test plant capabilities.

- ***Transfer of workload is not only subject to user security and confidentiality requirements***, specified as part of the SLA, but also to data policies and other legal constraints imposed by state actors, commercial organizations or international bodies (e.g., EU Privacy laws).

Given the current constraints on migration of VMs/tasks, DOLFIN will always assess first the possibility to save energy in the local DC through optimization and task shifting, before the possibility to run workload in another DC is considered.

3.3.12. SLA Renegotiation Controller (new DC functionality)

The basic idea of the SLA Renegotiation Controller is to use the technical specifications and guarantees in the SLA concept to steer the data centre operation towards environmentally sound behaviour.

As a starting point for the more detailed analysis we take into account that in DOLFIN a possible SLA between DC and End User will offer an extended scope of eco-optimization to the service provider by relaxing traditional performance parameters, introducing novel energy performance parameters as classifying elements, introducing parameters that determine the level of collaboration and offering incentives to the customers in the form of dynamic pricing schemes that are tied to the level of collaboration, the achieved system-wide environmental benefits, and the achieved reduction of energy consumption in the DC.

This definition has a set of implications for the analysis in this context, so we need to:

- Define the nature of the trade-off between performance degradation and energy savings as a basis for the compensation
- Define energy performance parameters that are relevant for the EP motivated version of a new green criteria to evaluate the quality / level of the service
- Suggest a dynamic pricing scheme as an incentive for the DC customers

Especially in case of cross-DC optimization, the contracts signed between DC are similar in nature with the traditional SLA. The QoS specified for workload outsourcing depend on the categories and range of services provided, the energy sources available at each DC's disposal, the QoS agreed upon with the counterparts, and on the amount of services contracted. Big customers have therefore more priority for handling requests, even though the requests are of the same nature. This implies that the workload outsourcing requests coming from "larger" customers or having higher QoS targets will need to be treated with more priority. This prioritization mechanism introduces additional constraints on the order in which jobs can be scheduled, and needs to be taken into account by the DCF when computing the alternative solutions for scheduling the jobs.

The SLA Renegotiation Controller, besides SLA criteria, we'll also try to comply with constraints derived from the **Green criteria.** The Green "criteria" between DCs and customers specify general rules for relocation of services, in case of accidents. Example of a rather technical SLA statement:

"Data needs to be constantly aligned between the Server Farms of the infrastructure provided by DC to the customer."

As this contractual requirement is translated into an architectural decision, the statement can become:

"Storage to Storage replication needs to be activated, with an active/passive configuration, without new servers or other components. The monitoring solution is configured to automate data movement failover and to send an alarm to the central monitoring systems."

This gets further translated into a set of capabilities for the DC situation taking into account that the offered SLAs are depending on the realised architecture:

- DC supports storage replication
- DC storage configuration is active-passive

- Data movement is automated
- In case of failure of data movement, an alarm is sent to the central monitoring system

Examples of constraints derived from Green SLA with end users include typically QoS constraints, such as:

- ***Computation location constraints:*** *"VM @VMx should run in DC @DCy located in region Y"*
- ***Computation amounts:*** "The maximum number of instances of VMs from customer Y that are allowed to run on DC Z is X"
- ***Time constraints:*** computation is restricted in terms of availability, performance or Service Management
  - *Availability: "VM should be available within working days in region Y: 08:00 – 18:00 (of time zone UTC+1)"*
  - *Performance constraints:*
    - *"VM should be available (started) within xx minutes from any request to use it"*
    - *"VM machine should not be too loaded, i.e. its minimum Virtual CPU speed needs to be above the @threshold value."*
  - *Service Management constraints (QoS): "Incidents reported about VM should be solved within X minutes from submitting the request to solve the incident"*
- ***Greenness constraints:*** "Green tasks" are tasks (with lower priority) to be run outside peak load periods

3.3.13. Customer profile management **(revised DC functionality)**

DC and cloud infrastructure users exhibit a wide variety of profiles and requirements as they originate from different sectors such as industry, academia, education, research and are widely geographically and culturally distributed. To that extent, the customer profiles are greatly influenced by the technical requirements of the services requested by the DC of cloud infrastructure, but also by the financial, societal and even moral status of the customer. The following service characteristics and their specific prioritization for each individual customer constitute a solid base to build customer profiles relevant to DOLFIN innovation.

- **Cost:** A typical requirement for service users is to maintain their operation cost a low as possible. In several cases, when the rest of the characteristics are of low importance, cost can be a dominant factor in a customer's profile.
- **Performance:** In performance critical applications like HPC computing or low-latency computing performance in terms of throughput (operations/sec, requests served/sec) or latency (time to respond) can be very critical so that users request the highest possible sustained performance.
- **QoS:** Quality of service is of ultimate importance in several business scenarios where the provided service needs to be as stable as possible, minimizing threshold violations regardless external traffic or load.
- **Security:** Security can be for utmost importance for several classes of users and applications that need to have strict guarantees for business critical or personal information.
- **Environmental friendliness:** In several cases, users tend to be quite sensitive about environmental matters, especially when they are aware of the energy and CO_2 footprint of modern data centers.

3.3.14. Smart Grid Controller **(new DC functionality)**

The Smart Grid Controller (SGC) acts as a broker between the core DOLFIN infrastructure and subsystems, represented by the Energy Efficiency Policy Maker & Actuator, and the outside world, namely the Smart Grid (SG) and the Energy Providers (EPs).

The SGC undertakes to support the DOLFIN System awareness of its energy surroundings, namely energy prices and SG status in terms of energy availability and (smart) energy grid stability. In this framework, it is responsible for assuring that the current and near future energy consumption profile of the DC are known to the SG and the EPs at all times. Last, it is responsible of informing the SG and the EPs of the availability of the DC's own energy resources, if any (e.g. in the case the DC is supported by one or more RES). The information from the DC side is provided by the Energy Efficiency Policy Maker & Actuator. The same holds for the information originating from the SG/EPs side; all is redirected to the Energy Efficiency Policy Maker & Actuator.

The data that are relayed to the Energy Efficiency Policy Maker and Actuator and the SG/EPs are lightly processed by the SGC in order to detect possible suboptimal situations that could lead to either less cost- and energy- efficient operation of the DC, or less stable SG operation, respectively. In that case, the interested party is responsible of issuing a proper action plan to optimize its place in the energy consumption/production chain. The SGC is not responsible for issuing such a plan; instead, it is responsible for communicating this information between the interested entities.

In short, the main tasks of the SGC are to:

- Handle information related to the DC energy consumption and relay it to the SG/EPs;
- Communicate to the SG/EPs information related to the DC's own energy sources, if available;
- Handle information related to the energy tariffs from the SG/EPs side.

It should be noted that the SGC should be in place to support more than one communication protocols, in order to guarantee reliable communication with the set of cooperating SG Operators and the EPs. From an architectural perspective and to fulfil all the requirements accompanying the aforementioned desired functionality, the SGC could be modelled as depicted in Figure 3-8.

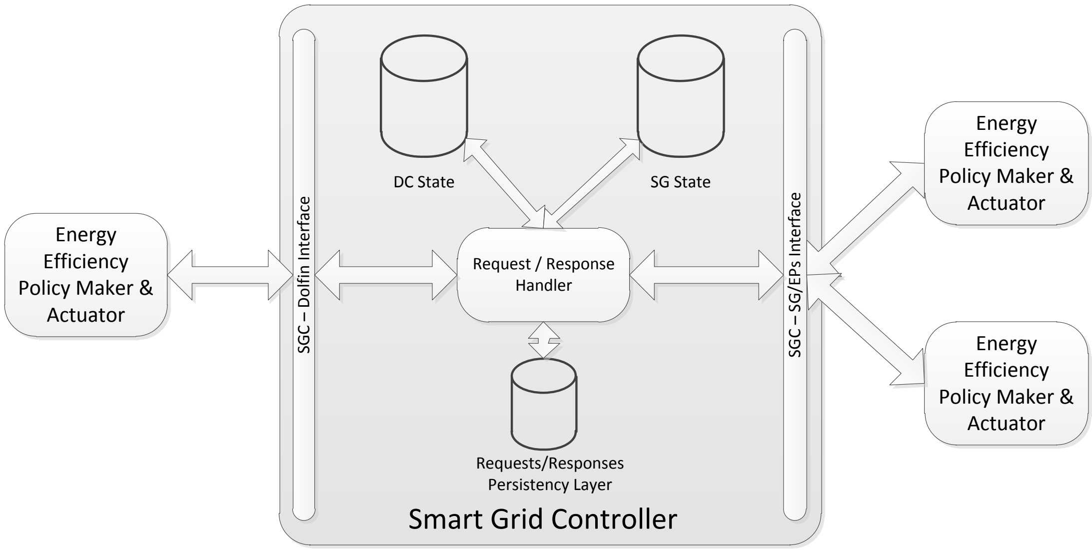

Figure 3-8: The SGC high level architecture

Every time a request is received by the SGC, it is translated in an internal SGC format (through the SGC-DOLFIN/SG/EPs interfaces) and is stored into the Requests/Responses Persistency Layer. Then, the request is retranslated and forwarded to the interested entity. Upon reception of an answer, any information related to the DC/SG state, if any, is stored in the appropriate registers. Note that the DC state may comprise information related to its current and near future computing load, the availability of its own energy sources (if any), whereas the SG state may include data relate to the current energy tariffs or their trend, the availability of energy in the SG along with some grid stability metrics indicating possible instability in the near future. Granted the previous DC/SG states, the SGC will evaluate whether an alert should be sent to any interested entity (DC/SG/EPs) along with the set of appropriate state information in order to render the receiver capable of performing sets of corrective actions to improve the DC cost- and energy-efficiency (from the DC side) or the SG stability (from the SG/EPs side). In any case, the SGC acts as a mediation layer between the Energy Efficiency Policy Maker & Actuator and the SG/EPs side, also providing minor caching services (regarding the DC/SG states) in order to limit overhead, unnecessary traffic to/from the DOLFIN core internal ecosystem.

3.3.15. Cooling Subsystem (Legacy DC functionality)

Traditionally, in most DCs, cooling at the server level is air based. Air cooling is simple and safe. On the other hand, air is not thermally efficient and thus impractical for cooling high-density implementations. The most attractive alternative to air cooling is liquid based cooling. One of the main draws of liquid cooling is the greater ability of liquids (relative to air) to capture and hold heat. Thus, a smaller amount of water can accomplish the same heat capture and removal as a relatively large amount of air, enabling a targeted cooling strategy. For high-density implementations, air cooling is often simply insufficient. In such cases, water (or, more generally, liquid) cooling offers an alternative in which not only can greater cooling efficiency be applied, it can be applied only where it is needed.

Liquids can serve as a medium for transporting heat in a number of different ways, ranging from broader cooling of the entire computer room to targeted cooling of particular racks/cabinets or even particular servers. Thus, we can categorize liquid based cooling systems to the following broad categories:

- **Computer-room liquid cooling**. This is the most basic option and is currently widely used in existent DCs. This approach involves the use of water for moving heat from the computer room to the outside environment. Computer room air handlers (CRAHs) use chilled water to provide the necessary cooling; the water then moves the collected heat through pipes out of the building, where it is released to the environment. This cooling approach is similar to the use of computer room air conditioners (CRACs).
- **Rack liquid cooling**. A more targeted approach involves supplying cool water to the rack or cabinet. This approach is similar to the whole-room case using CRAHs, except that only small spaces (the interiors of the cabinets) are cooled.
- **Server liquid cooling**. An even more targeted approach where liquid cooling is targeted to the servers themselves (CPUs and respective hardware components). Since hardware components operate at a much higher than ambient temperature, the liquid (water) used for transferring the heat away from the servers need not be cool. Actually, water as hot as 40 degrees or more can be used efficiently for cooling down the server components.
- **Liquid immersion cooling**. At the extreme end of liquid cooling are submersion-based systems, where servers are actually immersed in a dielectric liquid (i.e. refined mineral oil). A variation is to use the dielectric liquid inside the server only, avoiding the need for vats of liquid.

Depending on the particular implementation, liquid cooling obviously poses some risks and other downsides. In cases where water is involved, there is the risk of a water leak, which could be disastrous. Furthermore, water pipes can produce condensation, which can be a problem even in a system with no leaks. And with more-stringent requirements on infrastructure comes greater cost: infrastructure for liquid cooling requires greater capital expenses compared with air-based systems.

Liquid immersion cooling on the other hand, except from the high initial installation costs, there are also concerns on the maintenance of the servers and on the poor space utilization (since the “racks” are actually lying on their backs).

Within the project we will develop a “***water cooling server***” prototype. The prototype will consist of a water circulation system that will bring water directly to the server hardware components (CPU, GPU, memory, etc.). This system will communicate with the DC HVAC system in order to reuse the harvested thermal energy for various heating purposes (building heating, hot water system, etc.). The following figure depicts a rough diagram of the aforementioned prototype.

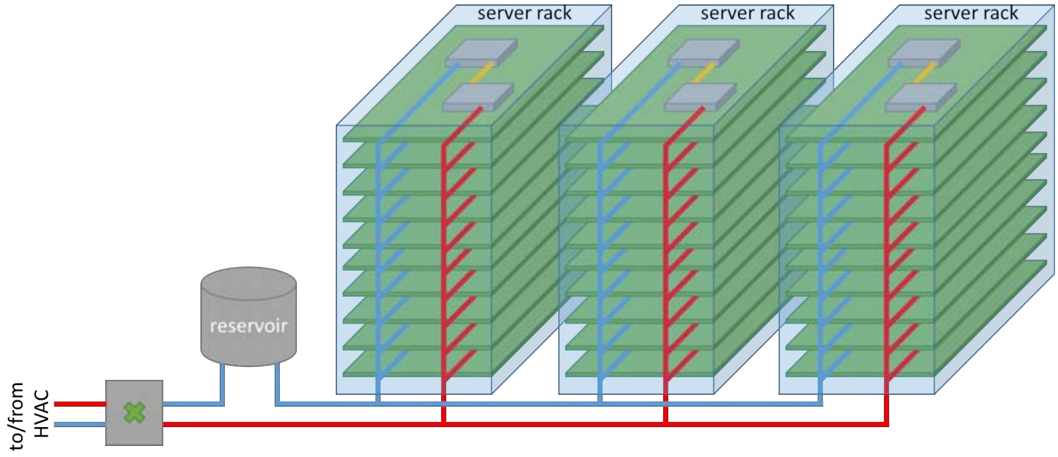

Figure 3-9: Water-cooling system diagram

The main constituents of the system will be the water-based coolers that will sit on top of the various hardware components, the water piping system, pumps that will circulate the water, and a reservoir that will be responsible for filtering and levelling the circulated water. Finally, the system will be interconnected to the building HVAC system.

The water based coolers should be a drop-in replacements for current air coolers for CPUs so that the installation to an existent system would be easy. Great care should be given to the water piping system and joints in order to eliminate the chance of water leakage and condensation. The pumps incorporated for the circulation of the water should be chosen with power efficiency and operational redundancy in mind. The reservoir, as in similar systems used for conventional heating and cooling, is required to maintain the correct amount of water inside the hydraulic system and also to remove air-bubbles that form. However, in the particular case, a careful filtration system may also be required to avoid mineral build-ups and algae formations in the long run.

Finally, the system must be interconnected with the DC HVAC system. The system will utilize warm water for cooling the processors. Although this seems oxymoron, it has been shown [55], [63] that water temperatures of 40 degrees Celsius and higher can efficiently be used for the cooling of servers, when the water reaches directly the hot hardware components (CPU, etc.) of the server. In addition to not having to chill the circulating water, which requires power hungry chillers, the returning water is hot enough to be used for various practical purposes (i.e. building heating).

DOLFIN water cooling subsystem will be evaluated in terms of power savings in typical usage scenarios. Deployment costs, in new and existent DCs will also be evaluated for various server densities and topologies. The system should incorporate management functions and processes so as to guarantee flawless operation in the long term and also diagnose possible failures at an early stage.

4. Testbeds in the DOLFIN project

This section presents the planned testbeds with the characteristics and functionality (e.g. 3 enterprise data centre test beds: WIND testbed, IRT testbed, GRNET testbed + 2 research data centre testbeds: UCL test bed, i2CAT testbed) aimed at interoperability and enhancing the DOLFIN architecture and its components.

4.1. Introduction

Data Centre - Large data centers are industrial scale operations using as much electricity as a small town and sometimes are a significant source of air pollution in the form of diesel exhaust. A data center is a logical centralized physical and / or virtual repository for the computation, storage, management, services and dissemination of data and information organized around a particular body of knowledge or pertaining to a particular business. The software of the data centres are realised and based on computing cloud principle and technologies. In general the hardware of the data centres would include redundant server and network nodes capabilities as well as backup power supplies, redundant data communications connections, environmental controls (e.g., air conditioning, fire suppression) and various security devices.

The following is a short overview of the main types of big data centres and their owner/operator/end users relationships:

- **Enterprise data centre:** In this case the main function is the delivery and management of services to the enterprise's employees and end users. In this case the owner, operator and end user of data centre is the same organisation, bearing all energy cost and having access to all relevant energy efficiency and environmentally relevant data.
- **Co-hosting data centre:** In this case multiple end users are provided with access to networks, servers, and storage equipment on which they operate their own services/applications. Both the information technology equipment and the support infrastructure of the building are provided as a service by the data centre operator, who bears all the initial energy costs, while users pay indirectly, depending on their contracts/tariffs, which are not related to energy consumption and often are flat rates. Energy efficiency and environmentally relevant data is available at the same organisation.
- **Co-location data centre:** In this case multiple carriers create a carrier dense network for end users to locate their own servers and storage equipment in a common structure. The support infrastructure of the data centre (i.e. power distribution, security and environmental control) is provided as a service by the data centre infrastructure operator, who bears all the initial energy costs. End users pay energy costs to data centre infrastructure operator, based on their contract which include actual energy consumption and a possible fee related to the additional energy costs such as cooling systems, UPS and other losses. Energy efficiency and environmentally relevant data is hence spread across different actors.
- **Network operator data centre:** In this case the main function is the delivery and management of broadband services to the operator's end users. The data centre operator bears initially all energy costs and the final users pay indirectly, depending on their contracts/tariffs, while these

are not related to energy consumption and often are flat rates; similar to "Co-hosting data centre". Energy efficiency and environmentally relevant data is spread across different actors.

- **Research Testbed data centre:** in this case this case the main function is the delivery and experiment with new management approaches which integrates energy efficiency of DC services to support the research and development community.

As there will be different types of big data centres in DOLFIN, a set of different solutions will be used. The following Data Centres infrastructures are planned for use as testbeds for the energy experiments and prototypes envisaged in the DOLFIN project.

4.2. WIND TestBed

4.2.1. Description of Wind Testbed

One of the DOLFIN testbed environments will be implemented within Wind's data centers to explore the advantage coming from the DOLFIN approach above all as far as concerns the energy reducing efforts by moving computation and services on a "federation" of IT data centres sites. In our case the federation will be based on two DCs having the same owner.

The following presents the developed solution: in the Wind federated infrastructure it will be possible to utilize two physically distinct data centres: the first one based in Ivrea (near Turin, Piedmont Region, North of Italy), the second one based in Molfetta (near Bari in Puglia, South of Italy). The two DCs are physically distinct and distant from each other (over 700 km) and are connected to two different Energy Providers. The facilities are designed and implemented to guarantee maximum flexibility and versatility of the testbed; the particular configuration of the federation reduces time to align infrastructure to business needs and allows Wind to:

- optimise the DC geographical distribution by HW resources utilization and manageability
- reduce incident impacts ensuring appropriate Disaster Recovery capabilities, increasing the quality of services.

The Wind Testbed will be included in this complex structure: the two data centre sites will provide virtualization tools ensuring both Linux and Windows environments. This will be useful to integrate the different DOLFIN components in the testing scenario. Moreover all the environment performance data, power consumption, CPU usage, RAM usage, HWs and VMs topology are stored into a RDBMS and available through dedicated interfaces based on SQL net protocol.

DOLFIN will offer tools to define a useful model for re-engineering of the systems taking into account different goals like the energy load control, and comparison of the best energy saving solutions according to the different proposals coming from the different Energy Providers.

We plan to follow a multi-parameter analysis able to:

- Analyse the servers and the applications utilization profiles
- Define the electricity consumption profile of the applications in order to determine a strategy for the optimized management of the same in order to reduce the electrical consumption

- Identify infrastructure innovative management plans, combining user SLAs with the data consumption of the infrastructure and the demands of the Energy Providers leading to the benefits in terms of savings in electricity consumption and reduction of CO_2 emissions
- place the elements according to energy consumption and according to the selected criteria
- provide additional information for datacentre control to support all management actions. (e.g. detection of abnormal consumption profiles)

4.2.2. Candidate technology

In the following we introduce the candidate technology (hardware and software) useful to test the DOLFIN solution taking into account the fundamental logic blocks:

1. The storage infrastructure, distributed and shared between the datacentres
2. The network infrastructure, extending layer 2 between the two datacentres
3. The local Business Intelligent platform for the infrastructure data consumption and infrastructure topology
4. The local infrastructure consisting of a couple of HP servers in each datacentre (to provide local redundancy in any case), on which certified applications are run, DOLFIN included

4.2.2.1. The distributed storage infrastructure

The distributed storage architecture between the two datacentres is deployed using the EMC2 VPLEX solution. By positioning this platform between the SAN (Storage Area Network) and the related storage inside the datacentres, data mobility, high availability and enhanced cooperation between the two sites are fostered. Data is therefore presented in a consistent manner within and between sites, even at remote distance thanks to the Access Anywhere technology. The SAN infrastructure is based on EMC2 Symmetrix storage units: Figure 4-1 depicts the configuration: the (remote) users are connected through a WAN.

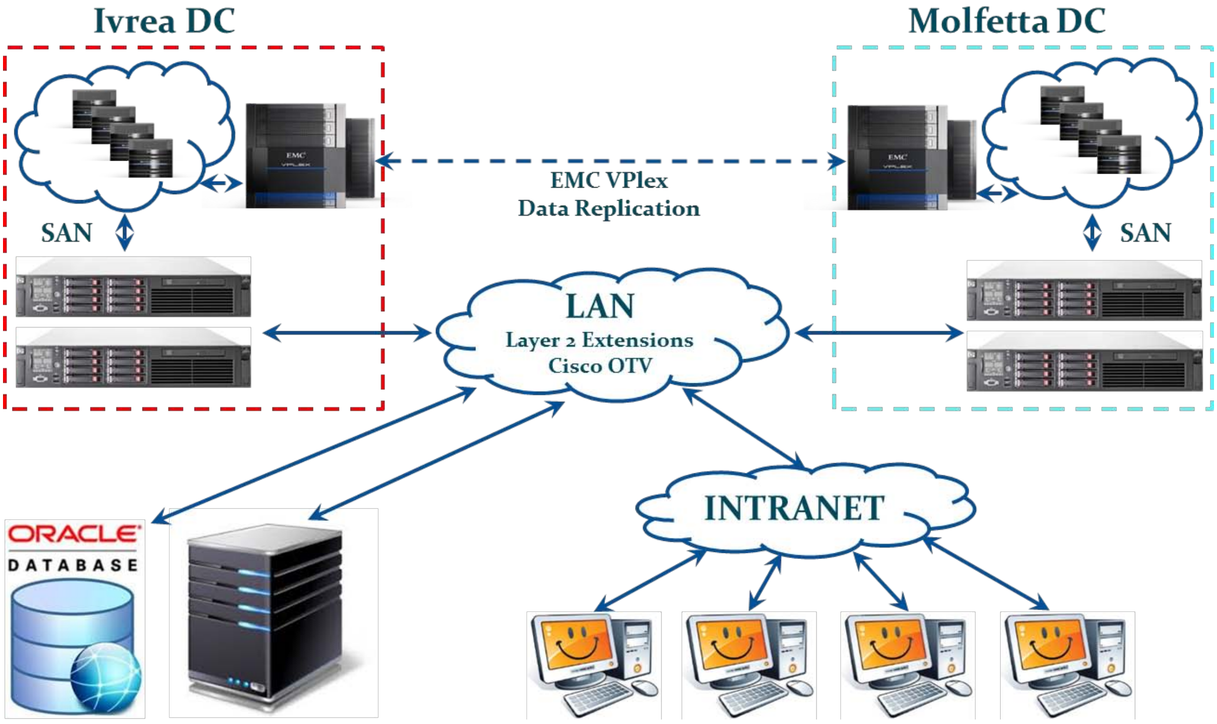

Figure 4-1: Main architecture: the model

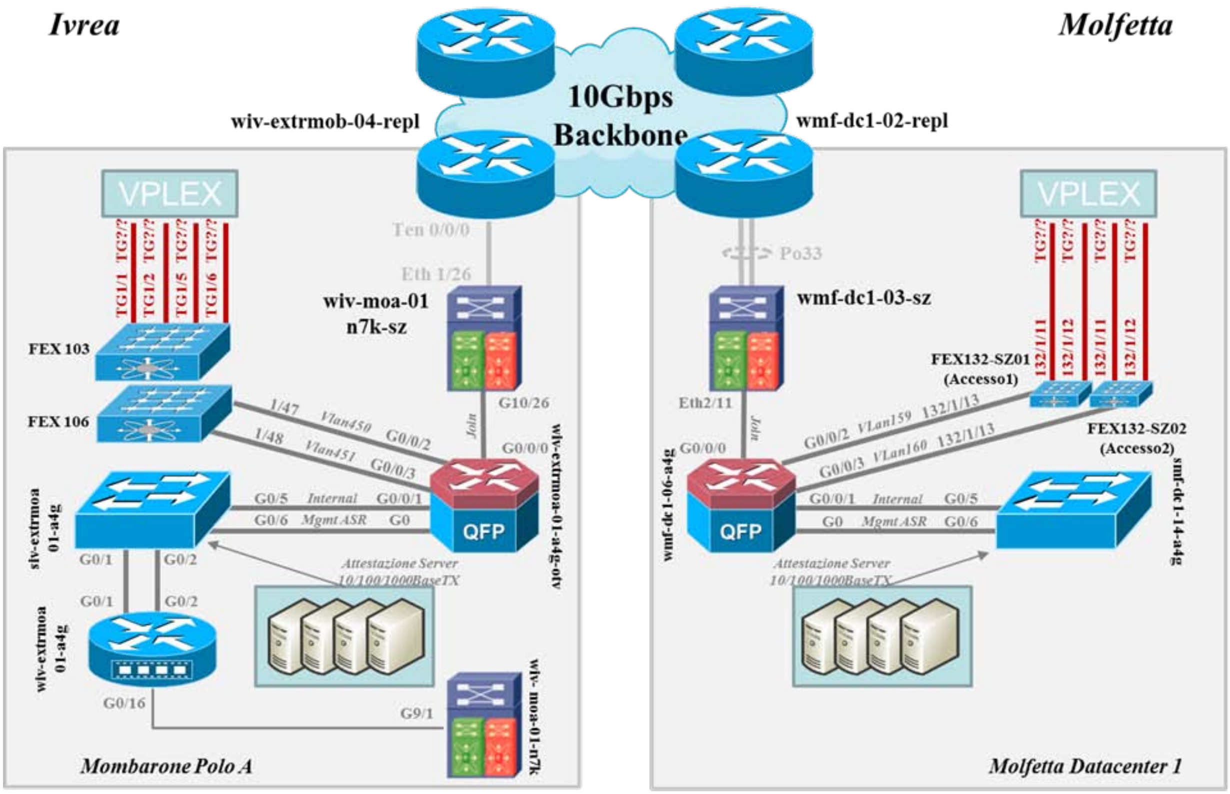

Figure 4-2: Main DCs architecture: the physical structure

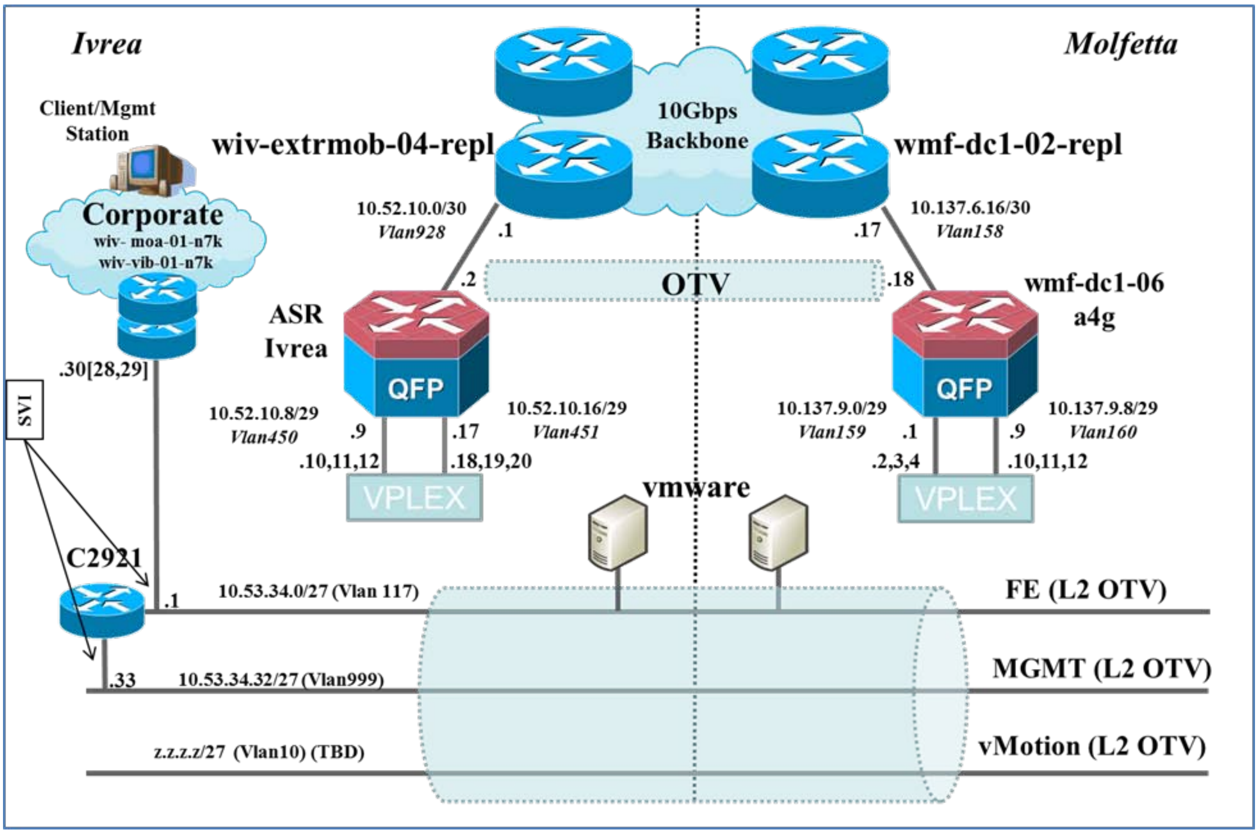

Figure 4-3: Layer 3 structure

The network infrastructure

The network architecture between the two datacentres must provide the following main capabilities:

- **Routing Network:** The routing network offers the traditional interconnection between remote sites and gives end-users access to the services. This component is improved using GSLB-based services such as DNS, HTTP redirection, dynamic host routes, and LISP
- **LAN (to WAN) Extension:** The technical solution for extending the LAN between two sites using dedicated fibres and Layer 2 over Layer 3 transport mechanisms for long distances.
- **Storage Services:** The storage services used to extend access between SAN resources. SANs are highly sensitive to latency.
- **Path Optimization Solution:** The path optimization solution improves the server-to-server traffic as well as the ingress and the egress workflows.

A dedicated fiber link between two datacentres is already provided. In our tests the performance capabilities will not be considered.

The local servers

Based on a dedicated infrastructure, the infrastructure provisioning end-users applications consists of a couple of HP servers on each DC. This dedicated testplant provides functionalities of centralized virtual storage and layer -2 extension between two DCs. The architecture is based on the following components:

- **Computing:** HP – DL380
- **Networking:** Cisco for IP switches and virtual switch; Brocade for storage switches (Fiber Channel)
- **Storage:** EMC storage with EMC Unisphere

Thanks to VMware virtualization software, it is possible run multiple operating systems (Windows or Linux, 32 or 64 bit), including Windows 7 and Chrome OS virtual machine, on blades provided by HP servers.

4.2.2.2. Application workload in TP Wind for DOLFIN

The application workload running on this platform is mainly consisting of a VDI (Virtual desktop infrastructure) platform. VDI is a server computing model enabling desktop virtualization, encompassing the hardware and software systems required to support the virtualized environment. Wind DCs adopt and support the VDI client model.

4.2.2.3. Virtualisation platform

The VMware Virtualization software platform allows access to multiple virtual servers even running with multiple different operating systems (Windows or Linux, 32 or 64 bit), including Windows 7 and Chrome OS virtual machine, on blades located in HP servers.

Virtualization platforms allow the user to simultaneously run multiple virtual machines on local hardware, such as a laptop, using hypervisor technology. Virtual machine images are created and maintained on a central server, and changes to the desktop VMs are propagated to all user machines through the network, thus combining both the advantages of portability afforded by local hypervisor execution and of central image management. This approach requires user hardware more capable of running the local VM images, such as a personal computer or notebook computer, and thus is not as portable as the pure client-server model. Internal users can access their desktops from any device connected which has the necessary permissions. Virtual machine images are created and configured on a central server, stored on the attached storage “ready to use” whenever a new VM is needed and changes to the desktop VMs are propagated to all user machines through the network.

The system CPU requirements are calculated using a specific vendor's reference architecture. It is possible to allocate 20 or 30 virtual desktop instances for a system with 8 or 16 cores.

The system memory (RAM) requirements are calculated utilizing vendor supplied sizing information. A common amount of RAM required for a virtual desktop instance is 1 or 2 GB of physical memory.

Client-hosted virtualization does not require VDI execution servers, because the virtual machines are executed on end-point machines. The networking requirements, both between servers and clients running VDI images, along with internal server-to-server networks and storage networks are configured properly. Storage plays a significant and strategic role in the overall performance of VDI. The use case proposed and implemented will make it possible to relocate the VDI services indifferently both in the datacentre of Ivrea, and in that of Molfetta with RTO (Recovery Time Objective) and RPO (Recovery Point Objective) near to zero.

This means that, in case of services relocation:

- Users get access to the Virtual Desktop of the active farm without interruption or at the most with an interruption limited to a VD re-login
- Users get access to documents and profile data without interruption or at the most with interruption limited to minutes
- Users don't lose documents or at the most lose only documents on which they were working at the time of the fault. Under normal conditions users always access on master Virtual Data Mover Storage replication via EMC Replicator, without new server and components are activated

4.2.2.4. Data Center management

The datacenter management, the automation frameworks and tools deployed into Wind's premises are designed to guarantee a full set of capabilities that allow full management of datacentres from every point of view even if at the moment we have to point out a general lack of tools and methodologies for measuring performances and energy efficiency indicators in a datacentre in a simple way. Standard metrics and measurement processes, agreed among vendors, datacenter managers and clients, would allow easier communication and better understanding.

Data centre management applications can be logically grouped into two main classes: the first one related to the tools that provide capabilities for the monitoring of infrastructures, processes, etc.; the second one related to the tools that provide functionality of Configuration Management and Datacentre Automation. More details are presented for each class as follows: there are 3 types of monitoring:

- Real-time Monitoring (Alarms & Events)
- Performance Management
- End-to-End Monitoring

The Wind monitoring platform is based on the framework GroundWork Open Source (GWOS).
To realise the "Virtual Management" the VMware suite is used.

4.2.2.5. Service Level Agreement

From a Telecom operator point of view, generally the workload allocation policies, driven by acting Service Level Agreements, require that no degradation be allowed. With this constraint in mind, it will investigate, using the DOLFIN capabilities, how to define a new efficient policy with less narrow thresholds for the SLA management exploring new potentialities coming from a stronger relationship with the different energy providers. It will be particularly useful to evaluate how to deal with the real market constraints when new capabilities will be offered in order to find new trade off points among the "service continuity " and the global cost of this strong requirement, or other coming from the new scenarios defined in the "smart environments".

The current SLAs and allocation policies are to guarantee Business Continuity; that means the RTO (Recovery Time Objective) and RPO (Recovery Point Objective) must be near zero, where and when it is possible. Using the testplant these capabilities will be verified as well.

4.2.3. The Federation

As indicated, the testbed environment as devised and implemented exactly meets the needs of the project: the two datacentres are physically distinct and distant from each other (over 700 km) and are powered by two different energy providers.

The federation is guaranteed by the architecture and infrastructure services envisioned and prepared that can lead to having two distinct scenarios of use of the datacentres:

- Active - Active: both datacentres offer the VDI service to users
- Active - Standby: one datacentre provides a full VDI service, while the other remains Standby.

The services provided can lead to a change of state of the Datacentres in case of needs:

- from active - active to active – standby
- from active - standby to active – active

To better clarify the concepts introduced above, more details are given in the following:

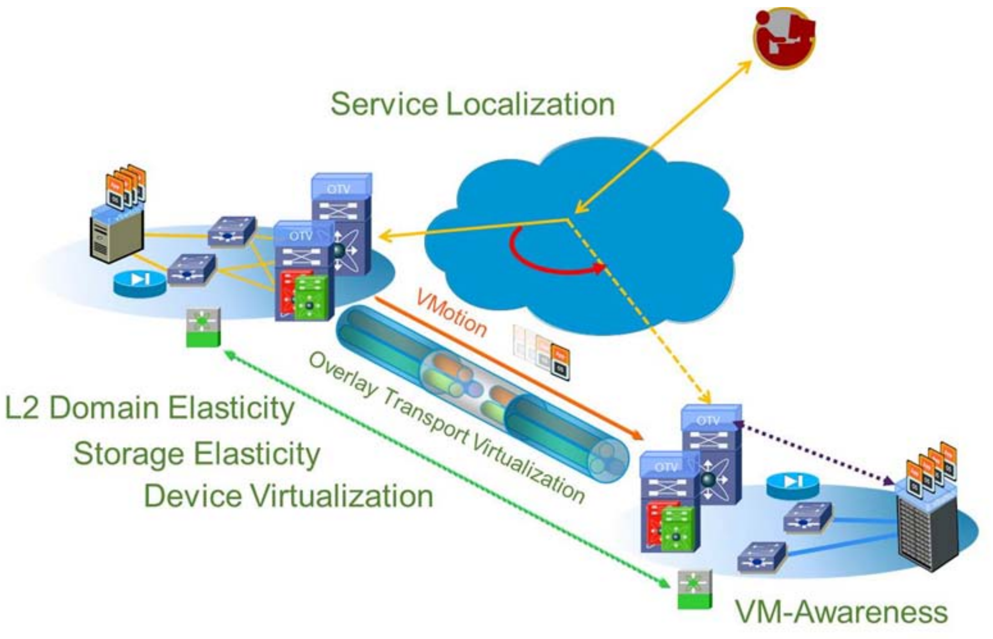

Figure 4-4: The implemented technical architecture

Considering Figure 4-4 the implemented technical architecture guarantees the service availability for the end users: all aspects are involved from the network to the service level; during the load switch this architecture allows full replication and seamless transaction: in this way any disruption of work will not occur. The networking technologies and architecture used where there are the main features of “Routing Network”, of “LAN Extension”, of “Storage Services” and “Path Optimization”, as represented in Figure 4-5

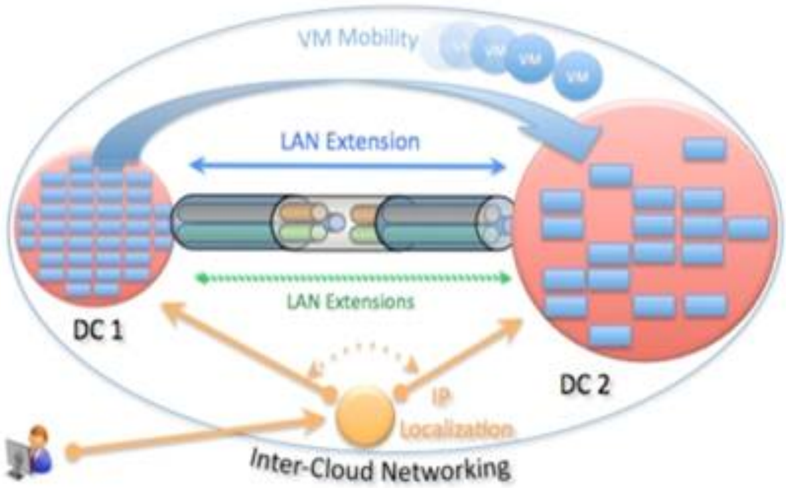

Figure 4-5: The networking schema

The advanced storage infrastructure and related management tools are based on EMC VPlex solution.

- the advanced virtualization software and "cloud management" software such as: VMWare View, VMWare HA, VMware vCenter, …

In case of the Active – Active scenario, during the load switch from the Datacentre1 to Datacentre 2 (VM move), the end user does not suffer any disruption of work. In fact:

- the application data on the storage of each user is always available in two datacentres
- for very few seconds (about 3 – the time to physically migrate the VM from source DC to target DC) the end-user VM will be frozen.
- After the VM relocation, the end-user restarts work with the same application session

In case of the Active – Standby scenario, during the load switch from the Datacentre1 to Datacentre 2 (VM move), the end user suffers some disruption of work. In fact:

- the end user application session will be frozen and stopped
- the target server in the target DC will be started-up
- the application data on the storage of the user will be synchronized and put available for the target server into target datacentres
- the end-user must login again into the system

With VDI, the end-user experience remains familiar. Their desktop looks just like their desktop and their thin client machine performs just like the desktop PC they've grown comfortable with and accustomed to. They have greater control over the applications and settings that their work requires

4.2.4. The CUSTOMER

From the end user perspective (the CUSTOMER) we have to deal with the necessity to save the energy cost: performance contracting often is thought of as receiving a guaranteed amount of energy savings that helps offset some or all of the costs of facility and infrastructure renewal projects: it is necessary to establish performance contracting. Traditionally, performance contracting allows users to make capital improvements, save energy, reduce emissions, improve sustainability and address tight budgets; besides using an integrated approach it is possible to use the expected utility and operational savings to offset the cost of the upgrades.

4.2.5. The Energy providers

It could be interesting to ask the energy provider to offer different levels of insurance; if the provider doesn't achieve the savings that are guaranteed, it has to pay the difference between what was guaranteed and what was actually achieved. The basic point is the "verification": the TELCO end user measures, manages and reports on the obtained progress and requires energy-saving, performance-based contracting with energy organizations. A detailed audit could identify opportunities to improve the energy efficiency of different elements for the data centres, e.g. lighting, heating, air conditioning, ventilation systems and all the other systems for data management; following this path, the DOLFIN system will offer elements for a technical review to define the best strategy and to confirm the ability to create benefits such as energy savings or improved accuracy. Automated measurement and verification tool as DOLFIN will provides optimization. DOLFIN will offer a comprehensive approach, identifying a variety of ways to achieve savings and operational efficiencies. It is necessary to select which ones best suit the involved organization based on business criteria such as payback or criticality.

4.2.6. Wind Testbed Summary

As described, the DOLFIN testbed environments implemented within Wind's data centers will make it possible to explore the advantage coming from the DOLFIN approach above all as far as concerns the energy reducing efforts by moving computation and services on a "federation" of IT data centres sites. All DOLFIN architectural components should be tested and analysed for each selected scenario according to the different Use cases evaluated in the D2.1.

4.3. IRT TestBed

4.3.1. Description of IRT Testbed

Interoute will provide a dedicated space for the implementation of the Testbed environment, for the purpose of evaluation and validation of the DOLFIN functionalities. The main objectives addressed by the Testbed may be summarized in four topics:

1. Provide a sets of devices and tools to collect information useful to evaluate both the energy consumed and the performances for servers (and virtual servers) and IT equipments. The Testbed will be used to analyse the DOLFIN behaviour during different stages of the test.
2. Monitor the quality of the services provided to the end customers in respect of contractual SLA. Also the system should provide control points (via dedicated interfaces) to change the SLA parameters and adjust consequently the assets of the deployed service and/or server instances.
3. Validate the energy efficiency obtained by DOLFIN approach in the context of single DC.
4. Validate the energy efficiency obtained supporting the integration of multiple DCs, implementing the "federated" DCs, where each DC will has to implement a own instance of DOLFIN system

Based on these primary requirement, Interoute will provide a co-location service within Rome or Milan DC. The Interoute Colocation Service will comprise of the installation and the support services associated with the provision of colocation facilities at Interoute premises.

The co-location service provides the availability of:

- Devices powering through an energy distribution system consisting of:
    - DC powering circuit, supported by dual 2000A backup energy stations.

  - AC powering circuit, protected by a twin UPS system.
  - A 1250 kVA genset.
- Forced air cooling system with double compressor chillers and a refrigerant capacity of 500 kW.
- Standard rack to house servers and other IT equipment.
- Access to Public network through the Interoute backbone. Also Interoute will provide the necessary public IP addresses for the servers and the other network devices that should be reachable from Public network.

A concrete setup of the Testbed infrastructure will be defined after the complete definition of the test cases that should be provided as formal activities in WP5. As general guideline the infrastructure will supports some main characteristics as briefly describe in the below section.

Physical servers

For what concerning the server's hardware, Interoute suggests to use a standardised industry leading equipment, such as facilities from HP, DELL, IBM. The physical servers should be differently dimensioned on the basis of the specific application requirements, for example the hosted operating system or the dimension of the virtualization system that should be operate on the physical environment.

Virtualized environment

For the virtualized environment Interoute recommends the use of open Virtual Machine hypervisors, such as KVM, XEN for both Linux and Windows operating systems. Virtual Servers are provisioned allocating resources from a shared infrastructure of processors, memory and storage. Onto this virtual infrastructure, the virtual servers will built using the same operating systems as their dedicated counterparts, and include the same hardware found in the physical server, such as CPU, RAM and other internal components.

A typical example of Virtual Server hardware specification (as defined in the Interoute portfolio) might include the following sizing:

- Small Virtual Server: 1 x vCPU, up to 8 GB of RAM, 50 GB of storage, Windows or Linux OS.
- Medium/Hi Virtual Server: 2 x vCPU, up to 16 GB of RAM, 500 GB of storage, Windows or Linux OS.

Monitor tool

A monitor tool will be utilized to collect performance data such as the consumed resources of CPU, RAM, Power, etc. Different standard tool should be provisioned depending on the monitored resources, for example:

- NMS tool to monitor physical hardware (i.e. Nagios, Zabbix, Cacti). This component is an automated system that uses MIBs (Management Information Base) in conjunction with SNMP (Simple Network Management Protocol) and a software agent installed on the servers and switches.
- Virtual Machine monitor. This component generally depends on the VM system environment installed on the physical servers. The role of the Virtual machine monitor tool is to collect all relevant information regarding the virtualized environment, such as the VM topology and the performance data for each VM instance.

All the legacy monitoring infrastructure will has to be integrated with the DOLFIN monitor module via dedicated APIs or mediated by some monitor events broker.

4.3.1.1. General DC infrastructure for the Testbed

The Testbed environment will be deployed over a physical DC infrastructure that is built using the Interoute standard specifications. The main characteristics of this specifications include both directives for the IT equipments and connectivity interfaces and also include a complete sets of requirements regarding the physical environment, such as rack unit sizing and location, powering infrastructure, cooling functions, etc.

Power requirements

Each datacenter have two separated power feeds, each feed is totally independent from the other and is backed up with either a UPS and or generators. All power is filtered to remove any power spikes and voltage fluctuations. The two power feeds should be known as Blue and Red, where Blue is the primary power supply and Red the secondary supply. Each Rack must be fitted with at least two power bars, one power bar is connected to the primary (Blue) feed and one to the secondary (Red) power feed.

All power bars must be 19" rack mountable and of the 250 VAC DI-Strip EIC 320 socket system. Each power bar must be of at least 10 amps (fused) and provide between 8 and 10 EIC 320 style power sockets. The power strips must connect to the mains power supply via IEC 60309 16 Amp 2P+E 230V sockets and plugs.

Air-conditioning & Airflow

The air-conditioning system within a datacenter is normally outside of the control of Interoute. However, the rack layout must be designed to provide the optimum airflow / cooling.

Rows of racks must be installed such that the gaps between the rows are either Hot or Cold. A Hot row is one into which the hot air is expelled. A cold row is one from which cold air is taken in. Equipment within the racks is arranged such that the front of the machine faces into a cold row, and the back of the machine faces a hot row.

Rack equipment installation

Equipment is installed so that its front faces to cold row of the rack, the back end of the equipment which normally contains that power and network interfaces go towards the hot row.

Equipment must never rest directly on top of another piece of equipment in a rack. All equipment must be installed with its own rack mounting kit or on top of a dedicated shelf.

All servers must be fitted with two power supplies, the first connected to the Blue power bar, the second to the Red. In some cases i.e. hubs, small switches and firewalls, the equipment may be fitted with only a single power supply in which case if the device is installed in the top half of the rack it should be connected to the Blue power bar, or the Red power bar if it is installed in the bottom half of the rack.

Network infrastructure

The generic target architecture for the DC networks is based upon the Interoute metro network architecture type 1 and type 2 POPs and features the following:

- An access router layer to provide IP-layer termination.
- A switch aggregation layer to provide aggregation of multiple edge switch functions.
- An access switch device to provide Ethernet-layer termination on 10/100/1000 Mbps Ethernet ports.

Each DCs can support up to 9 access switches with 48 10/100/100 Mbps ports, with a aggregated link from 1Gbps to 2Gbs aggregate bandwidth between metro switches.

Where feasible, the DC layer-2 Ethernet domain is isolated from any IP network equivalent in the same location because of the likely difference in MAC-address consumption. The DC layer-2 Ethernet domain operates the exact same standard for layer-2 networking as the IP network. The generic network architecture is described in the Figure 4-6: Generic POP network architecture.

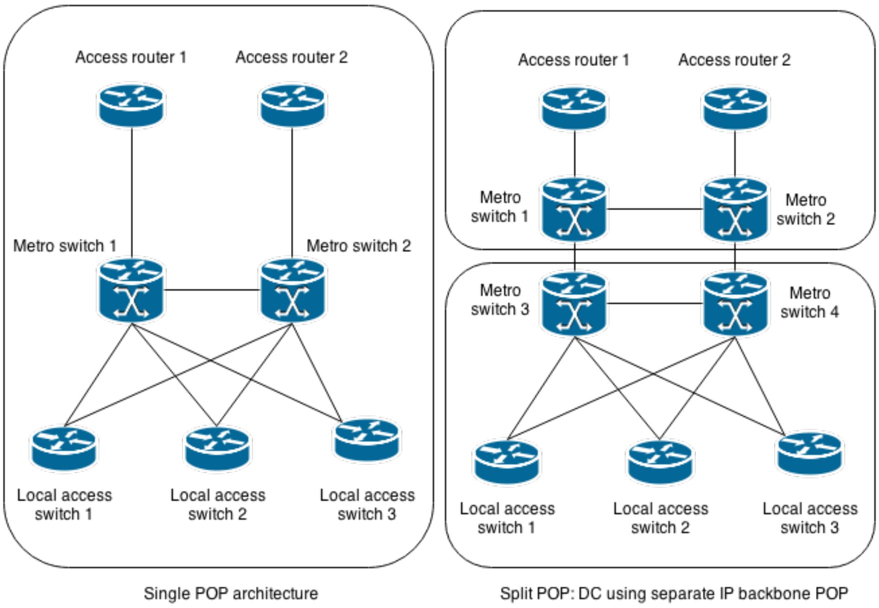

Figure 4-6: Generic POP network architecture

Inside the DC two main network are provided; the Public Internet network and the Out-Of-Band (OOB) used for internal and private scope.

The public network provide the access to Internet using a redundant connection to the Interoute Backbone. The connection to the backbone is sustained by two core switch, generally configured to works in fail over manner. The customer equipment are connected to a neighbourhood cell switch that receives a direct Gigabit fiber connection to each of the DC Core switches. Each customer is assigned their own VLAN on the cell switches; the VLAN is named with the customer service number. Each customer is also assigned their own subnet from the local datacentre's allocated public IP space, the first useable IP address from the subnet is attached to the customer VLAN or the customers managed firewall. In the Figure 4-7: DC Public Network Diagram is reported a basic diagram for the network connection to the Public Internet.

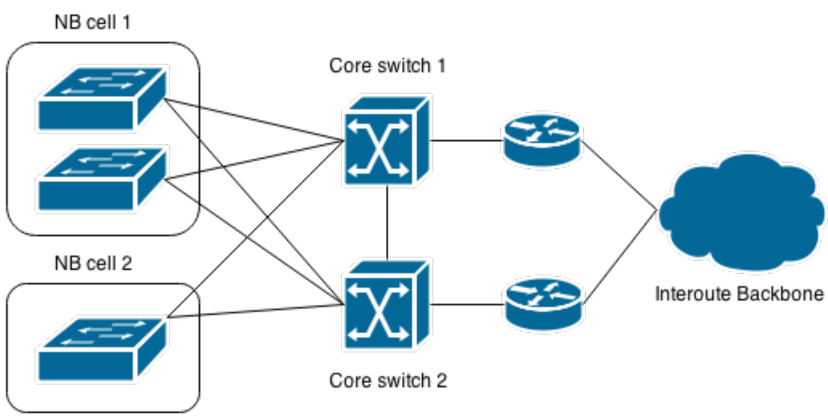

Figure 4-7: DC Public Network Diagram

In addition to the Public Internet network a OOB private network is provided for each DC. This network is designed to supports and carry the majority of management traffic, such as the access to server console, the access to the switch's management interfaces, etc. No customer traffic (included backups) must pass over the local DC OOB/Management networks.

4.3.2. IRT Testbed Summary

Interoute will provide a complete physical infrastructure and connectivity resource needed for hosting the DOLFIN Testbed environments. Moreover Interoute will offer the own expertise to prepare and consolidate the Testbed deployment. Also, Interoute will provide their supports during the test phases activities, such as the analysis, evaluation and validation of the test results of the DOLFIN functionalities, as will be defined in the DOLFIN architecture specification and use cases.

Extra efforts will be necessary to emulate some missing components identified in the DOLFIN architecture. More precisely a simulation of Smart Grid interfaces and DC facilities should be integrated in the Testbed to complete all test cases.

4.4. GRNET TestBed

4.4.1. Description of GRNET TestBed

The main GRNET DC is located within the premises of the Greek Ministry of National Education and Religious Affairs in Athens. The DC is currently equipped with 28 racks hosting servers and storage equipment reaching a total of 7132 logical CPUs and 1800 TB of storage. The average energy consumption of the equipment hosted at this DC is currently around 250 kW. The achieved PUE varies close to 1,8, while the energy consumption in 2011 was close to 850 MWh. This DC has been designed and implemented following high standards regarding the cooling efficiency. In-row cooling techniques are applied, while the equipment is fully virtualized. PUE is further optimized with free cooling techniques. The DC chillers are connected in parallel with air cooled heat exchangers. When the ambient air temperature drops to a set temperature, a modulating valve allows all or part of the chilled water to by-pass the existing chillers and run through the free cooling system, which uses less power and uses the lower ambient air temperature to cool the water in the system.

4.4.1.1. Hardware architecture

The hardware infrastructure of the GRNET TestBed in shown on **Figure 4-8**. The system comprises two kinds of servers, Virtual Machine Containers (VMC) and Storage Containers which serve the application execution and storage requirements respectively. In the following we provide details on the system's components.

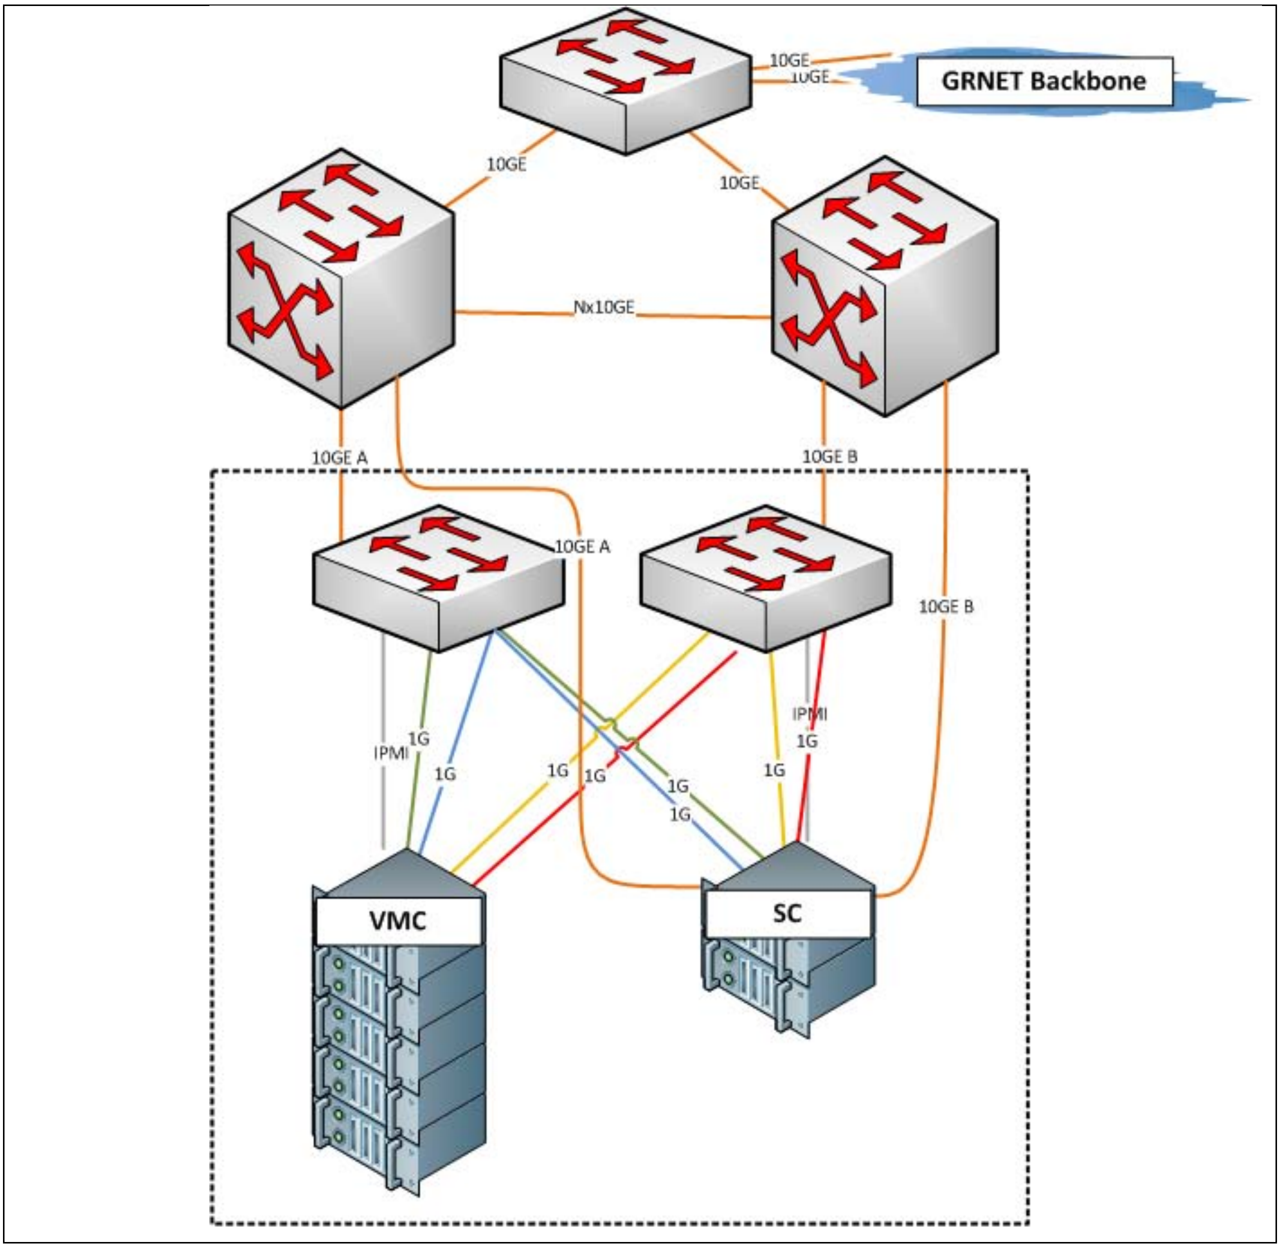

Figure 4-8: GRNET TestBed hardware architecture

Virtual Machine Containers (VMC): The GRNET Data Center hosts 234 HP DL385G7 servers to act as virtual machine containers. The technical characteristics of each server are summarized in the following table:

Processor model	AMD Opteron 6172
Number of processors	2
Number of cores per processor	12
Number of threads	12
Memory	192GB
Frequency	2.4GHz

Local Storage	480GB
Power consumption	80W
Size	2U

Table 7 - The technical characteristics of each server

Servers include a dedicated Ethernet port for remote administration over TCP/IP providing both remote console and VGA. Servers can be monitored remotely using protocols such as IPMI 2.0 and SMASH CLP providing information on the server's status including hardware health, power consumption, temperature etc. Server's power supply and cooling facilities can be remotely administered. Remote on and off powering is also supported.

Storage Containers (SC): The GRNET Data Center hosts 41 HP DL380G7 storage servers with the following characteristics:

Processor model	Intel Xeon E5645
Number of processors	2
Number of cores per processor	6
Number of threads	24
Memory	96GB
Frequency	2.8GHz
Local Storage	600GB
Power consumption	80W
Size	1U
Storage #1	HP StorageWorks D2600 Disk Enclosure
Storage volume	12x2TB SAS 7.200 RPMs
Size	2U
Storage #2	HP StorageWorks D2600 Disk Enclosure
Storage Volume	12x600GB SAS 15.000 RPMs
Size	2U

Table 8 - Servers characteristics: Remote monitoring and administration capabilities are similar to those of the VMC servers.

Network: The GRNET TestBed is well interconnected with Nx10 Gbps network links with the GRNET backbone network, ensuring the reliable and efficient provision of services. Dedicated circuits (even optical circuits) may be provided within specific locations in the GRNET network for testing purposes as well as for occasions where very high QoS requirements are imposed. In DOLFIN, several setups will be provided based on the type of the experiments with the dynamic allocation of end-to-end circuits.

4.4.1.2. Software architecture

The GRNET DC hosts an IaaS platform called ~okeanos operating as a public cloud for the Greek academic community. The goal of the ~okeanos project is to deliver production-quality IaaS to GRNET's direct and indirect customers, IT departments of connected institutions and students/researchers respectively. ~okeanos offers its users access to VMs, Virtual Ethernets, Virtual Disks, and Virtual Firewalls, over a simple web-based UI. ~okeanos was conceived to offer easy and secure access to GRNET's DCs, focusing on user friendliness and simplicity, while being able to scale up to thousands of VMs, users, terabytes of storage.

The ~okeanos service is a jigsaw puzzle of many pieces (see Figure 4-9):

- Compute/Network Service (codename: cyclades)
- File Storage Service (codename: pithos+)
- Identity Management (codename: astakos)
- Image Registry (codename: plankton)
- Billing Service (codename: aquarium)
- Volume Storage Service (codename: archipelago)

which are combined with a number of additional DC activities (monitoring, issue handling, helpdesk operations).

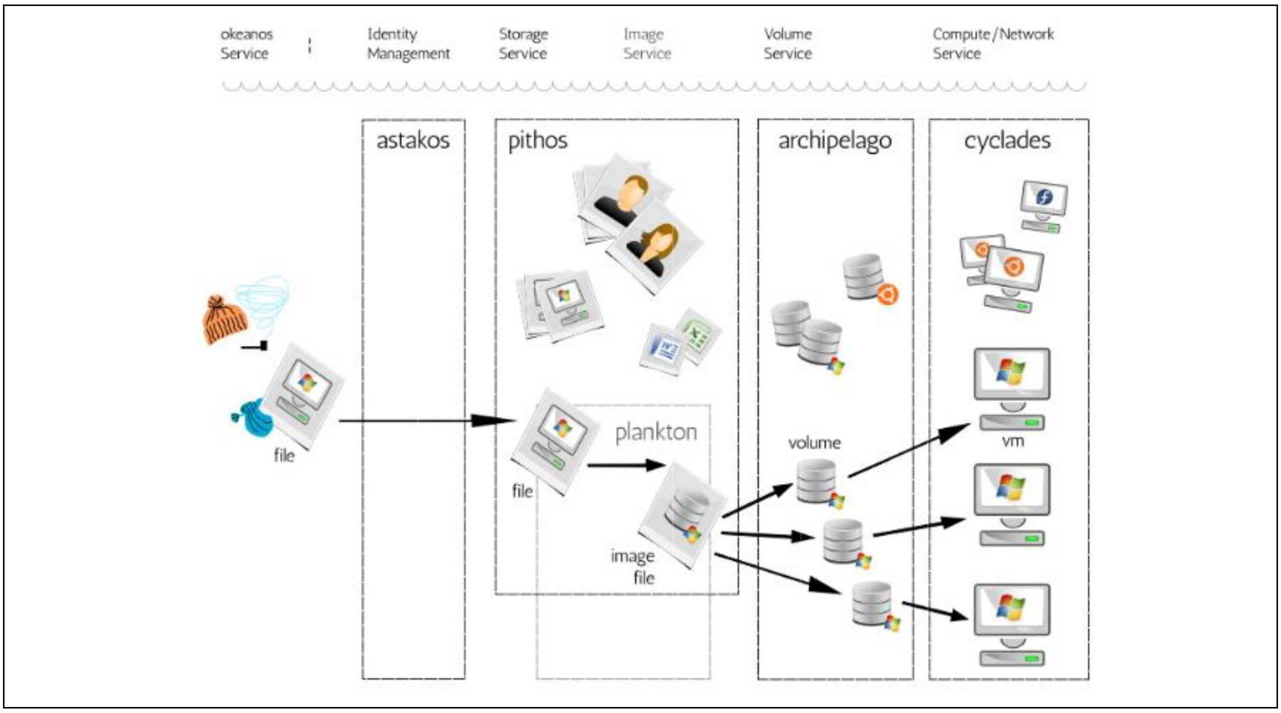

Figure 4-9: Synergy between ~okeanos components

The software underlying ~okeanos, called Synnefo, is custom open-source cloud management software. It encompasses a number of distinct components, all sharing a single installation and configuration mechanism, to streamline operations. Synnefo has a compatible API with OpenStack, and its software stack is shown in Figure 4-10

Figure 4-10 next to that of OpenStack.

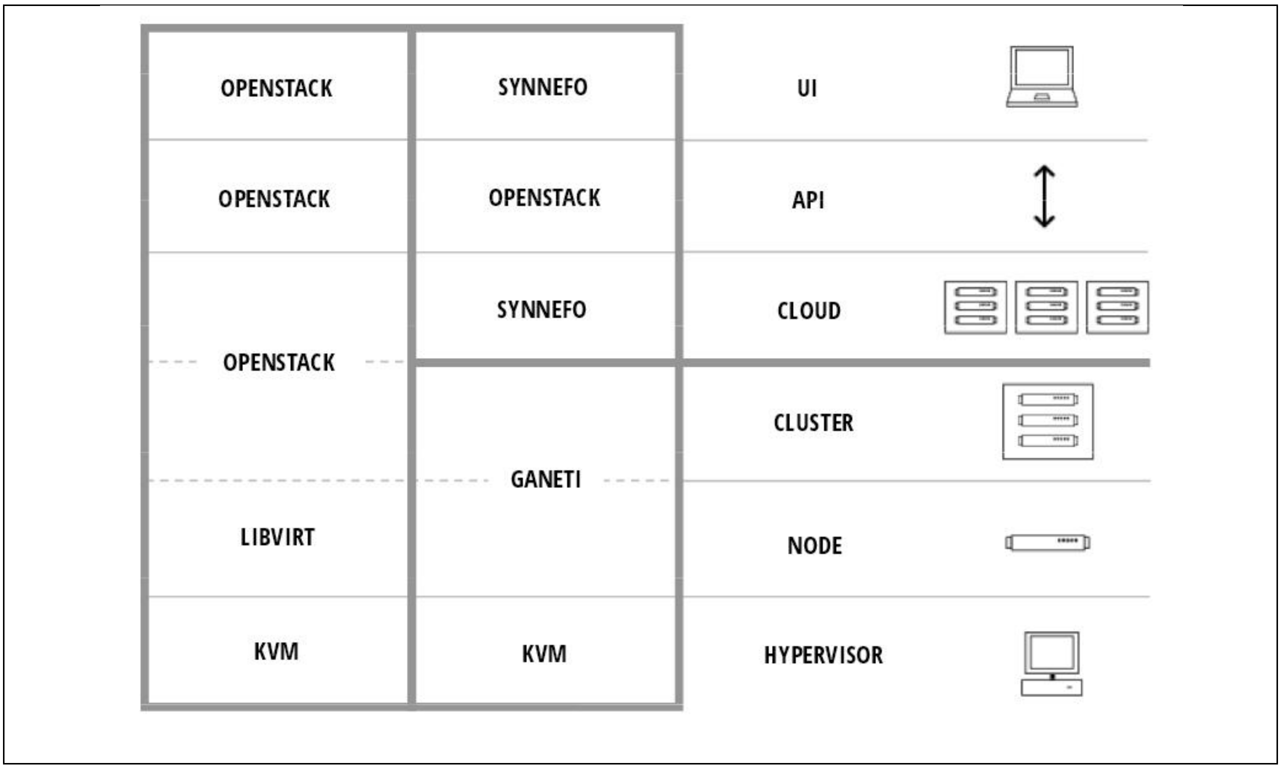

Figure 4-10: The synnefo software stack for GRNET ~okeanos IaaS platform

Cyclades is the Compute/Network part of ~okeanos. Its design combines a Google ganeti backend for VM cluster management (VM creations, migrations, etc.) with a Pythin/Django implementation of the user-visible API at the front-end. Ganeti organizes the VMs of the platform in Clusters and assigns one Ganeti master per cluster to manage VM creations, migrations and failovers (see Figure 4-11). The REST API for VM management is OpenStack compute v1.1 compatible and can interoperate with 3rd party tools and client libraries. GRNET has added custom extensions for unsupported functionality.

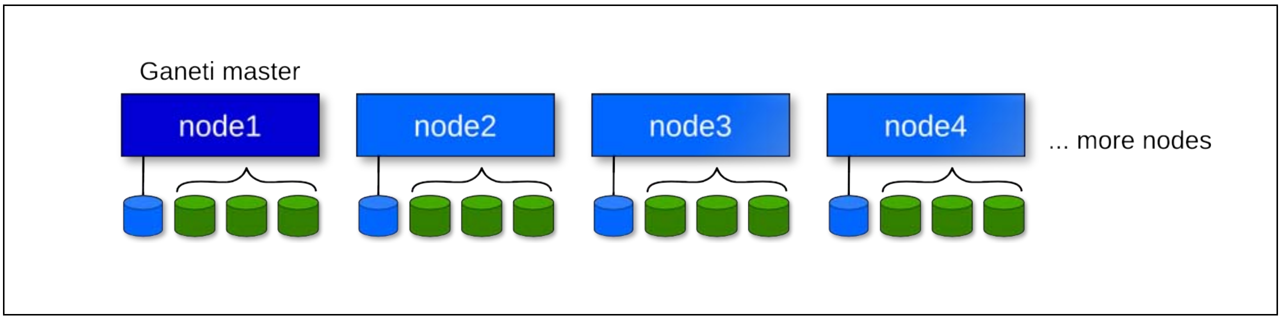

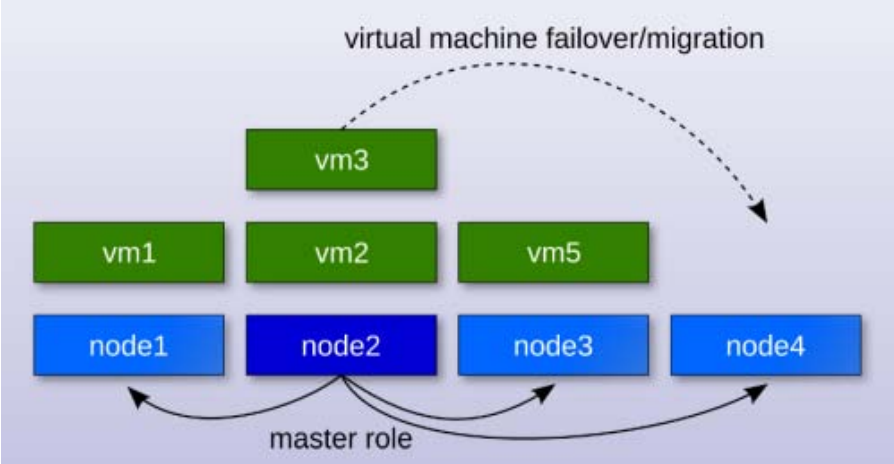

Figure 4-11: Basic Ganeti configuration

4.4.1.3. Monitoring facilities

Monitoring infrastructure is deployed supporting real-time monitoring of energy consumption per rack, as well as real-time estimation of PUE. This DC has been designed and implemented following high standards regarding the cooling efficiency. In-row cooling techniques are applied, while the equipment is fully virtualized. PUE is further optimized with free cooling techniques. More specifically, the DC chillers are connected in parallel with air cooled heat exchangers. When the ambient air temperature drops to a set temperature, a modulating valve allows all or part of the chilled water to by-pass the existing chillers and run through the free cooling system, which uses less power and uses the lower ambient air temperature to cool the water in the system. Energy consumption is monitored with specialized metering devices situated in the electrical panels, so that the total power consumption, the chillers and the rest infrastructure components can be also measured. In each rack there are installed power distribution units with monitoring capabilities to allow fine grain monitoring of the equipment inside the racks. Also, modern server hardware technologies are deployed with an on-board management controller that allows power monitoring in a per server basis. All this information can be accessed through standard protocols like HTTP, FTP or SNMP and can be easily integrated into a third-party system.

Figure 4-12 presents the output of the online PUE monitoring in GRNET TestBed (for more information visit https://mon.grnet.gr/powermeter/)

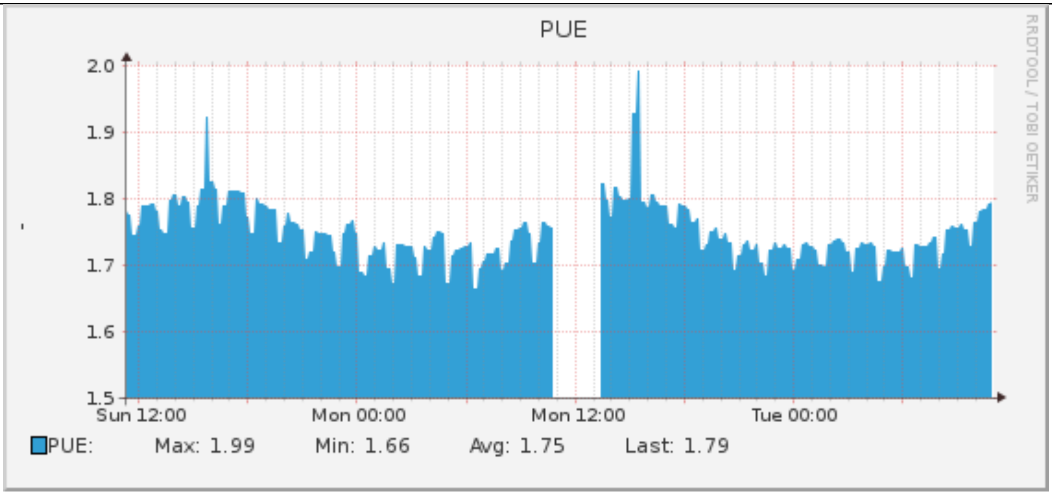

Figure 4-12: Online PUE monitoring in GRNET TestBed

4.5. UCL TestBed - A Software Defined DC Infrastructure

The following presents the developed Very Lightweight Software Driven Network and Service Platform (VLSP) testbed which is continuously developed at UCL and released as open-source software under the LGPL licence http://clayfour.ee.ucl.ac.uk. It is currently hosted on 12 H/W servers. Additional approx. 150 H/W Servers and networking nodes of the Heterogeneous Experimental Network (HEN) experimental network (http://www0.cs.ucl.ac.uk/research/hen/) are available for experimental purposes.

4.5.1. VLSP Overview

The Very Lightweight Software Driven Network and Services Platform (VLSP) test-bed includes a new lightweight network hypervisor, a novel infrastructure for the management and manipulation of virtual networks on top of the hypervisor and facilities for handling information / knowledge in the SDN environment.

The VLSP software stack consists of three layers:

- the Lightweight Network Hypervisor,
- the Virtual Infrastructure Management, and
- the Virtual Infrastructure Information Service (VIS) management components.

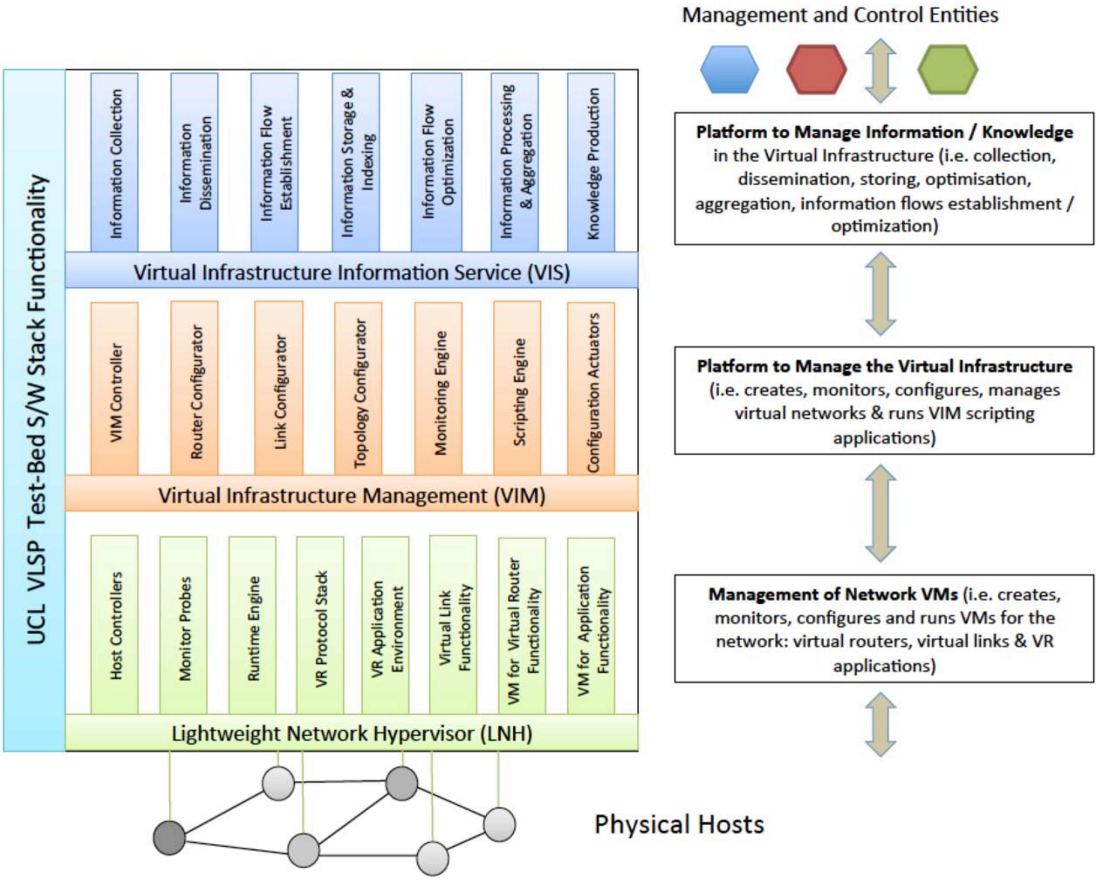

Figure 4-13: Overview of VLSP Test-bed Software Stack – University College London

An architectural overview of the software stack is shown in Figure 4-13. The VLSP test-bed software consists of over 700 java classes and more than 150 k-lines of code. In our experiments with VLSP, we have executed over 100 virtual routers on each of 12 dedicated physical servers. We detail the three main layers of the VLSP test-bed below.

4.5.2. The Lightweight Network Hypervisor (LNH) Component

The Lightweight Network Hypervisor includes a fully operational lightweight virtual router (VR) combined with virtual network connectivity. These elements can be combined in order to build any network topology required. The created virtual network is designed with the goal of transmitting and routing datagrams from any source to any destination. It behaves like a lightweight virtual network, but it has management facilities to start and stop virtual routers on-the-fly, together with the ability to create and tear down network connections between virtual routers dynamically. Furthermore, these lightweight routers have an application layer interface that provides the capability to start and stop Java software applications. These applications use a transport protocol API which can send and receive datagrams or packets, and thus act as the service elements within the platform.

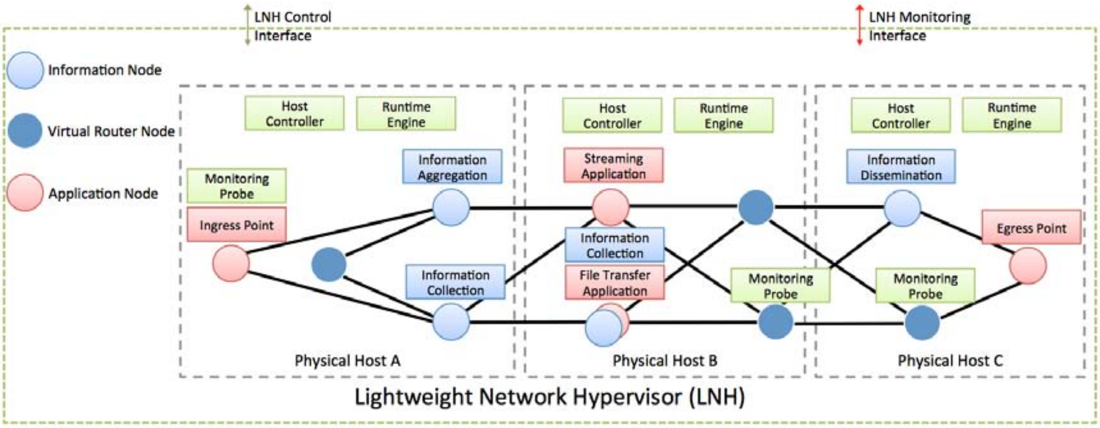

Figure 4-14: The Lightweight Network Hypervisor

The Host Controllers (shown in green in Figure 4-14) are executed on every machine that can host virtual routers. Their main job is to actually start a virtual router, stop a virtual router, start a virtual link, and stop a virtual link. Other tasks undertaken by the Host Controllers are to configure the routers once they have started, or to pass on commands to specific virtual routers, as needed.

The virtual network topology consists of virtual routers (shown as coloured circles) and virtual links (show in black). Each virtual router is instrumented with the VLSP monitoring system in order to gather data on each of the network interfaces of each virtual router. The data includes information on traffic volumes coming in and going out of each interface. The monitoring system collects the raw data and passes in onto the Monitoring Engine function of the above layer.

The main LNH functions specification can be found in the following table.

Name	Description
Host Controllers	The host controllers execute on every physical machine and manipulate & configure virtual routers, links and virtual router applications.
Monitor Probes	The monitor probes are tiny configurable applications probing the software or hardware for monitoring data.
Runtime Engine	It is responsible for the runtime operation of the LNH, including support for event-based notifications and time scheduling.
Virtual Router Protocol Stack	The lightweight network protocol stack of the VRs.
Virtual Router Application Environment	The application environment that hosts VR applications.

Virtual Link Functionality	The functionality of the virtual links, including link weighting and other configuration options.
Virtual Machine for Virtual Router & Application Functionality	A virtual machine with the virtual router and the relevant applications functionality.

Table 9 - The main LNH functions

4.5.3. The Virtual Infrastructure Management (VIM) Component

In this section we describe the Virtual Infrastructure Management component, highlighting its purpose and its architecture. The Virtual Infrastructure Management (VIM) component is responsible for the management and lifecycle of the virtualized elements that will be used within a network, particularly virtual network elements. As the virtual elements are not physical themselves, but exist on top of physical elements, their lifecycle and their management needs to be approached carefully to ensure continued operation and consistency. An overview of the VIM architecture on top of the LNH is shown in Figure 4-15.

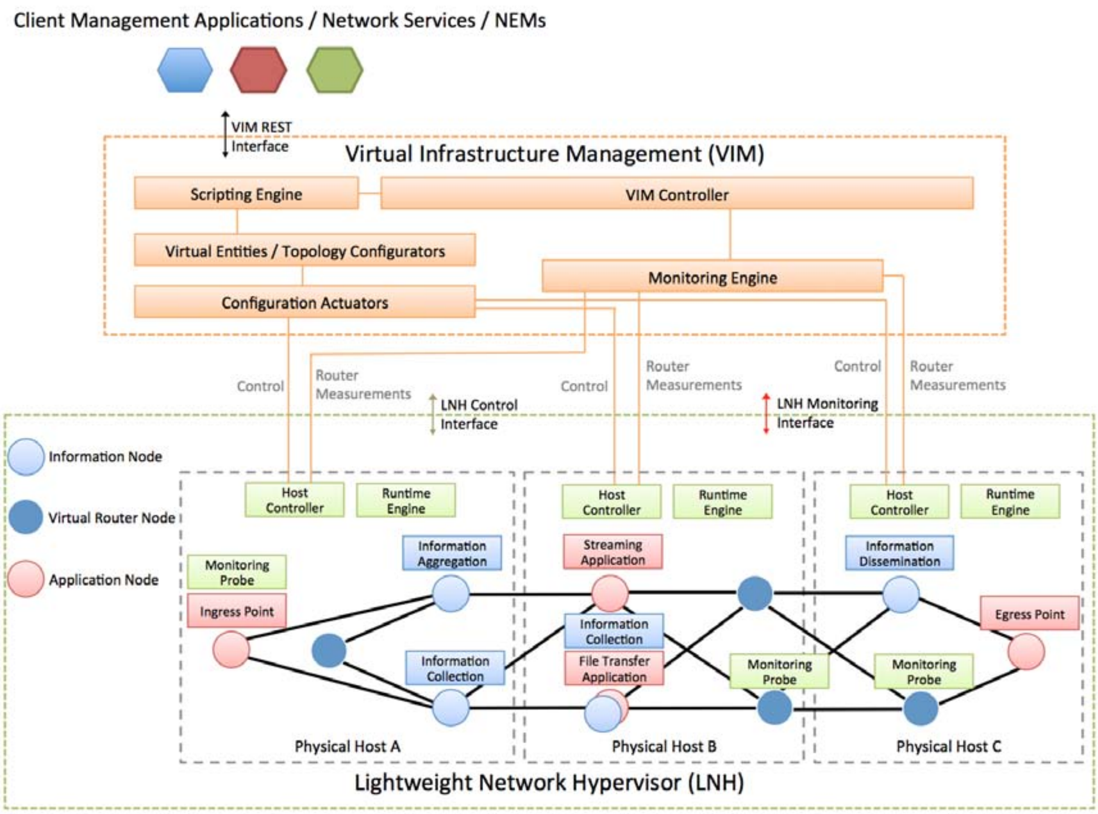

Figure 4-15: The Virtual Infrastructure Management Component

The virtual network elements, which exist on top of physical networks, can be setup with arbitrary topologies and with an arbitrary number of end-points. The virtual links in a virtual topology are eventually mapped onto physical links in the underlying network. A virtual link may span multiple physical links, and cross many physical routers, or it may span a single physical network link. New virtual links can be added or can be removed from a virtual network dynamically at run-time.

The virtual networks are very flexible and adaptable, and generally have few limitations, except that a virtual link cannot support more traffic and higher-data rates than the underlying physical links. Furthermore, a whole virtual network can be shutdown as needed, if the applications that use it no longer need the network. Such a shutdown frees resources from the underlying physical network.

The full management of virtual networks on physical networks requires the matching and analysis of the flow rates on the virtual links to the flow rates of the underlying physical links. It is important to ensure that the physical links are not congested with too many virtual links. Also, the allocation and mapping of virtual links must take into account the current state of the physical network and the current virtual networks. However, if a situation arises where a re-configuration is required, the virtual network management should be capable of mapping a virtual link across different physical links at run-time, but leave the virtual topology intact.

The VIM component has a seemingly simple task, but in reality the management requires continual monitoring, analysis, and adaption of the virtual elements to the physical elements. As all of these virtual elements are distributed, the management is a complex task.

The diagram in Figure 4-15 shows how the VIM component interacts with the virtual network elements that will be present in a running virtual network. All of the elements of the VIM component constitute a fully distributed system, whereby an element or node can reside on any host. A full virtual network can be instantiated on a single machine, for demonstration or testing purposes, or instantiated across multiple servers, in a full deployment situation.

The VIM directly controls the lifecycle of each virtual element, by collecting knowledge on the status of physical resources in order to determine where a virtual element can be created. The virtual network element will be created, managed, and shutdown by lifecycle phase of this component.

Due to the dynamic nature of virtual elements and because they can be disassociated from the physical elements they are mapped to, it is possible to do a live adaption of a virtual element from one physical host to another physical one, at run-time.

The VIM controller acts as a control point for managing the virtual elements. This block accepts all its input via the VIM REST interface from other management applications / network services. The monitoring engine acts a collection point for the monitoring data needed to keep the management functions running. Control commands are being sent to the VIM and they are either acted upon immediately or are passed to the corresponding Host Controllers of the LNH.

The main VIM functions specification follows.

Name	**Description**
VIM Controller	It is the heart of the component, providing the central control of the VIM operations.
Scripting Engine	VIM can be configured using Closure scripting.
Monitoring Engine	It is the main monitoring component of the infrastructure, i.e., collecting & manipulating measurements from the

	monitoring probes residing at the LNH.
Virtual Entities / Topology Configurators	These functions are responsible for the configuration of virtual routers, links and topologies, supporting different levels of abstraction.
Configuration Actuators	The Virtual Entities / Topology configurators communicate with the configuration actuators which in turn enforce the configuration changes through the LNH's host controllers.

Table 10 - The main VIM functions

4.5.4. The Virtual Infrastructure Information Service (VIS) Components

The Virtual Infrastructure Information Service component is a critical part of the VLSP and UMF since it plays the role of information / knowledge collection, aggregation, storage/registry, knowledge production and distribution across all the functional components, management applications, network services in the network environment. It can run on top of VIM, since it is fully integrated within the virtual network, e.g., the virtual routers have embedded information / knowledge manipulation capabilities. Furthermore, it is used by any client management application /network services/core service. As we have shown above, the VIS is fully integrated within the VLSP but acts as a standalone component as well. In this context, the VIS is the knowledge core service and supports a wide-range of network environments beyond SDNs.

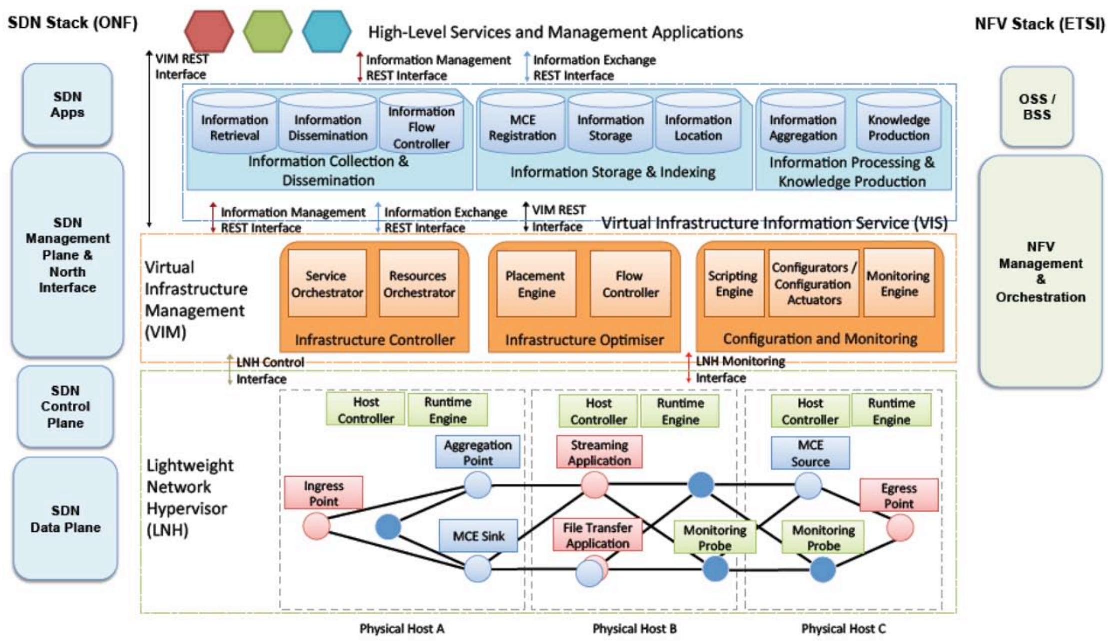

Figure 4-16: The Virtual Infrastructure Information Service Component

The main VIS functions can be found in the following table.

Name	**Description**
Information Collection & Dissemination	This function is responsible of information retrieval, sharing and dissemination.
Information Storage & Indexing	This function is a logical construct representing a distributed repository for registering information-enabled entities, indexing (and optionally storing) information/knowledge.
Information Flow Establishment & Optimization	This function regulates the information flow based on the current state and the locations of the participating entities and nodes.
Information Processing & Knowledge Production	The Information Processing and Knowledge Production function is responsible for operations related to information processing (e.g., aggregation) and knowledge production.

Table 11 - Main VIS functions

4.6. I2CAT TestBed

4.6.1. EXPERIMENTA platform

I2CAT will provide as in kind contribution to the project the service-oriented EXPERIMENTA platform, which aims at improving the competitiveness of enterprises so that they can accelerate their developments and validate prototypes and services in an optimal manner. Disruptive and high-end services are very expensive and private companies cannot effort them just for validation purposes. In addition, new technologies demand economies of scale to make validation and experimentation viable. EXPERIMENTA was born as a service–oriented platform deployed to solve the lack of places that offer advanced validation services. The EXPERIMENTA platform offers storage, connectivity, experimental validation and testing services on i2CAT experimental infrastructure. The facility provides last-generation optical transport equipment, together with test-beds of optical fibre access networks, and mesh networks, as well as connection to the different i2CAT international links.

The platform consists in an innovation metropolitan facility to test multi-layer devices on a real environment. It has more than 60 km-optical ring transport network with 3 ROADM-based optical nodes and 10GBps and 2.5GBps interfaces in layer 1, 2 and 3. It contains multiples Cloud Services and uses Juniper MX480 routers that allow a high efficiency, reliability and scalability to commercial and service application providers. The infrastructure presents an ideal environment to validate prototypes for Future Internet technologies, architectures, protocols and services. The different DOLFIN components could be tested in the EXPERIMENTA data centre. In addition it offers the capability for the end users to reconfigure their assigned resources according to their testing, validation and research needs. It is user-empowered platform, where users can establish and control their network resources. It is an experimental testbed for server for several finished and un-finished FP7 projects: FP7 GEYSERS, FP7 BONFIRE and FP7 OFERTIE. The description of the optical ring that is linked to the i2CAT infrastructure can be viewed in the Figure 4-17.

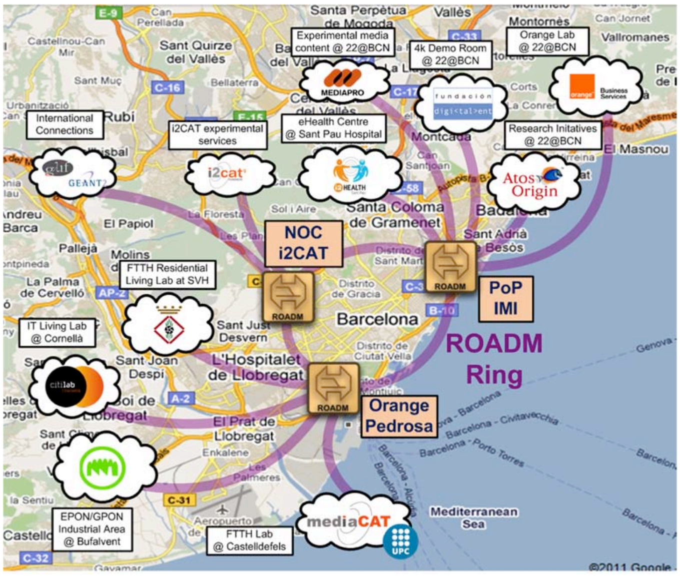

Figure 4-17: Optical ring linked to the i2CAT infrastructure

4.6.2. Equipment description

EXPERIMENTA connects more than 30 centres and the connectivity goes with the main national and international researched networks and has a commercial output of 1GB. In the Figure 4-18, the infrastructure of EXPERIMENTA is presented.

Table 12 - EXPERIMENTA hardware describes the hardware devices and

Table 13 - EXPERIMENTA software the software components. The facility has connectivity to GEANT testbed of 1/10G to UnivBris.

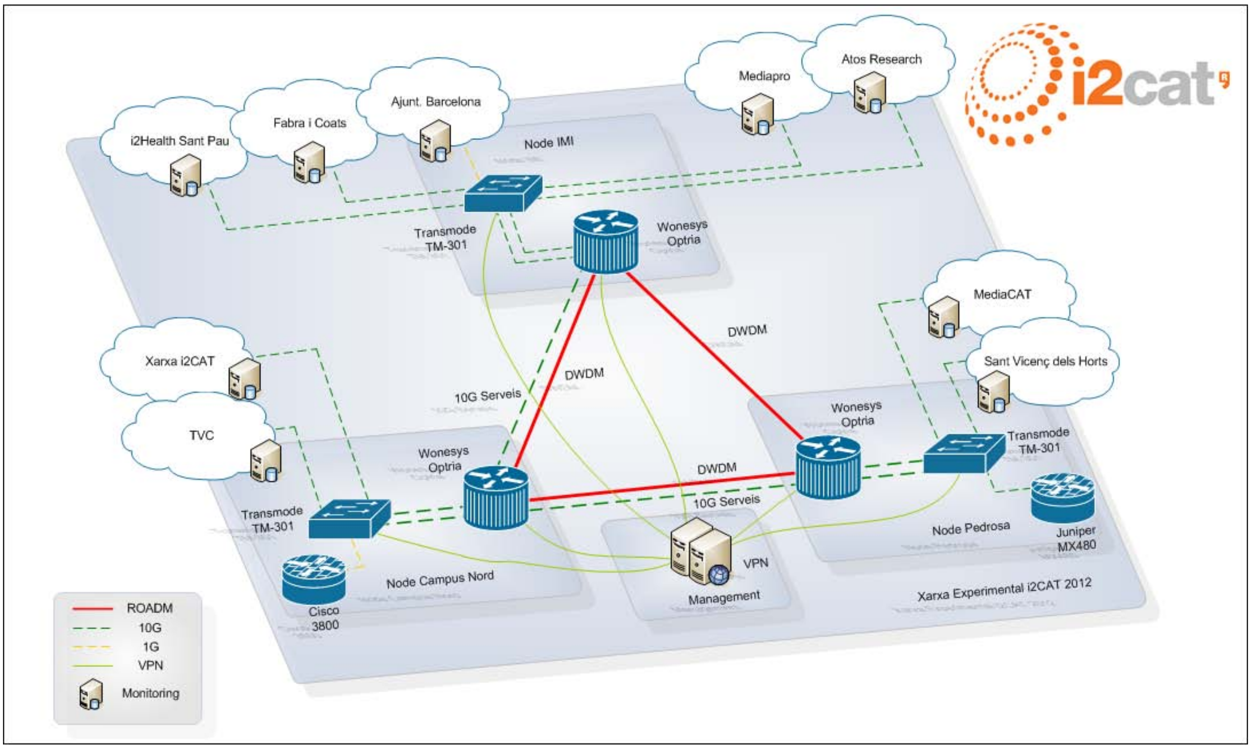

Figure 4-18: EXPERIMENTA network

Hardware:
5x OpenFlow switches (specification: NEC IP8800/S3640-24T2X2, 24 ports (1/10GE)
3x OF-enabled PRONTO switches TN3290, included two SFP+ optical transceivers for each switch
3x HP E3500-48G-PoE+ yl OpenFlow-Enabled switches
3x SuperMicro SYS-6016T-T NEHALEM E5506 2,13 GHz 12GB RAM 2TB disk
2x SuperMicro SYS-6016T-T WESTMERE 2,4 GHz 12GB RAM 2TB disk
5x SuperMicro SuperServer 6016T-T with x2 XEON E5620 2,4GHz 48GB RAM 2TB disk

Table 12 - EXPERIMENTA hardware

Software:
15VMs running the OFELIA OpenFlow reference implementation
Storage sharing (NFS, FTP, and SMB)
Code Repositories (SVN & GIT)
Source Code Management and Issue Tracker (JIRA, GitHub, FishEye)
Automated Testing Platform (Bamboo)
Infrastructure Management Platform: OpenNaaS
Cloud Management Platform: OpenNebula / OpenStack

OFELIA control framework: Resource aggregates (OF resource, virtual machines), AAA & Policy engine and user interface (Expedient).
OpenNaaS framework

Table 13 - EXPERIMENTA software

4.6.3. EXPERIMENTA SDN Services

OpenFlow in Europe: Linking Infrastructure and Applications (OFELIA) is a collaborative project within the European Commission's FP/ ICT Work Programme [71]. The project creates a unique experimental facility that allows researchers to not only experiment "on" a test network but to control and extend the network itself precisely and dynamically. The OFELIA facility is based on OpenFlow, a currently emerging networking technology that allows virtualization and control of the network environment through secure and standardized interfaces. OpenFlow is an open standard that allows to run experimental protocols in production networks. OFELIA allows controlling and extending the network itself precisely and dynamically.

OFELIA facility exists as a federation of heterogeneous experimental facilities with a homogeneous control framework. OFELIA is composed by several islands. One of them belongs to i2CAT. The i2CAT OFELIA island is integrated in EXPERIMENTA platform and composed by 5 NEC IP8800/S3640-24T2XW, 3 HP E3500-48G-PoE+ yl OF-Enabled switches and 3 servers. The switches are deployed outside the island at edges connected to the ROADMs so they can dispose of OpenFLow. The Figure 4-19 is a description of how the interconnection between the OpenFlow switches and the DWDM ring is connected.

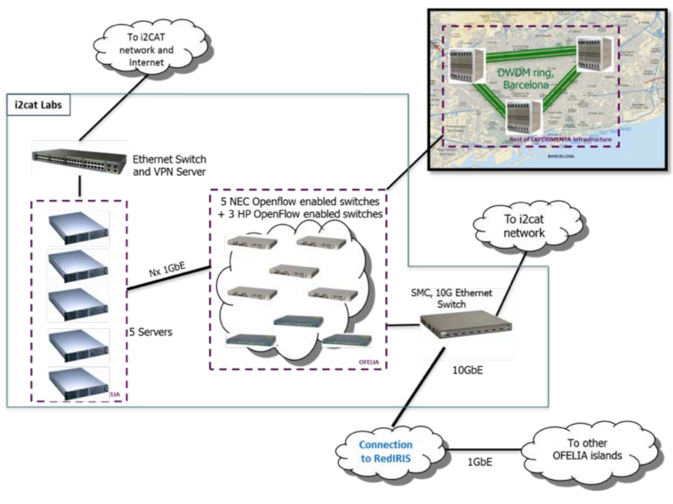

Figure 4-19: Interconnection between OFELIA and the DWDM ring

4.6.4. Experimental Management Tools

4.6.4.1. OFELIA Control Framework (OCF)

The OFELIA experimental control framework is capable of provisioning of experimental facility in single and federated OFELIA islands. OFELIA architecture is a framework for the specification of the components of a federated experimental facility and their functional organization and configuration, its operational principles and procedures, as well as data formats (data types and information exchange models) used in its operation.

The OFELIA Control Framework (OCF) is a set of software tools for test-bed management created in the OFELIA project [71]. It supports OpenFlow resources, Virtual Machines (XEN) reservation and bootstrapping. OCF allows the configuration of network slices for experimentation and also the related management to each OFELIA Island Managers (IM). In this way, OCF assures the control of the overall experimentation during all its life-cycle including the reservation, instantiation, configuration, monitoring, and un-instantiation. The OCF architecture, uses the Global Environment for Network Innovations (GENI)'s [72]. Expedient that, with its pluggable architecture, allows adding different Aggregate Managers (AMs) for each supported resource. In particular, the graphical user interface (GUI) provided by Expedient allows the underlying FlowVisor configuration guaranteeing the set-up of slices in an experimental network.

OCF has extended the original Expedient architecture (which included only Opt-in manager for the OpenFlow resources) by adding two new modules: Virtualisation AM (VT-AM) enabling users to allocate VMs on physical servers and a centralized Lightweight Directory Access Protocol (LDAP) which is the responsible for the user authentication. Further AMs are being developed currently to control and configure new type of resources such as wireless access points and optical switches.

4.6.4.2. OpenNaaS

OpenNaaS (Open Network as a Service) is an open platform for resource management [73]. The OpenNaaS framework was created in order to deploy and operate innovative NaaS paradigms. It allows, among others, on-demand provisioning of network resources using web services and abstracting resources from vendor and model details. The software is implemented based on two key components: resources and capabilities. The resources contain the information regarding the logical model of the device. The capabilities represent its features.

Currently, OpenNaaS supports different types of resources: router, BoD (Bandwidth on Demand), optical switch, mac bridge, network, vCPE, OpenFlow switch and SDN Network. The modular architecture of OpenNaaS allows the implementation of new types of devices. OpenNaaS is management platform that offers the essential tools for implementing extensions to support DOLFIN. The OpenNaaS GUI is a web application where a user can manage the network resources and shows information related to the switches configuration. The GUI can be extended to include DOLFIN monitoring data visualization.

4.6.4.3. Experimental Validation Services and Monitoring Tools

Lightweight IP Measurement System (LIMS) is a validation service provided by Tecsidel [61]. It is a solution based on a light ad-hoc development, structured in a LEAs (Linux Embedded Applications Systems) environment to obtain, collect, store and present active quality continuous measurements on user access by means of probes and their testing servers.

In addition, EXPERIMENTA is provided by a set of tools to measure the transmission speed, latency, packets loss and jitter, as well as alarms to network failure management.

4.6.5. I2CAT Testbed Summary

At following the most important EXPERIMETA concepts and objectives:

- **Open Testbed Facility:**

- i2CAT EXPERIMENTA open testbed allows access to network resources as a Service at different network layers. Users can configure the network to fulfil their needs. The optical validation platform is designed to support multivendor equipment.

- **Advanced Experimental Services:**
  - i2CAT EXPERIMENTA platform offers multi-layer services to test network applications and devices at different OSI layers on a real environment.
- **Infrastructure as a Service:**
  - Physical network substrate is virtualized to offer virtual isolated slices to final users.
  - Users can manage and configure their virtual substrate through a dedicated software tool which deals with hardware particularities and manages concurrency.
- **Parallel disruptive validations and experiments:**
  - By virtualizing the network, parallel disruptive tests are supported running above the same physical infrastructure.
- **Promote an innovative environment to empower technology validation:**
  - The collaborative platform environment promotes research, development and innovation between public and private (PPP) entities.

5. Summary and Conclusions

Large data centers are industrial scale operations using as much electricity as a small town and sometimes are a significant source of air pollution in the form of diesel exhaust. A data center is a logically centralized physical and virtual repository for computation, storage, management, services and the dissemination of data and information organized around a particular body of knowledge or pertaining to a particular business. The software elements of these data centres are realised and based on computing cloud principle and technologies. In general the hardware of the data centres includes redundant server and network node capabilities, as well as backup power supplies, redundant data communications connections, environmental controls (e.g., air conditioning, fire suppression) and various security devices.

DOLFIN's primary objective is to design, develop, and validate a Data Centre platform capable of monitoring the energy usage of the Data Centre and react accordingly for efficient energy management. The design of the DOLFIN DC framework would explicitly accommodate energy management and it is aimed at:

- Improving capital and operational efficiencies for DC operators through the use of a common organization, automation, and operations of all energy functions across the different domains
- Migrating from an ecosystem of separate energy management functions towards a coordinated arrangement of energy management functions as represented in the following figure:

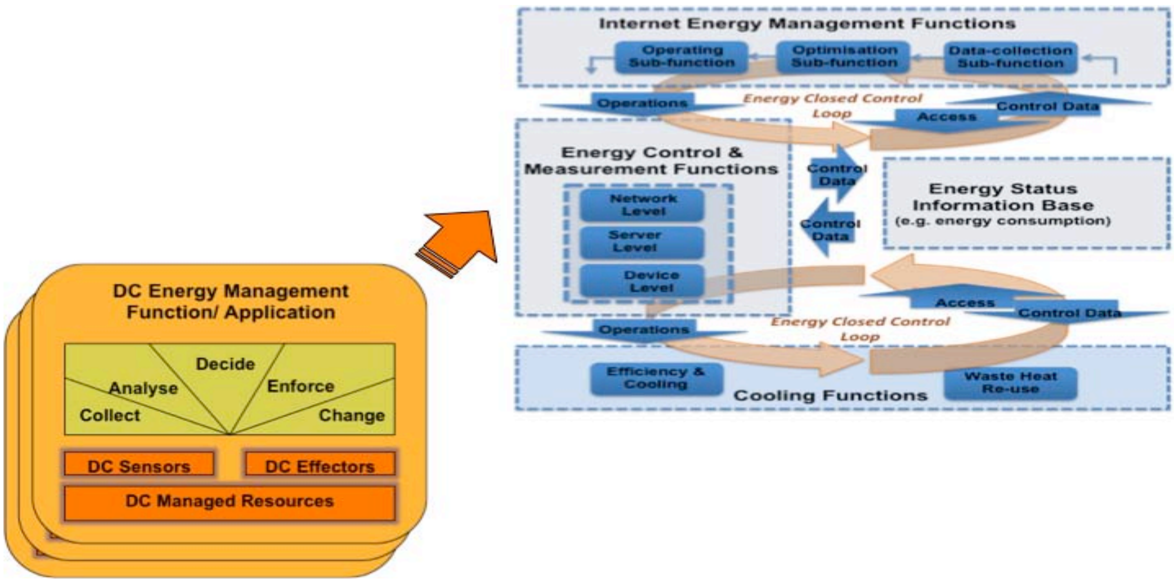

Figure 5-1: Migration from separate energy control loops to a coordinated arrangement of multiple DC energy control loops

The overall DOLFIN requirements list and design goals are derived individual project partners expertise, as well as the general vision and research directions for Future DC and Future Internet. The DOLFIN requirements list has two axes:

- a "***bottom-up/row requirements***" expressed through 24 use-case problem specific requirements addressing DC operators day-to-day problems identified in live DC and on existing DC architectures. The use-cases as presented in deliverable D2.1 were prioritised via a QFD analysis. 478 row functional requirements were identified in terms if interactions between key actors, functional role and links with the relevant components.
- "***vertical requirements/aggregated requirements***" synonymous with high-level functions, new/updated functional blocks, and inter-working interfaces in a DC. A similarity and consolidation exercise was performed on the 478 row functional requirements with a view to generate a higher level aggregation of the requirements with the following characteristics: Consistent and Atomic, Complete, Dependable, Current and Verifiable. From this process, 39 aggregated requirements were identified.

The first approach, "bottom-up requirements", aims at addressing the set of requirements elicited for the 24 use cases defined and developed so far within deliverable D2.1 – Business scenarios and use case analysis. The second approach, "vertical requirements", aims at elaborating the new management functionality in DCs or updating the existing management functionality in DCs.

The requirements together as a set, and not necessarily per individual requirement, describe what distinguishes DOLFIN from earlier DC management technologies and what the DOLFIN project intends to design and deliver.

The following is a synthesis of the main DC energy requirements and characteristics which are the basis of the DOLFIN architecture.

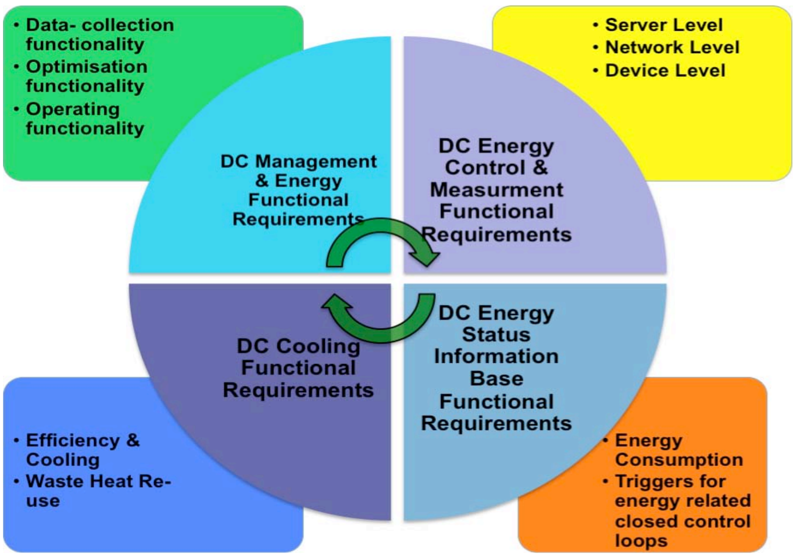

Figure 5-2: DOLFIN Requirements Synthesis

A detailed description of relevant energy efficiency metrics and KPI (Key Performance Indicators) have been produced, with the main purpose of providing background information on how to monitor and measure the achievement of DOLFIN deployment within Data Centre infrastructure.

For each key actor in the DOLFIN project (i.e. DC Operator initiator, DC Operator responder, Aggregator, End Customer, Utility Operator, Authority) we evaluated correlations between the aggregated requirements and business objectives w.r.t. DOLFIN project expected outcomes. In other words we evaluated how each aggregated requirement is contributing to solve/enable each business objective / expected project outcome as seen from each key actor. These are categorised as:

- ***Business Objectives:***

  BO1.1 - Reduce Server Costs,

BO1.2 - Reduce Infrastructure Costs,

BO1.3 - Reduce Power Costs,

BO1.4 - Reduce Other Costs,

BO1.5 - Reduce all operational costs,

BO2 - Increase the performance of offered services,

BO3 - Increase customer satisfaction,

BO4 - Improve the overall efficiency of the energy system,

BO5.1 - Shift processing power,

BO5.2 - Migrate processing power,

BO6 - Increase the availability /reliability of DC federation.

- ***DOLFIN Project expected outcomes:***

  O1 - Model, monitor, and measure energy consumption,

  O2 - Energy optimisation and trigger events for seamless and autonomic movement of VMs between servers of the same DC or across a group of energy-conscious synergetic DCs respecting/renegotiated SLAs,

  O3 - Optimize DC energy consumption, by dynamically changing the percentage of active/standby servers and load per server,

  O4 - Optimise the cumulative energy consumption in a group of DCs (policy-based VMs allocation),

  O5 - Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs, following the electricity demand-response approach),

  O6 - Smart grid energy stabilisation, by dynamically changing the energy consumption/production requirements of DCs

- **Use Cases:**

  UC1 - Energy efficient workload redistribution,

  UC2 - Multi tariffs from the Utility companies,

  UC3 - Multi tariffs from the Utility companies,

  UC4 - SLA Renegotiation with end customers,

  UC5 - Optimize benefits/incentives from national/European authorities,

  UC6 - Smart City

- **Actors involved in DC operations in each use case***:*
  DC Initiator,
  DC Responder,
  DC Aggregator,
  End Customer,
  Utility Operator,
  Authority

The resulting aggregated requirements were ranked for importance and impact on the business goals and DOLFIN project expected outcome using the Quality Function Deployment (QFD) method. The ordering of the aggregated requirements will be used in the detailed design and implementation of the DOLFIN functional components.

The importance of aggregated requirements for realising business objectives based on combined all actors view points is presented in the following Figure 5-3.

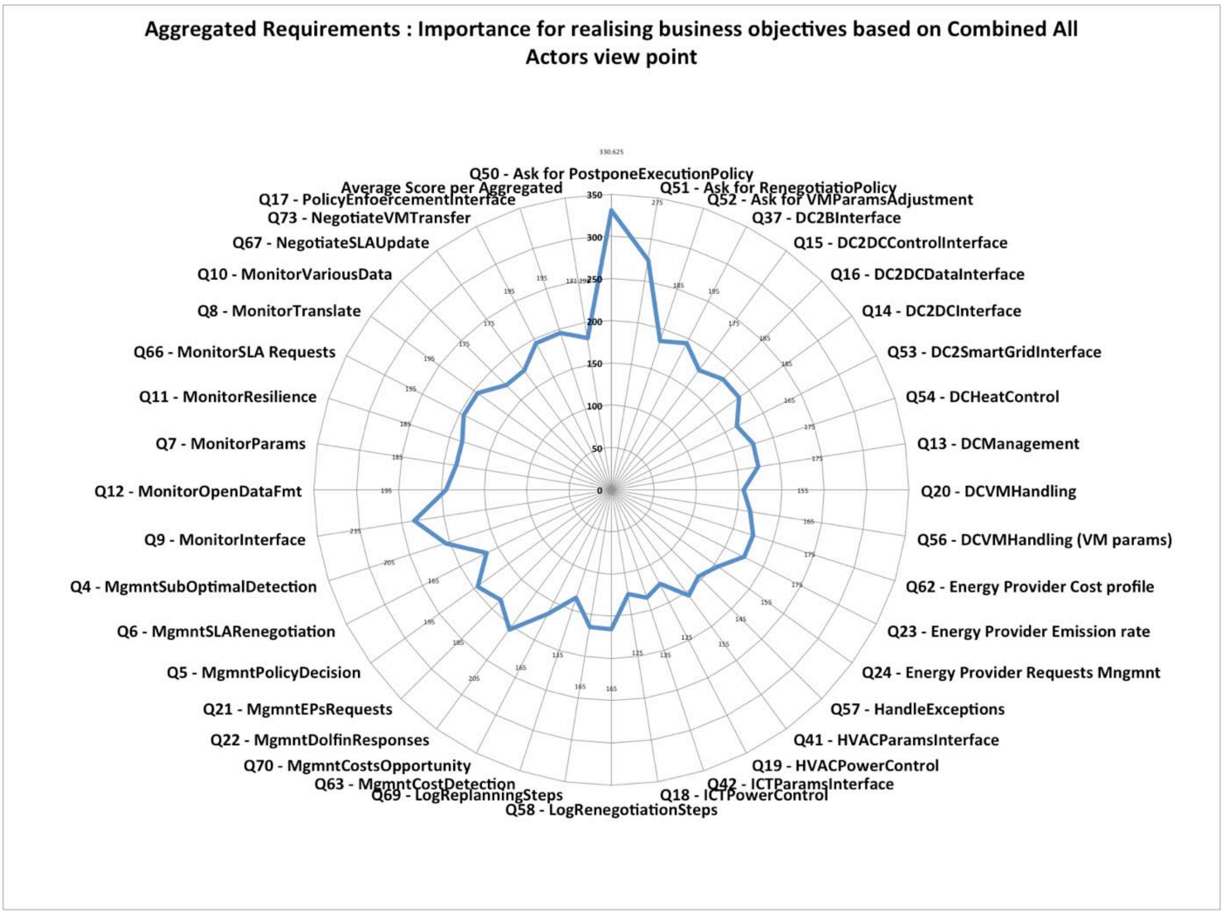

Figure 5-3: Importance for realising business objectives – combined all actors view points

The importance of aggregated requirements for realising project outcomes based on combined all actors view points is presented in the following **Figure 5-4**.

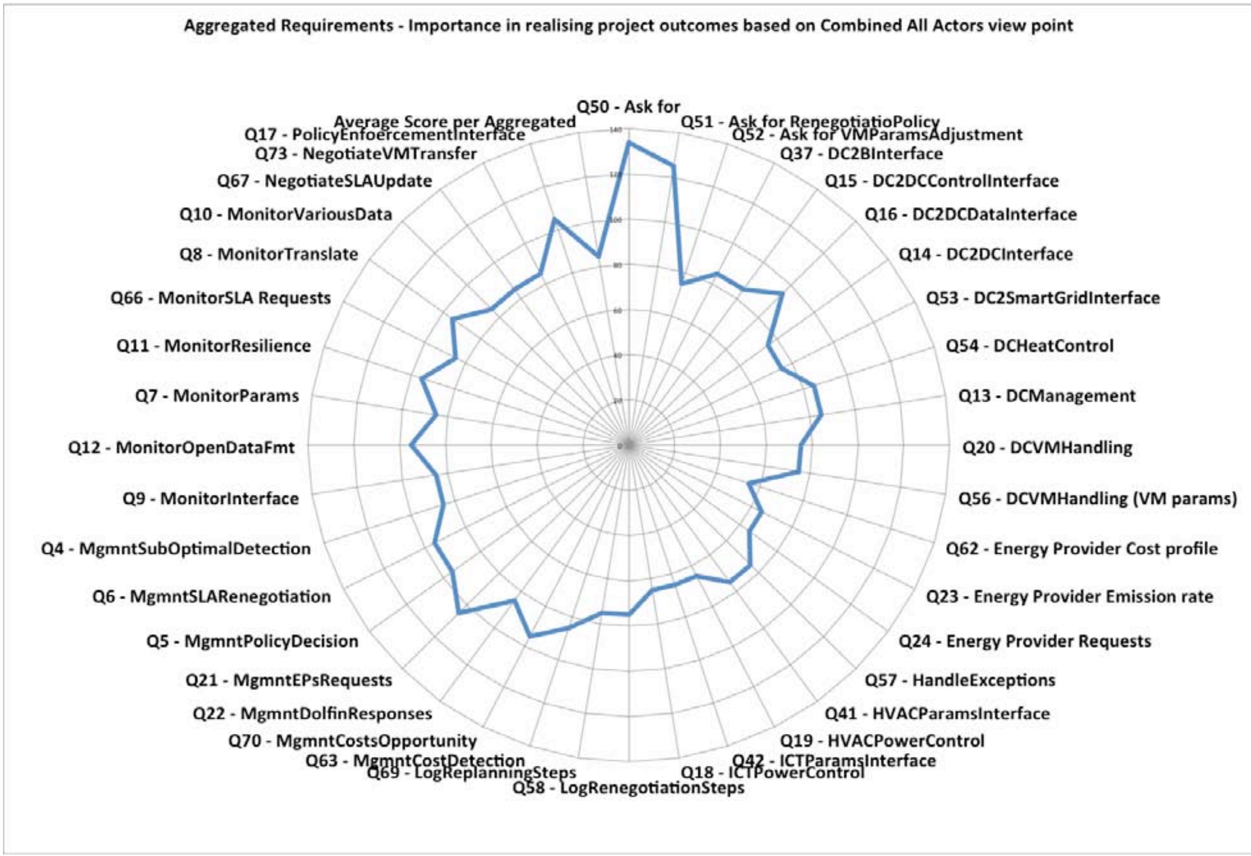

Figure 5-4: Importance for realising project outcomes – combined all actors view points

This report presents the first version of the DOLFIN architecture. It will be updated in deliverable D5.2 based on the validation analysis and experimental results planned for the next work periods of the project.

In this report the DOLFIN DC functional architecture is defined in terms of new DC functional blocks, revised DC functional blocks, and inter-working interfaces needed to realise different energy closed control loops and their interactions with the normal operation of the DC. The description of the DOLFIN architecture and its main functional blocks are derived from the energy efficiency requirements and it is targeted at future DCs with the following key new characteristics:

- SLAs, Orchestration, and Energy control loops per DC and/or group of DCs
- Optimisation of DC energy consumption, by dynamically changing the percentage of active/standby servers and load per server
- Optimisation of the cumulative energy consumption in a group of DCs (policy-based VMs allocation)
- Optimisation of the energy consumption at the DS / Smart City level and provide energy stabilization, by distributing VMs across a group of DCs, following the electricity demand-response approach
- Smart grid energy stabilisation, by dynamically changing the energy consumption/production requirements of DCs

The main DOLFIN functional components are:

- The **Energy Consumption Optimisation Platform (eCOP)**, which is the DOLFIN core platform for energy consumption optimization at Data Centre level. It provides facilities for continuous monitoring, energy benchmarking, dynamic control and adaptive optimisation of the Data Centre infrastructure, including ICT and HVAC equipment.
- The **energy-conscious Synergetic Data Centres (SDC)** module, which provides a dynamic, service-effective and energy-efficient allocation of resources across a distributed network of co-operating DCs, and
- The **Energy Broker**, which extends the **Smart-Grid Energy stabilisation module**, and controls the interconnection with the smart grid network, providing responses on the changing demands for energy, with a module that controls the legacy (brown) Energy Providers.

The functional architecture refers only to the software platform in a DC and it is depicted in the following Figure 5-5.

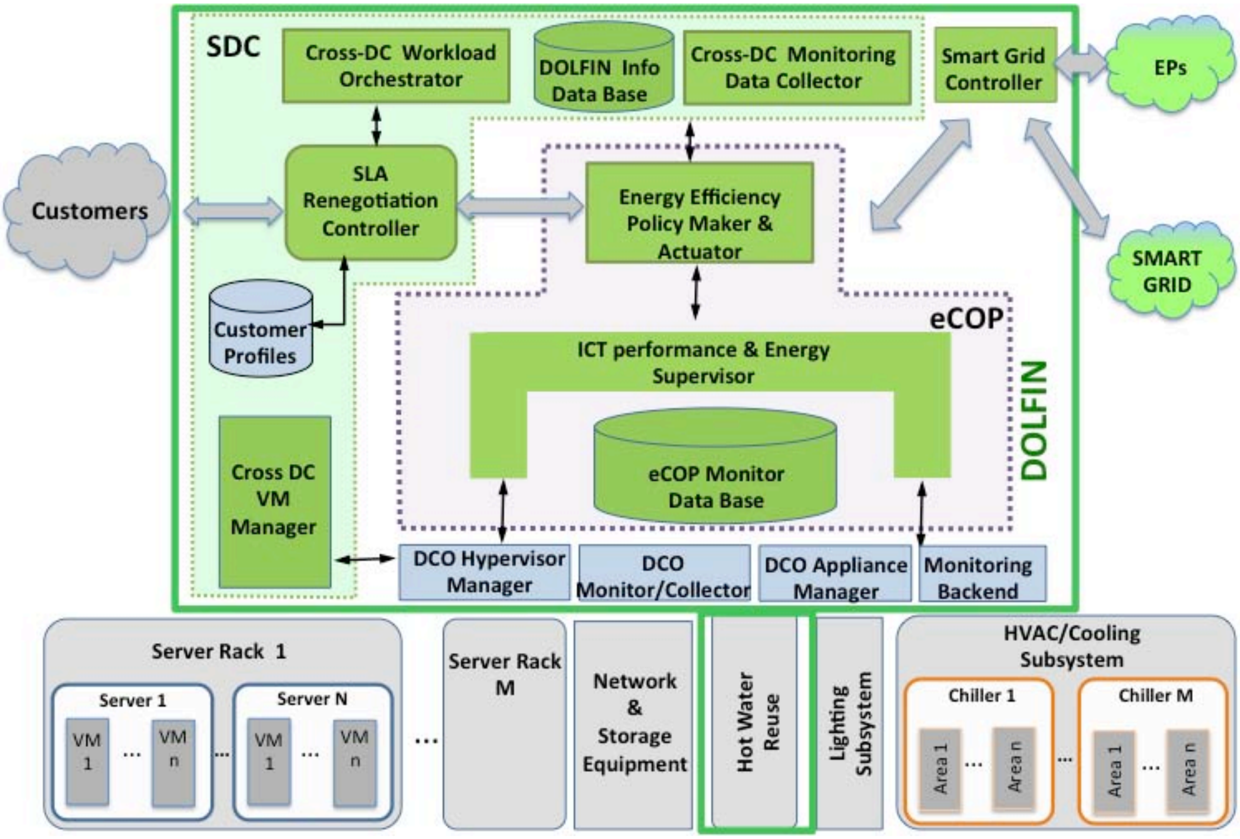

Figure 5-5: Detailed DOLFIN Functional Architecture

The new DC functional blocks, the revised DC functional blocks, and the inter-working interfaces that are needed to realise different energy closed control loops, together with the WP responsibility, are summarised and presented in Figure 5-6 as follows:

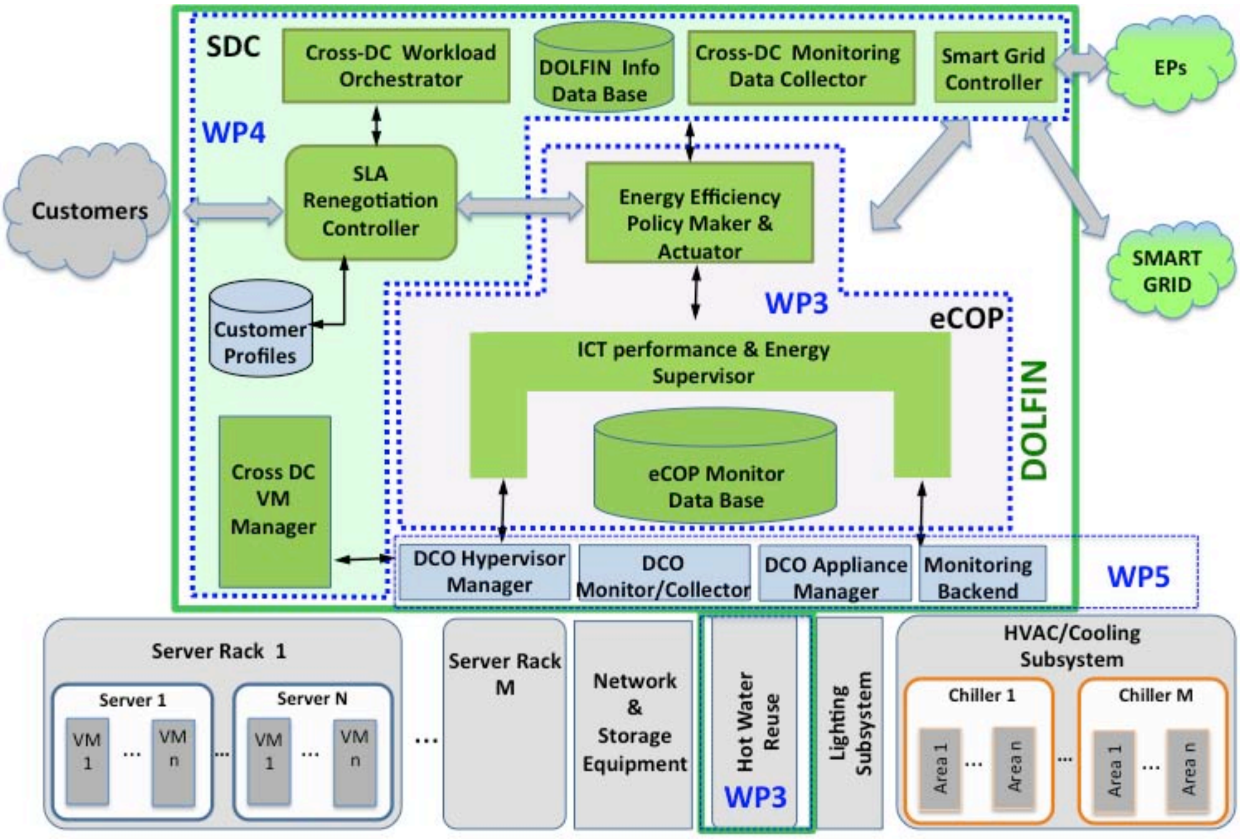

Figure 5-6: DOLFIN S/W Components with revised /new DC functionality

To summarise, the main blocks are:

- **DCO Hypervisor Manager (revised DC functionality)** – The DCO Hypervisor Manager is an adaptation layer that maps the high-level decisions taken by the DOLFIN system into low-level technology-dependent commands which control the DC ICT infrastructure.
- **DCO Appliance Manager (revised DC functionality)** – The DCO Appliance Manager adapts and guarantees the execution of strategies elaborated by DOLFIN expert system on the non-ICT infrastructure by translating high level commands into low-level technology-agnostic actions, and then realises the real operations, specific to the DC control framework, either directly or via the relevant BMS system.
- **DCO Monitor/Collector (revised DC functionality)** – The role of the DCO monitor/collector is to interface with both the ICT and the non-ICT DC infrastructure and collect all operational and energy related information to be stored in the eCOP Data Base.
- **eCOP Monitor Data Base (new DC functionality)** – The eCOP Monitor DB has the sole purpose of storing all real-time and historical energy related data collected from DCs. This information is used for energy efficiency decision and VM load prediction, along with data from all the components within a DC architecture. The eCOP Monitor database interfaces with the DCO Hypervisor and the DCO Appliance Managers, which store the relevant information to the database. The eCOP Monitor DB may be also accessed by a monitoring backend component, which can be connected with a ***dashboard*** to offer a Graphical User Interface (GUI) monitor tool on top of the ICT Performance

and Energy Supervisor system. The main objective of this interface is to provide an easy way to consult all the relevant information related to ICT Performance and Energy Supervisor.

- **ICT Performance & Energy Supervisor (new DC functionality)** – The ICT Performance and Energy Supervisor focuses on the analysis of performance monitoring data and energy data. The objective of this component is to provide information on the actual performance of the applications (typically VMs) utilizing the resources of the DC devices and the energy consumed by the non-ICT components.
- **Energy Efficiency Policy Maker & Actuator (new DC functionality)** – This is the most intelligent part of the eCOP, that makes the decisions realizing the requested policy of the DC.

To achieve energy efficiency at the cloud level, the DOLFIN eCOP collaborates with the ***Synergetic Data Centres (SDC)*** module, which consists of:

- **Cross-DC Monitoring Data Collector (new DC functionality)** – This collects knowledge not only from the SDC resources, but also from the network routers and in general from a network resources point of view. This is an important prerequisite to achieve energy efficiency across DCs.
- **Cross-DC Workload Orchestrator (new DC functionality)** – This is a distributed software element which gets the decisions and does most of the different types of resource optimisation including energy optimisation and management of the trade-offs in a cross DC optimization scenario. It is in charge of managing the full lifecycle of the virtual routers in the network and the allocation of the applications running on the virtual nodes
- **Cross-DC VM Manager (new DC functionality)** – This provides a DC Interconnect (DCI) interface and performs the actual migration of VMs between DCs. It will typically apply a set of standard alternatives for coping with high/peak workloads, more precisely with allocation of VMs, data, services and tasks.

The Cross DC Monitoring Collector, Workload Orchestrator, and VM Manager (also referred together as ***Workload and VM Manager***) closely collaborate in order to offer a flexible, reliable, and fast communications solution to handle the increased network traffic arising from the DCs operation.

- **DOLFIN Information Data Base (new DC functionality)** – This offers abstracted and logically-centralised information manipulation (including information collection, aggregation/processing, storage / indexing and distribution) across all DOLFIN architectural components. Moreover, it includes information of the local energy requirements e.g. Demand/Response requests from the local Smart Grid operator.
- **SLA Renegotiation Controller (new DC functionality)** – The SLA Renegotiation Controller guarantees to steer the data centre operation towards environmentally sound behaviour. This module will take into account the existing approaches in modelling SLA criteria (e.g. overall cost of an offered service) and augment them with the use of energy related criteria.
- **Customer profile management (revised DC functionality)** – DC and cloud infrastructure users exhibit a wide variety of profiles and requirements as they originate from different sectors such as industry, academia, education, or research and are widely geographically and culturally distributed. To that extent, the customer profiles management are greatly influenced by the technical requirements of the services requested by the DC of cloud infrastructure, but also by the financial, societal and even moral status of the customer.

The DOLFIN System includes the **Smart Grid Controller (SGC)** which is a Smart City Component. The SGC interconnects the DC with the Smart city electricity network. This most critical part of the system acts as a broker between the core DOLFIN infrastructure and subsystems, represented by the Energy Efficiency Policy Maker & Actuator, and the outside world, namely the Smart Grid (SG) and the Energy Providers (EPs).

This report presents the testbeds (3 enterprise data centre test beds: WIND testbed, IRT testbed, GRNET testbed + 2 research data centre testbeds: UCL test bed, i2CAT testbed) aimed at interoperability experiments and evaluation and being enhanced with the DOLFIN architecture and its components. Detailed results of QFD analysis are in terms of ranked aggregated requirements as they contribute most/least to the realisation of business objectives and of project outcomes are also presented in this report. These results will be used as input to the detailed design and implementation of the DOLFIN system components.

6. References

[1] Data Centre Dynamics, 2011 “Global industry census 2011” available on-line at: http://www.dataCentredynamics.com/research/market-growth-2011-2012

[2] Open Data Center Alliance: Carbon Footprint and Energy Efficiency Rev. 2.0, 2013.

[3] S. Pelley, et al., “Understanding and Abstracting Total Data Center Power,” Workshop on Energy Efficient Design (WEED), 2009.

[4] L. Minas, B. Ellison, Energy Efficiency for Information Technology: How to Reduce Power Consumption in Servers and Data Centers, Intel Press, USA, 2009.

[5] EMERSON NETWORK Power, White paper, "Energy Logic: Reducing Data Center Energy Consumption by Creating Savings that Cascade Across Systems", 2008.

[6] Beloglazov, R. Buyya, Y. C. Lee, A. Zomaya, “A Taxonomy and Survey of Energy-Efficient Data Centers”, Advances in Computers, Vol. 82, Elsevier, 2011.

[7] E. Le Sueur and G. Heiser, “Dynamic voltage and frequency scaling: The laws of diminishing returns,” in Proceedings of the 2010 International Conference on Power Aware Computing and Systems, ser. HotPower’10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–8. [Online]. Available: http://dl.acm.org/citation.cfm?id=1924920.1924921

[8] T. Burd and R. Brodersen, “Energy efficient cmos microprocessor design,” in Proceedings of the Twenty-Eighth Hawaii International Conference on System Sciences, vol. 1, Jan. 1995, pp. 288–297

[9] G. Warkozek, E. Drayer, V. Debusschere, S. Bacha, “A new approach to model energy consumption of servers in Data Centers”, in Proc. of IEEE ICIT 2012, 2012.

[10] U. Mandal, M. F. Habib, S. Zhang, B. Mukherjee, and M. Tornatore, “Greening the cloud using renewable-energy-aware service migration.” IEEE Network, vol. 27, no. 6, pp. 36–43, 2013.

[11] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, “Elastictree: saving energy in data center networks,” USENIX NSDI, pp. 17–17, 2010.

[12] X. Wang, Y. Yao, X. Wang, K. Lu, and Q. Cao. Carpo, “Correlation-aware power optimization in data center networks,” IEEE INFOCOM,pages 1125–1133, 2012.

[13] US Department of Energy, http://www.doe2.com

[14] EnergySoft website, http://www.energysoft.com/

[15] eQUEST webpage, http://www.doe2.com/equest/

[16] http://www.pge.com/includes/docs/pdfs/mybusiness/energysavingsrebates/incentivesbyindustry/hightech/data_center_baseline.pdf

[17] Pedram, M., "Energy-Efficient Datacenters," Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on, vol.31, no.10, pp.1465,1484, Oct. 2012

[18] J. Moore, J. Chase, P. Ranganathan and R. Sharma, “Making scheduling “cool”: Temperature-aware workload placement in data centers.” ATEC '05 Proceedings of the annual conference on USENIX Annual Technical Conference, Anaheim, CA, 2005.

[19] Lin Wang; Fa Zhang; Arjona Aroca, J.; Vasilakos, AV.; Kai Zheng; Chenying Hou; Dan Li; Zhiyong Liu, "GreenDCN: A General Framework for Achieving Energy Efficiency in Data Center Networks," Selected Areas in Communications, IEEE Journal on, vol.32, no.1, pp.4,15, January 2014

[20] Hao Shen, Ying Tan, Jun Lu, Qing Wu, Qinru Qiu: Achieving autonomous power management using reinforcement learning. ACM Trans. Design Autom. Electr. Syst. 18(2): 24 (2013)

[21] CoolEmAll project, http://www.coolemall.eu/

[22] ISO/IEC JTC 1/SC 39 Sustainability for and by Information Technology, http://www.iso.org/iso/standards_development/technical_committees/other_bodies/iso_technical_committee.htm?commid=654019

[23] Standards and projects under the direct responsibility of ISO/IEC JTC 1/SC 39 Secretariat, http://www.iso.org/iso/home/store/catalogue_tc/catalogue_tc_browse.htm?commid=654019&development=on

[24] ETSI GS OEU 001, Operational energy Efficiency for Users (OEU), Technical Global KPIs for Data Centres, http://www.etsi.org/deliver/etsi_gs/OEU/001_099/001/01.02.02_60/gs_oeu001v010202p.pdf

[25] ITU-T Recommendation Y.3021, "Framework of energy saving for Future Networks", June 2012 (http://www.itu.int/rec/T-REC-Y.3021/en).

[26] ITU-T Focus Group on Future Networks, http://www.itu.int/en/ITU-T/focusgroups/fn/Pages/Default.aspx.

[27] BSC DC Specialist Group, http://dcsg.bcs.org

[28] http://dcsg.bcs.org/sites/default/files/protected/data-centre-energy.pdf

[29] GreenIT webpage, http://www.greenit.net/

[30] See EC DC Cluster Collaboration initiative Tasks 1-4.

[31] The Green Grid, PUE: A Comprehensive Examination of the Metric, http://www.thegreengrid.org/en/Global/Content/white-papers/WP49-PUEAComprehensiveExaminationoftheMetric

[32] Green IT Promotion Council, New metrics for Data Center Energy Efficiency - http://home.jeita.or.jp/greenit-pc/topics/release/pdf/dppe_e_20120824.pdf

[33] The Green Grid, WP#29-ERE: A Metric for Measuring the Benefit of Reuse Energy from a Data Center, http://www.thegreengrid.org/en/Global/Content/white-papers/ERE#sthash.siFnYQfj.dpuf

[34] See EC DC Cluster Collaboration Task 3.

[35] SuperMUC supercomputing system, http://www.lrz.de/services/compute/supermuc/

[36] Ebay project Mercury, http://www.thegreengrid.org/~/media/Case%20Studies/CS3-Breaking_New_Ground_on_Data_Center_Efficiency.pdf

[37] S. Andreozzi, N. De Bortoli, S. Fantinel, A. Ghiselli, G. L. Rubini, G. Tortone, and M. C. Vistoli, "GridICE: A monitoring service for grid systems," Future Gener. Comput. Syst., vol. 21, no. 4, pp. 559–571, 2005.

[38] S. Clayman, A. Galis, C. Chapman, G. Toffetti, L. Rodero-Merino, L. Vaquero, K. Nagin, and B. Rochwerger. Monitoring Service Clouds in the Future Internet. In Towards the Future Internet - Emerging Trends from European Research, pages 115–126, Amsterdam, The Netherlands, The Netherlands, 2010. IOS Press.

[39] S. Clayman, A. Galis, L. Mamatas, Monitoring virtual networks with Lattice. In: Management of Future Internet - ManFI 2010, URL http://www.manfi.org/2010/

[40] Cooke, A. J. G. Gray, L. Ma, W. Nutt, et al., “R-GMA: An information integration system for grid monitoring,” in Proceedings of the 11th International Conference on Cooperative Information Systems, pp. 462–481, 2003.

[41] M. L. Massie, B. N. Chun, and D. E. Culler, “The ganglia distributed monitoring system: Design, implementation and experience,” Parallel Computing, vol. 30, p. 2004, 2003.

[42] “Nagios.” http://www.nagios.org/.

[43] H. Newman, I. Legrand, P. Galvez, R. Voicu, and C. Cirstoiu, “MonALISA : A distributed monitoring service architecture,” in Proceedings of CHEP03, La Jolla, California, 2003.

[44] Xen, http://www.xen.org/

[45] VMware, http://www.vmware.com/

[46] KVM, http://www.linux-kvm.org/

[47] "ORACLE DATABASE APPLIANCE: Delivering “easy” to midsize Oracle database environments", http://www.oracle.com/us/products/database/database-appliance/deliver-easy-to-midsize-env-wp-2202033.pdf

[48] IBM, WebSphere Appliance Management Center for WebSphere Appliances, http://www-03.ibm.com/software/products/en/wamc

[49] ABB Decathlon, http://new.abb.com/cpm/decathlon-software/decathlon-datacenters

[50] Intel, “Data Center Manageability Interface (DCMI) Specification”, White paper http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/data-center-manageability-interface-paper.pdf

[51] Intel, HP, NEC, Dell, “Intelligent Platform Management Interface (IPMI) 2.0,: White Paper, Oct. 2013, http://www.intel.com/content/www/us/en/servers/ipmi/ipmi-second-gen-interface-spec-v2-rev1-1.html

[52] CISCO EnergyWise Suite, http://www.cisco.com/c/en/us/products/switches/energywise-technology/index.html

[53] IETF Energy Management (EMAN) WG, https://datatracker.ietf.org/wg/eman/charter/

[54] ERWIN Data Modeler http://erwin.com/

[55] The Green Grid Data Centre Power efficiency Metrics: PUE and DCiE, White Paper, The Green Grid, 2007.

[56] ITU-T Recommendation Y.3021, “Framework of energy saving for Future Networks”, June 2012 (http://www.itu.int/rec/T-REC-Y.3021/en).

[57] ITU-T Focus Group on Future Networks, http://www.itu.int/en/ITU-T/focusgroups/fn/Pages/Default.aspx.

[58] BSC DC Specialist Group, http://dcsg.bcs.org.

[59] http://dcsg.bcs.org/sites/default/files/protected/data-centre-energy.pdf.

[60] ETSI Operational energy Efficiency for Users: Technical Global KPIs for DCs, http://www.etsi.org/deliver/etsi_gs/OEU/001_099/001/01.01.01_60/gs_OEU001v010101p.pdf

[61] CISCO White paper, ”Distributed Virtual Data Center for Enterprise and Service Provider Cloud”, http://www.cisco.com/c/en/us/products/collateral/routers/asr-9000-series-aggregation-services-routers/white_paper_c11-694882.pdf

[62] CISCO, Virtual Private LAN Services (VPLS), http://www9.cisco.com/c/en/us/products/ios-nx-os-software/virtual-private-lan-services-vpls/index.html

[63] CISCO, Advanced VPLS, http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization /data-center-interconnect/design_guide_c07-647422.html

[64] CISCO, Overlay Transport Virtualization (OTV), http://www.cisco.com/c/en/us/solutions/data-center-virtualization/overlay-transport-virtualization-otv/index.html

[65] CISCO, Ethernet over MPLS, http://www.cisco.com/c/en/us/td /docs/switches /lan/catalyst6500/ ios/ 12-2SY/configuration/guide /sy_swcg/eompls.html

[66] IETF, E-VPN working group, http://tools.ietf.org/html/draft-sajassi-raggarwa-l2vpn-evpn-req-00

[67] IETF, Transparent Interconnection of Lots of Links (trill) working group, https://datatracker.ietf.org/wg/trill/documents/

[68] NEC's High Temperature Ambient Server, http://www.nec.com/en/global/prod/ express/related/design_technology/collateral/Technologies_Behind_NECs_HTA_Server.pdf

[69] http://www.datacenterknowledge.com/archives/2012/06/25/hot-water-cooling/

[70] GLIF project website: http://www.glif.is/

[71] OFELIA Project website: http://www.fp7-ofelia.eu/

[72] GENI Project website: http://www.geni.net/

[73] OpenNaaS website: http://www.opennaas.org/

[74] Tecsidel http://www.tecsidel.com/en/products/monitoring-and-oss-ict/measurement-of-quality-in-ip-lims

7. Abbreviations and Definitions

7.1. Definitions

Term	**Definition**
Utility Operator	This Actor provides electricity using fixed or multi-tariff pricing schemes. In case of multi-tariff, the cost per electricity unit could vary depending on the day of the week, the time of day, certain electricity consumption limits or, possible, other parameters. The objective of this Actor is to provide electricity to consumers, maintaining stability between demand and supply.
Authority	This Actor represents any entity that, from a regulatory or policy perspective, is posing rules on the way that the DOLFIN system should operate or interact with any external system. In this context, environmental, energy, data security and processing related rules might be applicable. The objective of an Authority is to protect the public interest and to promote a healthy market competition. To this respect it can influence the framework of operation of DCOs, Aggregators and Customers, determining also the way they interact and exchange information.
Customer	This actor consumes ICT services offered by the DCO. Usually, this entity utilises the ICT services so as to store, retrieve, process and transfer business information under a specific contract and SLA guarantees provided by the DCO. The objective of a Customer is to receive ICT services with a guaranteed SLA that correspond to a reasonable and fair price, or even favourable, in case such an opportunity could arise. A Customer may also opt to participate in a demand-response curtailment program and receive some cost benefits and contribute to the implementation of policies opting for green-powered ICT services and Data Centers.
Aggregator	This actor acts as a service provider by enlisting electricity consumers to participate in demand-response curtailment programs and sell the combined load reduction to electricity retailers or Distribution System Operators (DSO). Typically, the Aggregator takes a percentage of the demand-response incentive as compensation, passing the rest on to the consumer. To this respect, the Aggregator could represent a number of different entities, such as Utility Operators, DSOs, Distributed Energy

	Resources (DER) or Grid Operators, in an effort to avoid electricity consumption peaks by offering demand-response programs to DCOs. In the simplest scenario, the role of the Aggregator could also be played directly by the entities that it represents. The objective of the Aggregator is to maintain a balanced electricity supply and demand within the energy grid network that is under its immediate responsibility, by offering incentives to DCOs, so as to lower their electricity demands.
Data Center Operator	This actor is responsible for operating a DC or a group of Data Centers that are part of the same administrative domain. It is responsible for providing ICT services to customers, adhering to specific Service Level Agreements (SLA). The primary objective of a DCO is to offer ICT services under the promised SLA, maximising the utilisation of the ICT infrastructure, while maintaining the electricity consumption within levels that are proportional to the ICT processing work delivered. Furthermore, the DCO would like to take advantage of any demand-response curtailment incentives that would be available, so as to decrease its operational cost and provide greener ICT services. In order to benefit from the incentives, DCO might be required to take immediate actions or even schedule them for the near future.
Data Center	A data center is a centralized repository physical and / or virtual, for the storage, management, services and dissemination of data and information organized around a particular body of knowledge or pertaining to a particular business. The S/W of the data centres are realised and based on computing cloud principle and technologies. In general the H/W of the data centres would includes redundant server and network nodes capabilities as well as backup power supplies, redundant data communications connections, environmental controls (e.g., air conditioning, fire suppression) and various security devices. Large data centers are industrial scale operations using as much electricity as a small town and sometimes are a significant source of air pollution in the form of diesel exhaust.
Functional Requirement	It is a description of what a system is supposed to do and it defines of a function, or a feature of a system, or its components, capable of solving a certain problem or replying to a certain need/request. The set of functional requirements present a *complete* description of how a specific system will function, capturing every aspect of how it should work before it is built, including information handling, computation handling, storage handling and connectivity handling.
System Design	It is a plan for implementing functional requirements.
Non-functional requirement	It is a specification criteria that can be used to judge the operation of a system, rather than specific behaviours; it is a description of how well a system performs its functions; it represents an attribute that a specific

	system must have. The non-functional requirements are controlled by other aspects of the system.
Business objectives	It is a description in business terms of what must be delivered or accomplished to provide value.
System boundaries / limits	It defines the constraints and freedoms in controlling the system. Limits can be determined by analysing how the behaviour of the system depends on the parameters that drive the system. Some limits would lead to unexpected and significant behaviour changes of the system, for example the unpredictable boundaries or changes in the scale of magnitude. Some other limits are determined by non-common behaviour interactions between the components of a system.
Architecture	It is a plan for implementing non-functional and functional requirements within the system limits/boundaries. It is conceptual model that defines the structure, behaviour, and a number views of a system within the system limits
Quality Function Deployment (QFD)	It is a method for developing a design quality translating the user's demand into design targets and major quality assurance points to be used throughout the implementation phase.
Use Case (UC)	It is a descriptor of a set of precise problems to be solved. It describes steps and actions between stakeholders and/or actors and a system, which leads the user towards a value added or a useful goal. A UC describes what the system shall do for the actor and/or stakeholder to achieve a particular goal. Use-cases are a system modelling technique that helps developers determine which features to implement and how to gracefully resolve errors.
DC Governance	It is a framework, which enables operators to describe their goals and objectives, through high-level means and govern their network. Includes the derivation of DC policies from the business goals through the use of semantic techniques.
Actor	It is a person, group or organization with an interest in a specific viewpoint of a system.
Viewpoint	It is a representation of a whole system from the perspective of a related set of concerns.
Accessibility	It represents the degree to which a system, device, service, or environment is available to as many people as possible. Accessibility can be viewed as the "ability to access" and benefit from some system or entity.
Availability	It represents the degree to which a system is in a specified operable and committable state at the start of a task. It is the proportion of time a system is in a functioning condition.
Certification	It refers to the confirmation of certain characteristics of an object, element of system. This confirmation is often, but not always, provided

	by some form of external review, assessment, or audit.
Configuration	It is a function establishing and maintaining consistency of a system and/or its performance. It is changing system's functional and physical attributes with its non-functional requirements, design, and operational information throughout its life.
Compliance	It represents the conformance to a rule, such as a specification, policy, standard or regulation.
Extensibility	It represents the ability to extend a system and the level of effort and complexity required to realize an extension. Extensions can be through the addition of new functionality, new characteristics or through modification of existing functionality/characteristics, while minimizing impact to existing system functions.
Interoperability	*It represents the ability of diverse systems and subsystems to work together (inter-operate). It is also a characteristic of a system, whose interfaces are completely understood, to work with other systems, present or future, without any restricted access or implementation.*
Maintainability	It is a characteristic of design and installation, expressed as the probability that an element of a system will be retained in or restored to a specified condition within a given period of time, when the maintenance is performed in accordance with prescribed procedures and resources.
Operability	It is the ability to keep a system in a safe and reliable functioning condition, according to pre-defined operational requirements.
Performance	It describes the degree of performance of a system (according to certain predefined metrics, e.g. convergence time)
Privacy	It is the ability of system or actor to seclude itself or information about itself and thereby reveal itself selectively.
Resilience	It is the ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operations.
Reliability	It is the degree to which a system must work. Specifications for reliability typically refer to stability, availability, accuracy, and maximum acceptable/tolerable bugs.
Robustness	It is the ability of a system to cope with errors during execution or the ability of a system to continue to operate despite abnormalities in input or in environment context.

7.2. Abbreviations

BMS	Building Management System
CEN	European Committee for Standardization
CF	Carbon Footprint
CUE	Carbon Usage Effectiveness
DB	Data Base
DC	Data Centre
DCeP	DCP Data Centre energy Productivity
DCF	DC Federation
DCiE	Data Centre infrastructure Efficiency
DCIM	Data Centre Infrastructure Management
DCMM	Data Centre Maturity Model
DCO	Data Centre Optimization
DPPE	Data centre Performance Per Energy
DRS	Distributed Resource Scheduling
DVFS	Dynamic Voltage and Frequency Scaling
EC	European Commission
EMS	Environmental Management System
EP	Energy Providers
ERF	Energy Reuse Factor (= KPIREUSE)
ESO	European Standards Organization (i.e. CEN, CENELEC and ETSI)
ETS	Emissions Trading System
ETSI	European Telecommunications Standards Institute
EU	European Union
EUCoCDC	European Code of Conduct on Data Centre Energy Efficiency
FCoE	Fibre Channel over Ethernet
GEC	Green Energy Coefficient (= KPIREN)
GHG	Greenhouse Gases
GUI	Graphical User Interface

HVAC	Heating, Ventilation, and Air Conditioning
ICT	Information and Communication Technology
IDB	Information Data Base
IoT	Internet of Things
ISO	International Organization for Standardization
ITEE	IT Equipment Efficiency
ITEU	IT Equipment Utilization
ITU-T	International Telecommunication Union - Telecommunication standardization sector
KPI	Environmental Key Performance Indicator
KPIEC	Energy Consumption
KPIREN	Renewable Energy Use
KPIREUSE	Energy reuse (= ERF)
KPITE	Task Efficiency (= PUE)
PUE	Power Usage Effectiveness (= KPITE)
SDC	Synergetic Data Centres
SDO	Standards Development Organization
SG	Smart Grid
SGC	Smart Grid Controller
UPS	Uninterruptible Power Supply
VM	Virtual Machines
WF	Water Footprint
WUE	Water Usage Effectiveness

8. Appendix I -Requirements Analysis and Ranking

8.1.1. Quality Function Deployment - Introduction

Quality Function Deployment (developed by Y. Akao in Japan in 1966) is a systematic approach to design and develop a product (of any kind, including platforms/pieces of software) based on a close awareness of customer desires and requirements, coupled with the integration of functional groups (of a project team or company).

Quoting Y. Akao, QFD "*is a method for developing a design quality aimed at satisfying the consumer and then translating the consumer's demand into design targets and major quality assurance points to be used throughout the production phase. ... [QFD] is a way to assure the* ***design quality while the product is still in the design stage****.*"

In essence, quality is the barycentre of the methodology and the ultimate goal is to translate (often) subjective quality criteria into objective ones that can be quantified and measured, and which can then be used to design and develop the product. Basically, QFD allows determining how and where priorities are to be assigned in product development apriori of implementation. The intent is to employ objective procedures in increasing detail throughout the development of the product or project.

The three main goals in implementing QFD are:

1. Prioritize spoken and unspoken customer desires and needs;
2. Translate these needs into technical characteristics, requirements and specifications;
3. Develop and deliver a quality product (or service) by focusing on customer/user satisfaction whilst optimizing usage of internal costs, resources and teams (e.g. in a project, or in a Company).

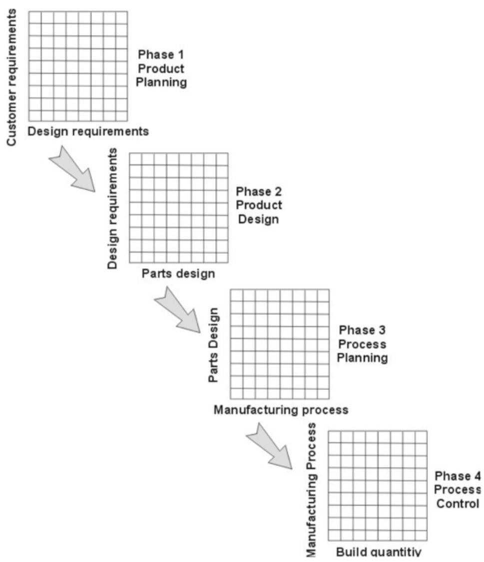

Figure 8-1: Flows of the QFD Analysis

QFD uses some principles from Concurrent Engineering in that cross-functional teams are involved in all phases of product development. As depicted in Figure 8-1**,** each of the four phases in a QFD process uses a matrix to translate customer requirements from initial planning stages through production control. These phases are:

- ***Phase 1: Product Planning***. It is also called The House of Quality. Main goals are: documenting customers' requirements, competitive opportunities, product measurements, competing product measures, and the technical ability of the organization to meet each customer requirement.
- ***Phase 2: Product Design***. Product concepts are created during this phase and part specifications are documented. Parts that are determined to be most important to meeting customer needs are then deployed into process planning, or Phase 3.
- ***Phase 3: Process Planning***. During this phase, development processes are planned and flowcharted and the parameters (or target values) are documented.
- ***Phase 4: Process Control.*** It concerns the control of the development processes, the maintenance of schedules, and skills training for developers. Also, in this phase decisions are made as to which process poses the most risk and controls are put in place to prevent failures.

For each key actor in DOLFIN project (i.e. DC Operator initiator, DC Operator responder, Aggregator, End Customer, Utility Operator, Authority) we evaluated correlations between the aggregated requirements and business objectives/ DOLFIN project expected outcomes. In other words we evaluated how each aggregated requirements is contributing to solve/enable each business objective / expected project outcome as seen from each key actor.

The business objectives and expected outcomes used in the QFD analysis are as follows:

- ***Business Objectives:*** BO1.1- Reduce Server Costs, BO1.2- Reduce Infrastructure Costs, BO1.3- Reduce Power Costs, BO1.4- Reduce Other Costs, BO1.5- Reduce all operational costs, BO2- Increase the performance of offered services, BO3- Increase customer satisfaction, BO4- Improve the overall efficiency of the energy system, BO5.1-Shift processing power, BO5.2- Migrate processing power, BO6- Increase the availability /reliability of DC federation.
- ***DOLFIN Project expected outcomes:*** O1- Model, monitor & measure energy consumption; O2 - Energy optimisation & triggers for seamless and autonomic movement of VMs between servers of the same DC or across a group of energy-conscious synergetic DCs respecting/renegotiated SLAs; O3 - Optimize DC energy consumption, by dynamically changing the percentage of active/standby servers and load per server; O4 - Optimise the cumulative energy consumption in a group of DCs (policy-based VMs allocation); O5 - Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs, following the electricity demand-response approach);O6 - Smart grid energy stabilisation, by dynamically changing the energy consumption/production requirements of DCs

The use cases identified in deliverable D2.1 and used in the QFD analysis are as follows:

- **Use Cases:** UC1-Energy efficient workload redistribution, UC2-Multi tariffs from the Utility companies, UC3-Multi tariffs from the Utility companies, UC4-SLA Renegotiation with end customers, UC5-Optimize benefits/incentives from national/European authorities, UC6-Smart City
- **Actors involved in DC operations in each use case***:* DC Initiator, DC Responder, DC Aggregator, End Customer, Utility Operator, Authority

Correlation scores and importance factors for each problem were provided by an expert team from DOLFIN partners. These correlation evaluations were done according to each key actor Quality context for each business objectives and project outcomes and also for all actors (i.e. *Overall value*: correlation values related to the average of each previous value).

8.1.2. Requirements Ranking Results of the QFD Analysis

The results of this QFD analysis are in terms of ranked aggregated requirements as they contribute most/least to the realisation of business objectives and of project outcomes. These results will be used as input to the detailed design and implementation of the DOLFIN system components.

8.1.2.1. DCO Initiator (as Actor) View Point On Business Objectives

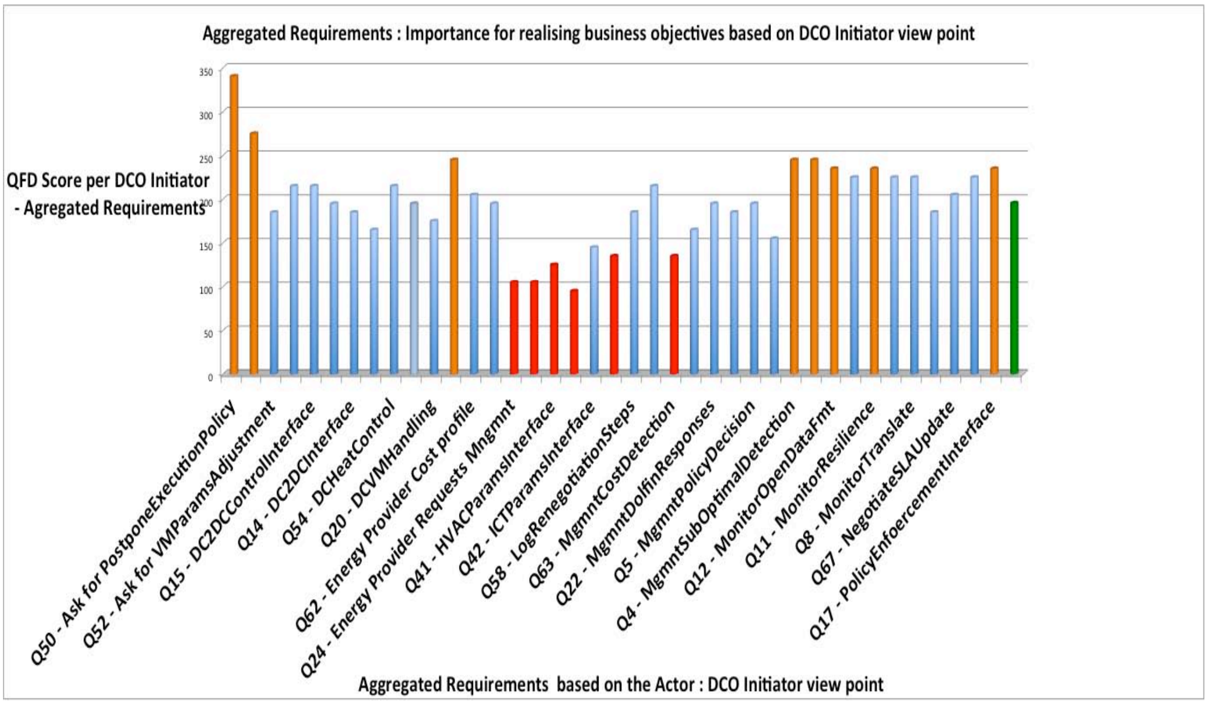

Figure 8-2: DCO Initiator View Point – Realising Business Objectives

- 8 (of 39) requirements marked in orange colour are emerging as very important in realising business objectives as far as DCO Initiator is concerned.
- 6 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising business objectives as far as DCO Initiator is concerned.
- The rest of 25 (of 39) requirements marked in blue colour are emerging as having average importance in realising business objectives as far as DCO Initiator is concerned.

8.1.2.2. DCO Initiator (as Actor) View Point On Expected Project Outcomes

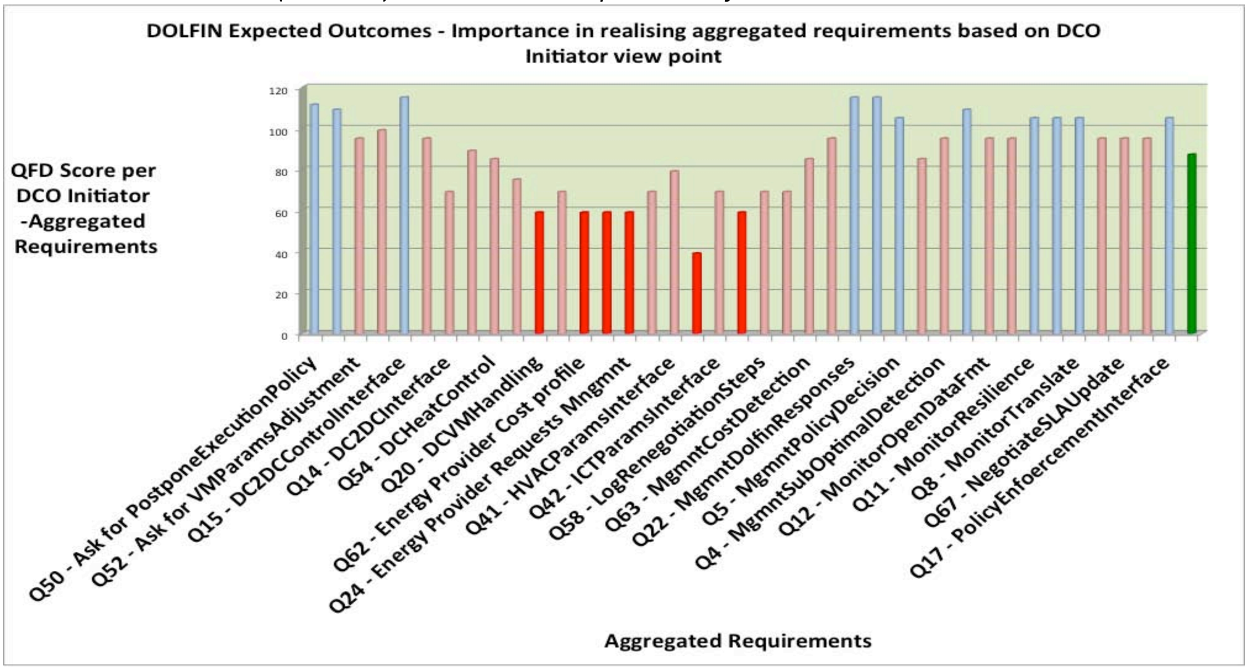

Figure 8-3: DCO Initiator View Point – Realising Project Outcomes

- 10 (of 39) requirements marked in blue colour are emerging as very important in realising project outcomes as far as DCO Initiator is concerned.
- 6 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising project out comes as far as DCO Initiator is concerned.
- The rest of 23 (of 39) requirements are emerging as having average importance in realising project outcomes as far as DCO Initiator is concerned.

8.1.2.3. DCO Responder (as Actor) View Point On Business Objectives

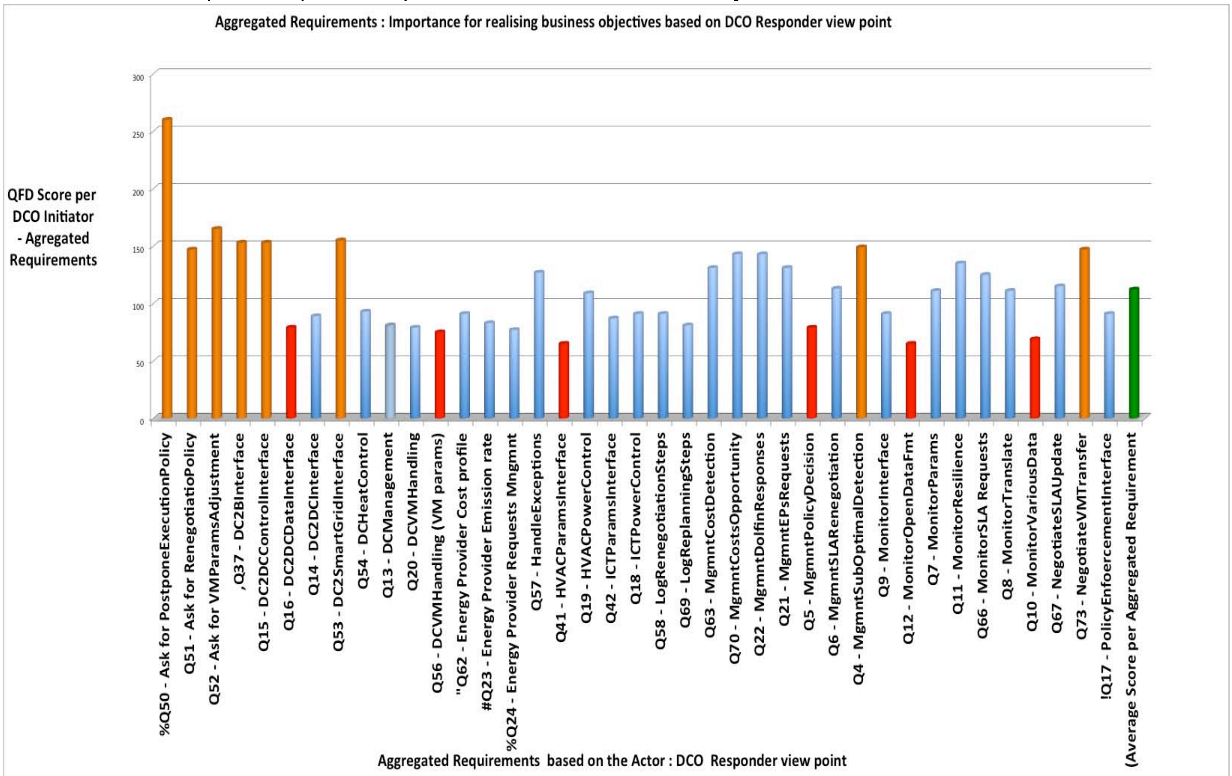

Figure 8-4: DCO Initiator View Point – Realising Project outcomes

- 8 (of 39) requirements marked in orange colour are emerging as very important in realising business objectives as far as DCO Responder is concerned.
- 6 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising business objectives as far as DCO Responder is concerned.
- The rest of 25 (of 39) requirements are emerging as having average importance in realising business objectives as far as DCO Responder is concerned.

8.1.2.4. DCO Responder (as Actor) View Point On Expected Project Outcomes

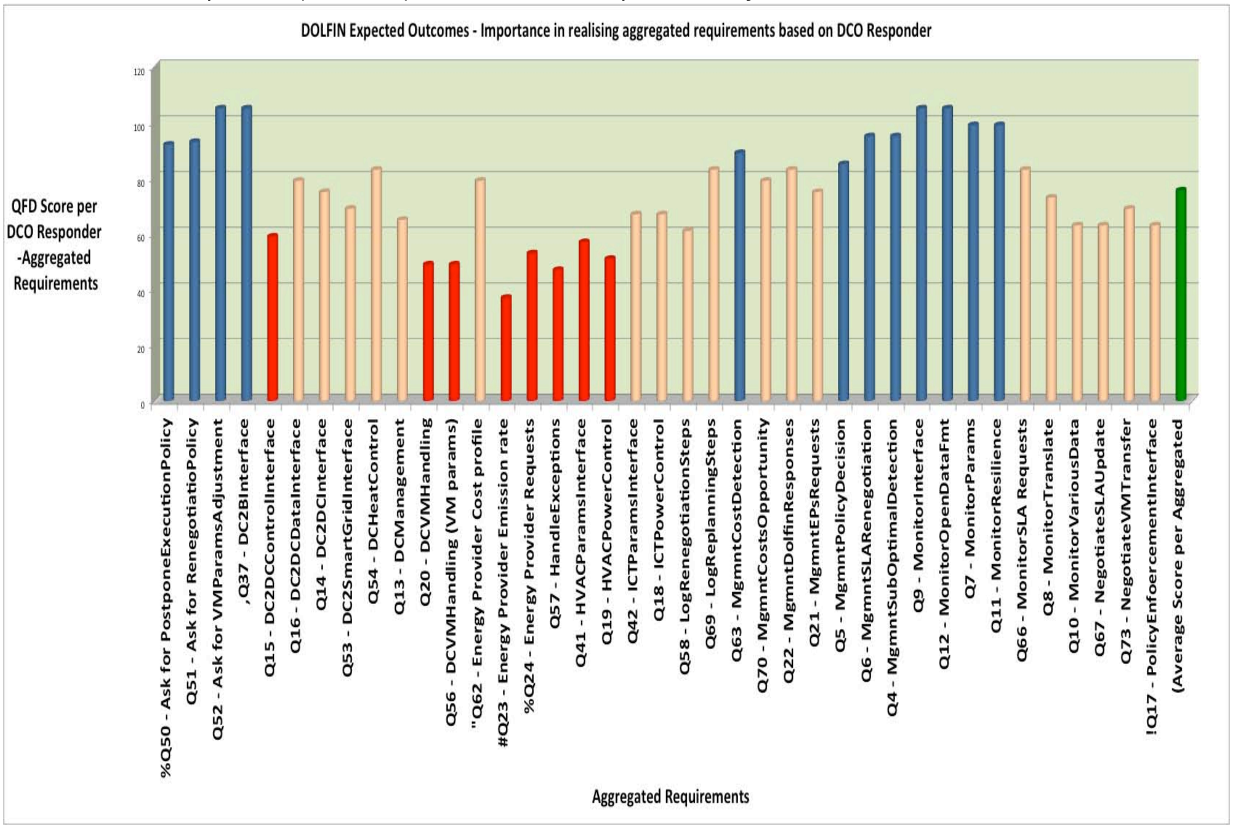

Figure 8-5: DCO Responder View Point – Realising Project Outcome

- 12 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as DCO Responder is concerned.
- 8 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as DCO Responder is concerned.
- The rest 19 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as DCO Responder is concerned.

8.1.2.5. DCO Aggregator (as Actor) View Point On Business Objectives

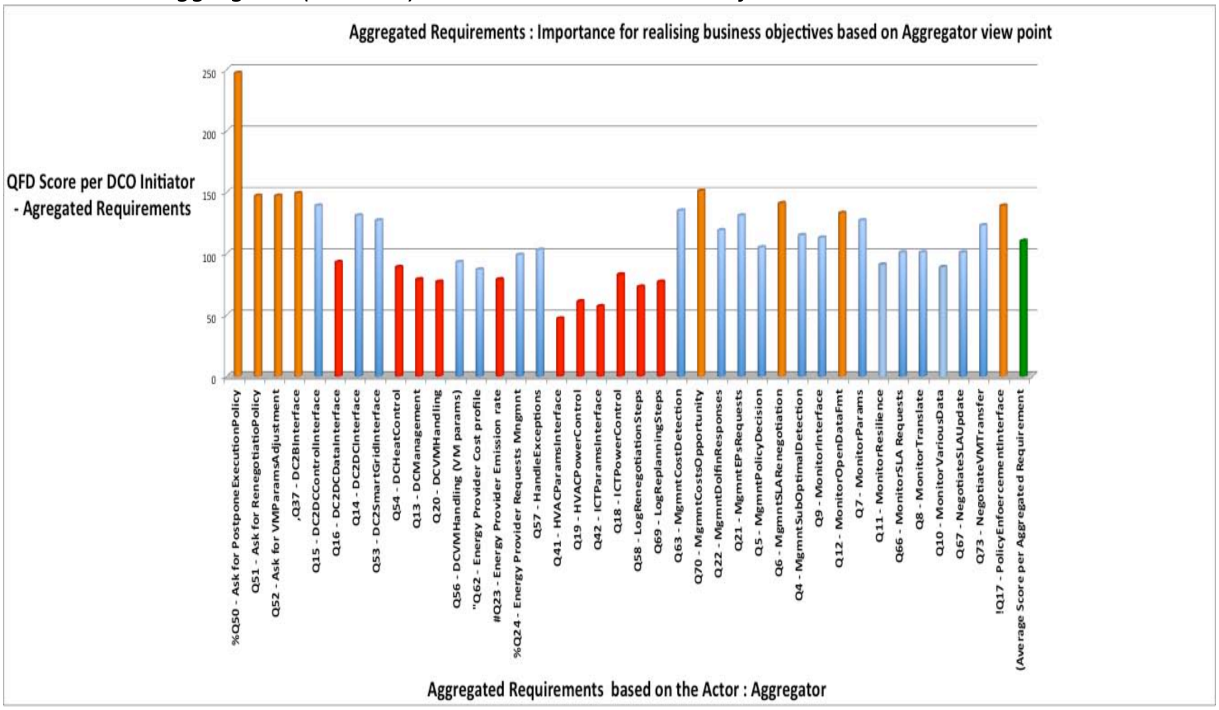

Figure 8-6: DCO Aggregator View Point – Realising Business Objectives

- 8 (of 39) requirements marked in orange colour are emerging as very important in realising business objectives as far as DCO Aggregator is concerned.
- 11 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising business objectives as far as DCO Aggregator is concerned.
- The rest of 20 (of 39) requirements are emerging as having average importance in realising business objectives as far as DCO Aggregator is concerned.

8.1.2.6. DCO Aggregator (as Actor) View Point On Expected Project Outcomes

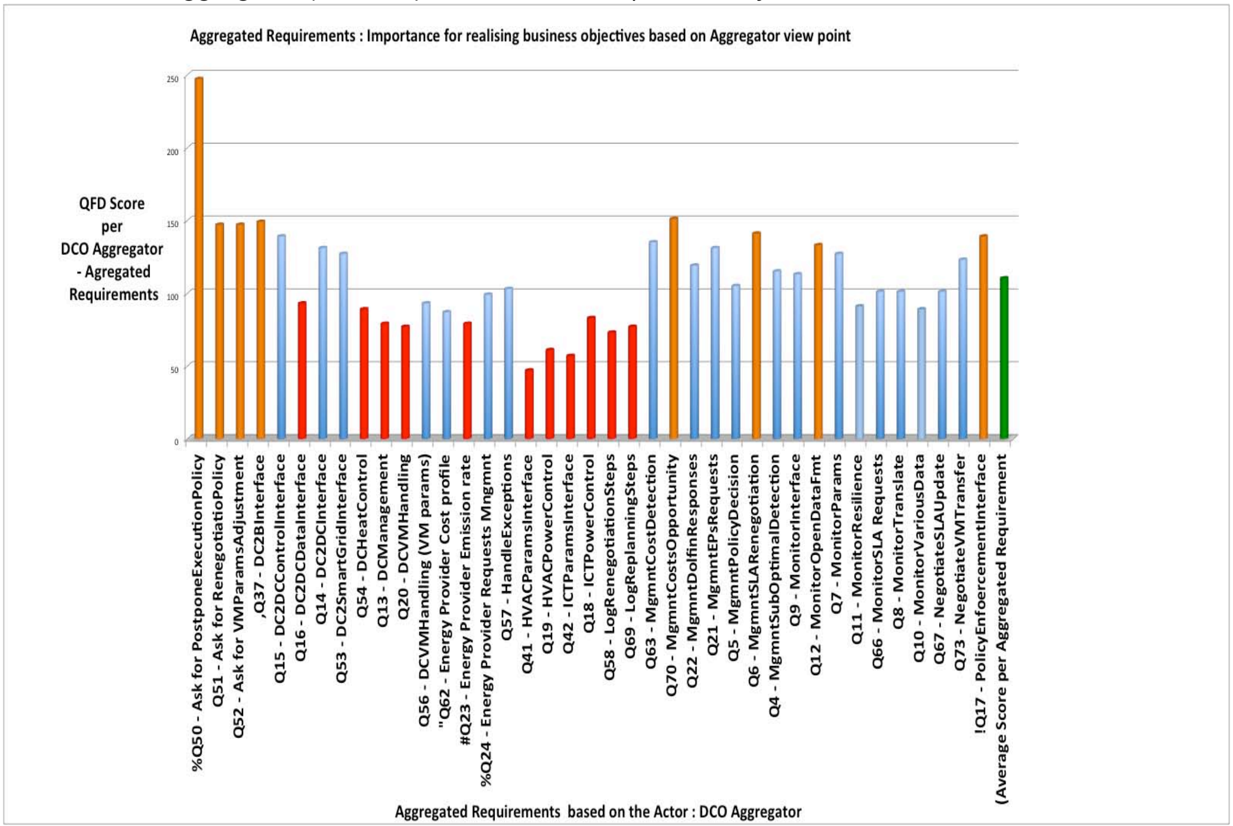

Figure 8-7: DCO Aggregator View Point – Realising Project Outcome

- 12 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as DCO Aggregator is concerned.
- 8 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as DCO Aggregator is concerned.
- The rest 19 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as DCO Aggregator is concerned.

8.1.2.7. End Customer (as Actor) View Point On Business Objectives

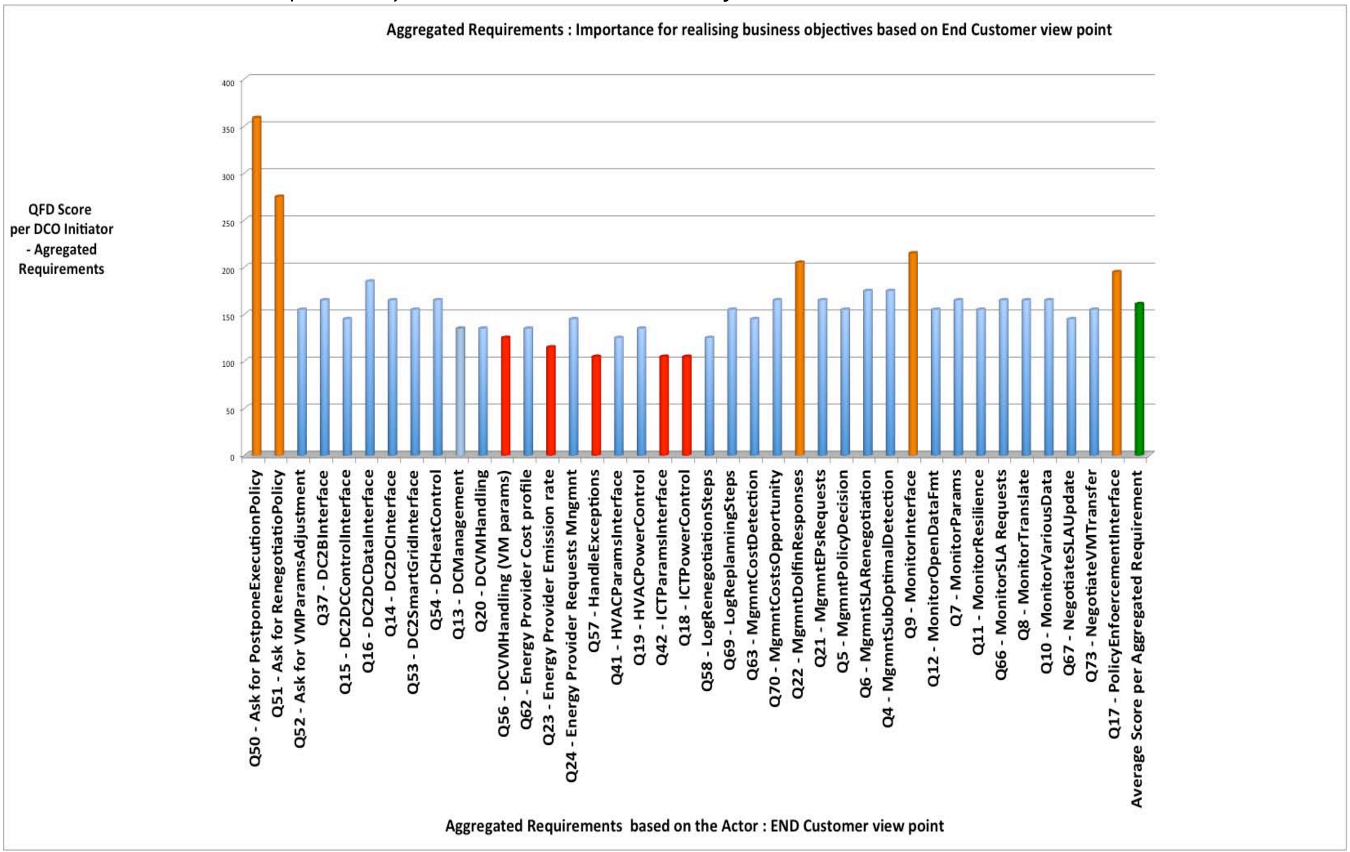

Figure 8-8: End Customer View Point – Realising Business Objectives

- 5 (of 39) requirements marked in orange colour are emerging as very important in realising business objectives as far as End Customer is concerned.
- 5 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising business objectives as far as End Customer is concerned.
- The rest of 29 (of 39) requirements are emerging as having average importance in realising business objectives as far as End Customer is concerned.

8.1.2.8. End Customer (as Actor) View Point On Expected Project Outcomes

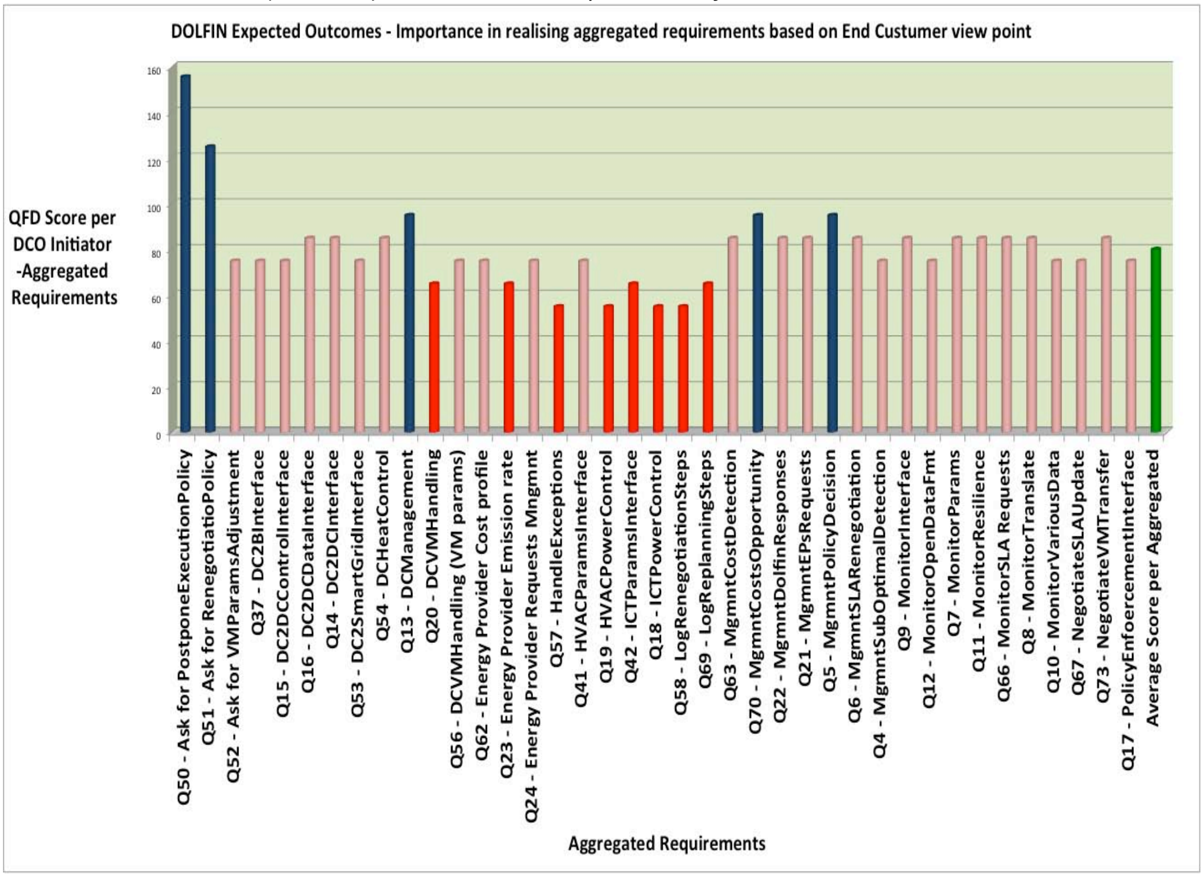

Figure 8-9: End Customer View Point – Realising Project Outcome

- 5 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as End Customer is concerned.
- 8 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as End Customer is concerned.
- The rest 26 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as End Customer is concerned.

8.1.2.9. Utility Operator (as Actor) View Point On Business Objectives

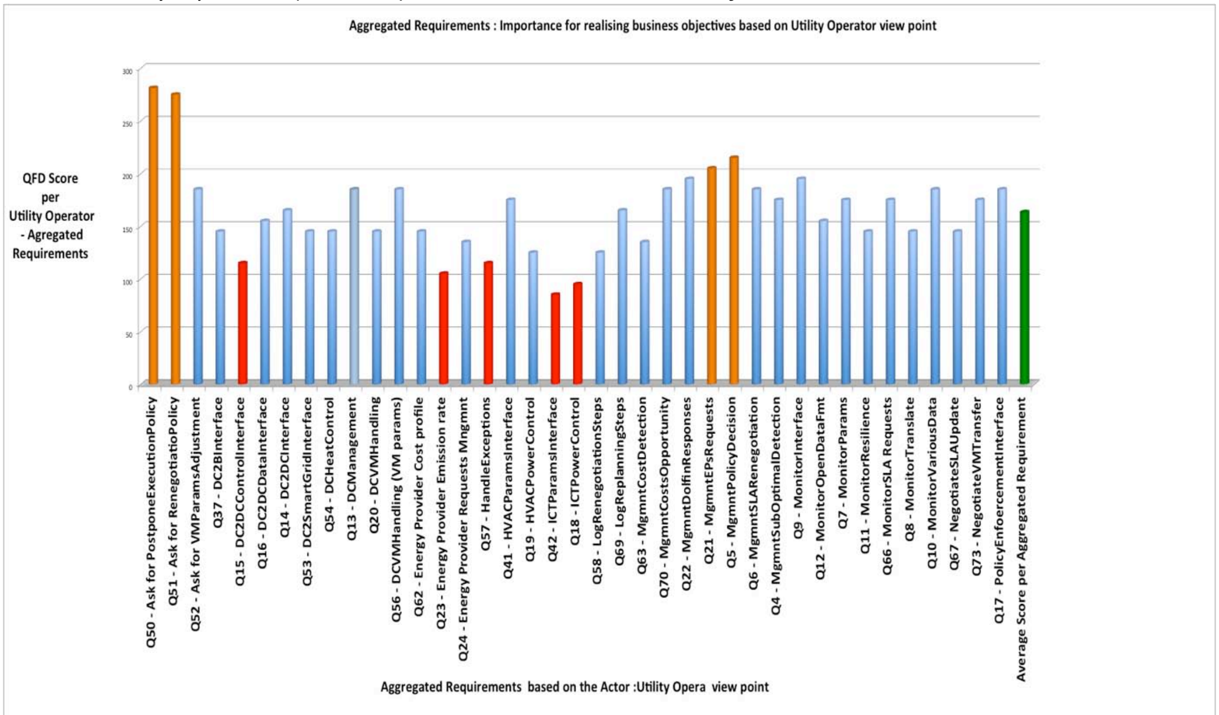

Figure 8-10: Utility Operator View Point – Realising Business Objectives

- 4 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as Utility Operator is concerned.
- 5 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as Utility Operator is concerned.
- The rest 30 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as Utility Operator is concerned.

8.1.2.10. Utility Operator (as Actor) View Point On Expected Project Outcomes

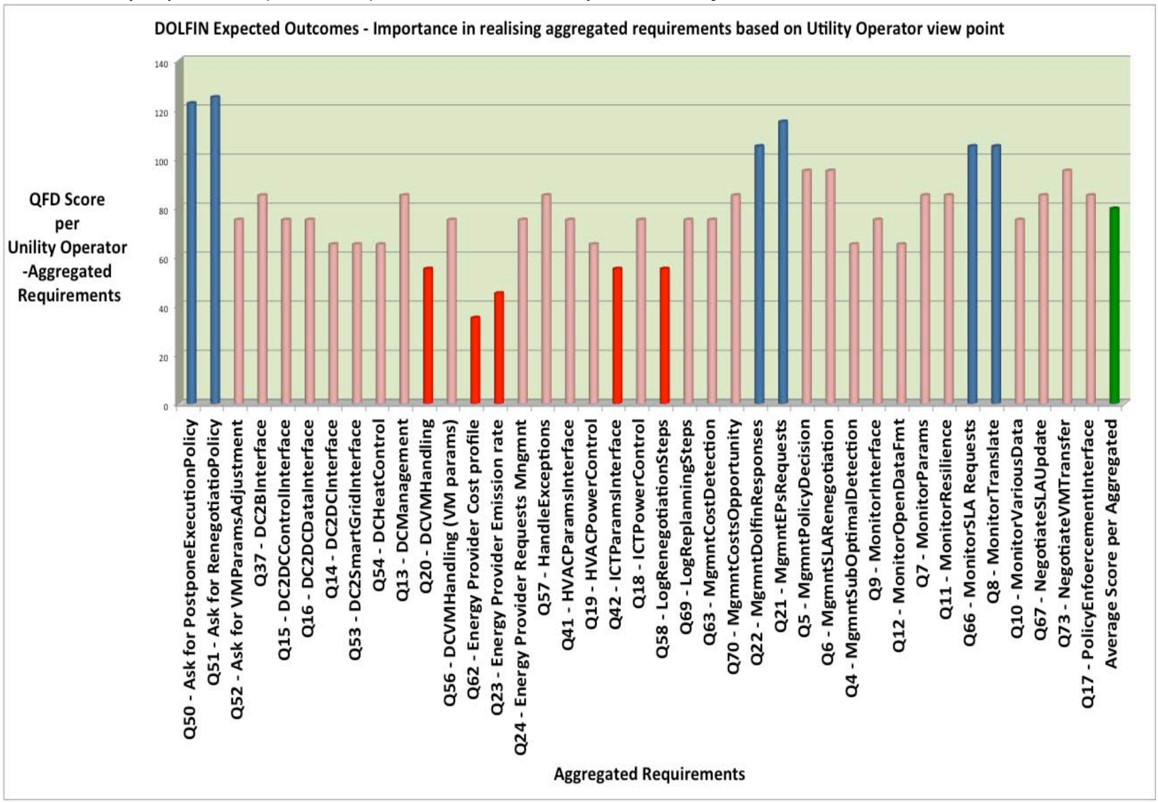

Figure 8-11: Utility Operator View Point – Realising Project Outcome

- 6 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as Utility Operator is concerned.
- 5 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as Utility Operator is concerned.
- The rest 28 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as Utility Operator is concerned.

8.1.2.11. Authority (as Actor) View Point On Business Objectives

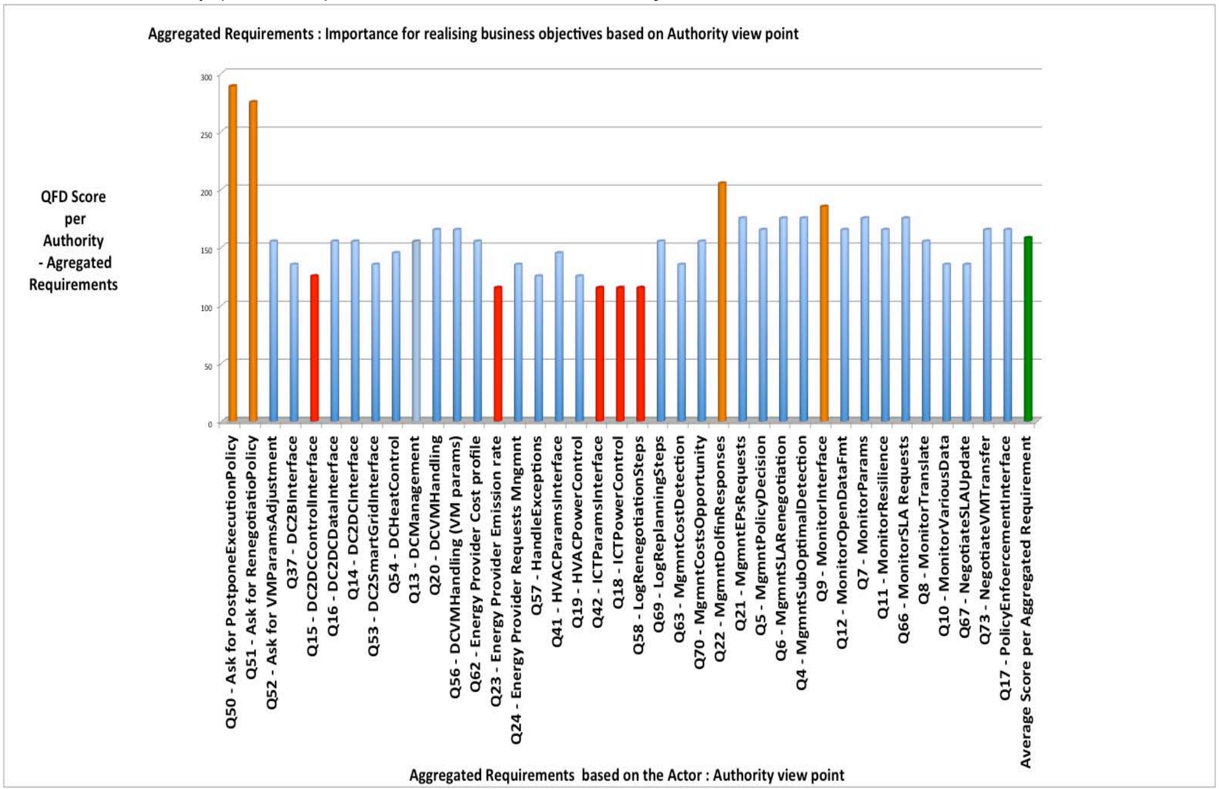

Figure 8-12: Authority View Point – Realising Business Objectives

- 4 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as Authority is concerned.
- 5 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as Authority is concerned.
- The rest 30 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as Utility Operator is concerned.

8.1.2.12. Authority (as Actor) View Point On Expected Project Outcomes

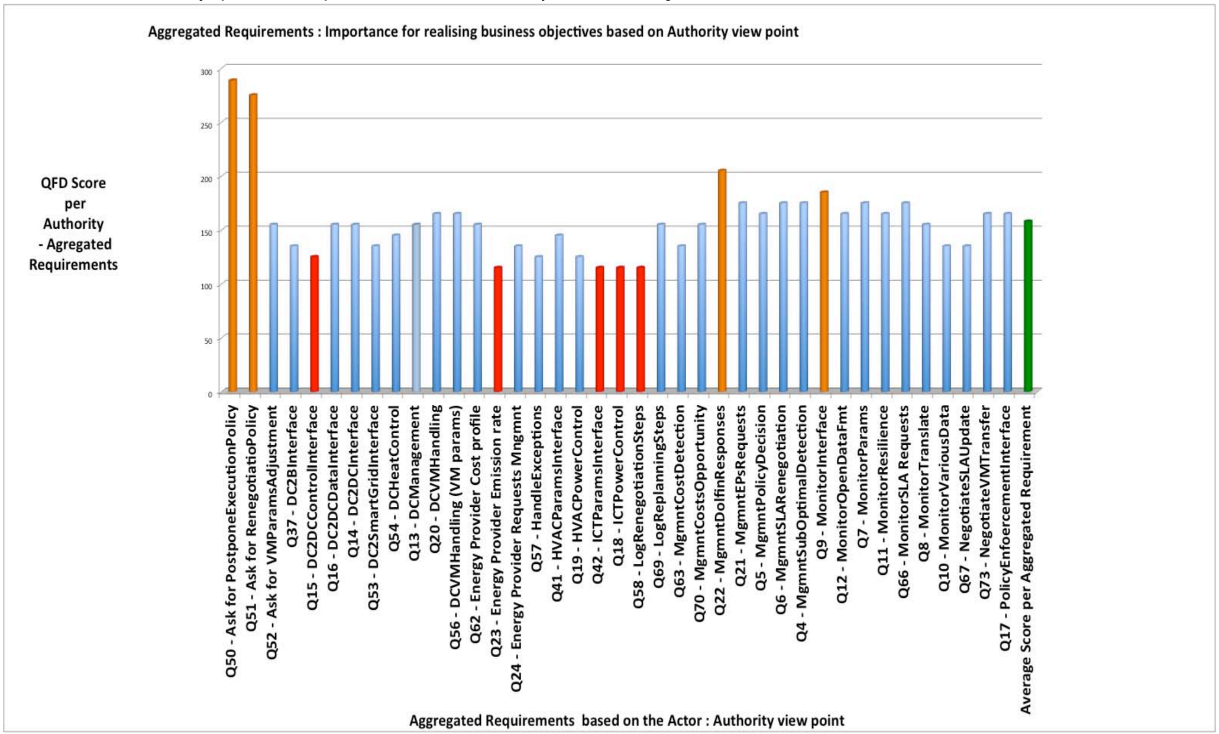

Figure 8-13: Authority View Point – Realising Project Outcome

- 6 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as Authority is concerned.
- 7 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as Authority is concerned.
- The rest 26 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as Authority is concerned.

8.1.2.13. Combined All Actors View Point On Business Objectives

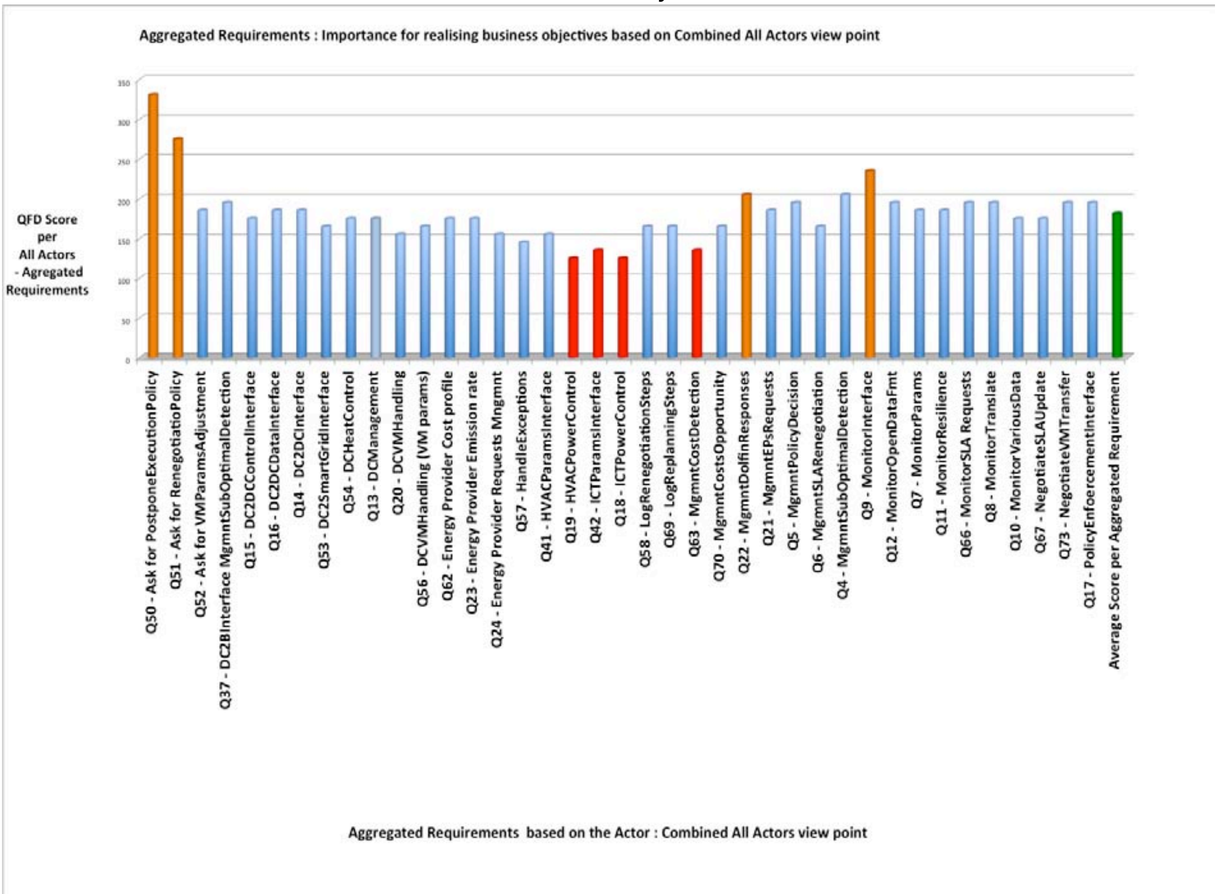

Figure 8-14: Combined All Actors View Point – Realising Business Objectives

- 4 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as Combined All Actors are concerned.
- 4 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as Combined All Actors are concerned.
- The rest 31 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as Combined All Actors are concerned.

8.1.2.14. Combined All Actors View Point On Expected Project Outcomes

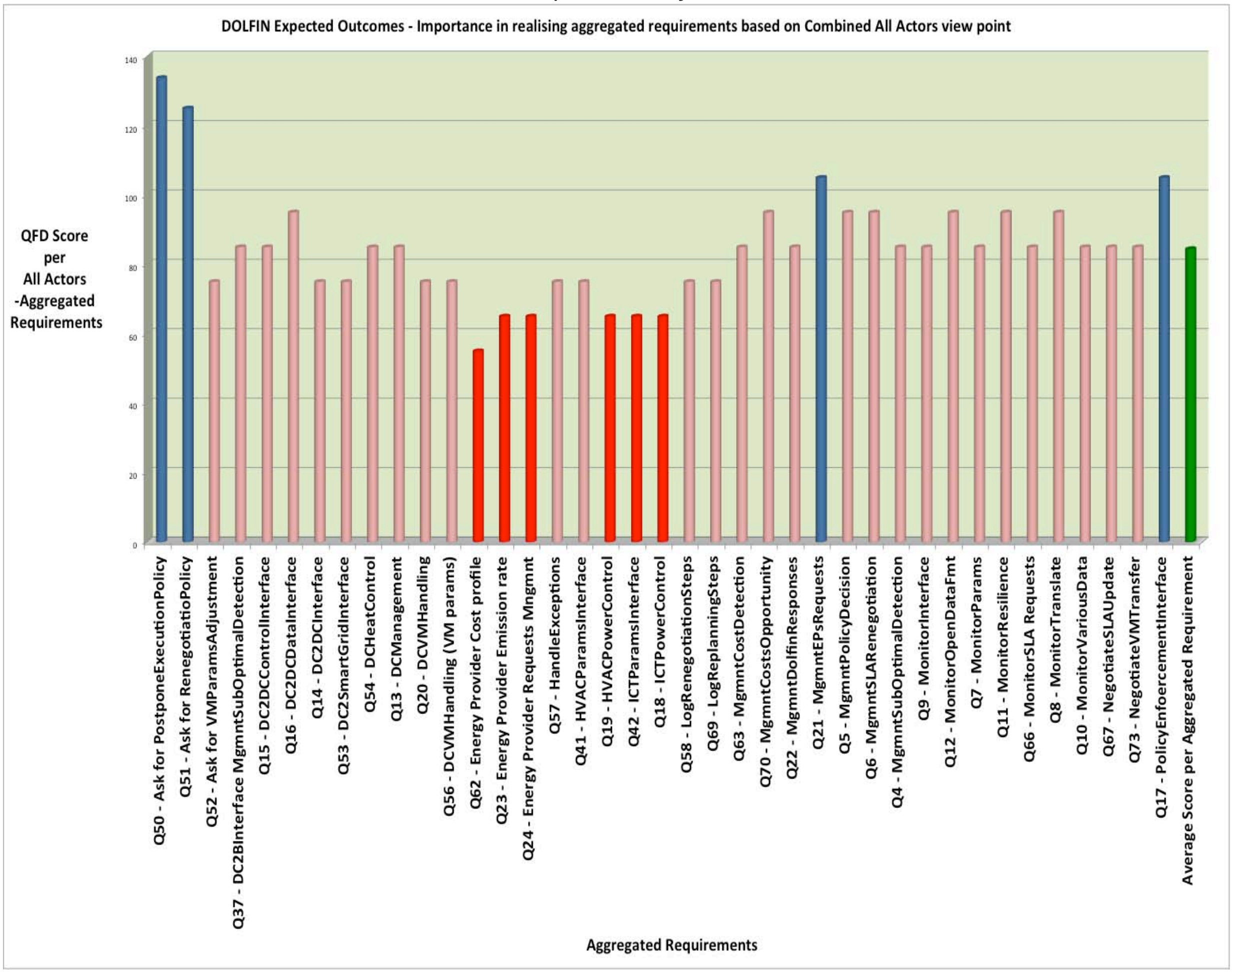

Figure 8-15: Combined All Actors View Point – Realising Project Outcome

- 5 (of 39) requirements marked in blue colour are emerging as very important in realising Project Outcomes as far as Combined All Actors are concerned.
- 6 (of 39) requirements marked in red colour are emerging as having a lower than average importance in realising Project Outcomes as far as Combined All Actors are concerned.
- The rest 28 (of 39) of requirements are emerging as having average importance in realising Project Outcomes as far as Combined All Actors are concerned.

8.1.3. Actors Prioritisation Results using QFD Analysis

Wrapping up the conducted QFD analysis, the rest of this section summarizes the respective results regarding the evaluation of both the business objectives and the expected project outcomes in hand. Specifically, a prioritization of the various Actors per business objective and project outcome is attempted. Finally, for each of the business aspects and project outcome examined, an overall value analysis is presented, indicative of the design quality of the project as a whole.

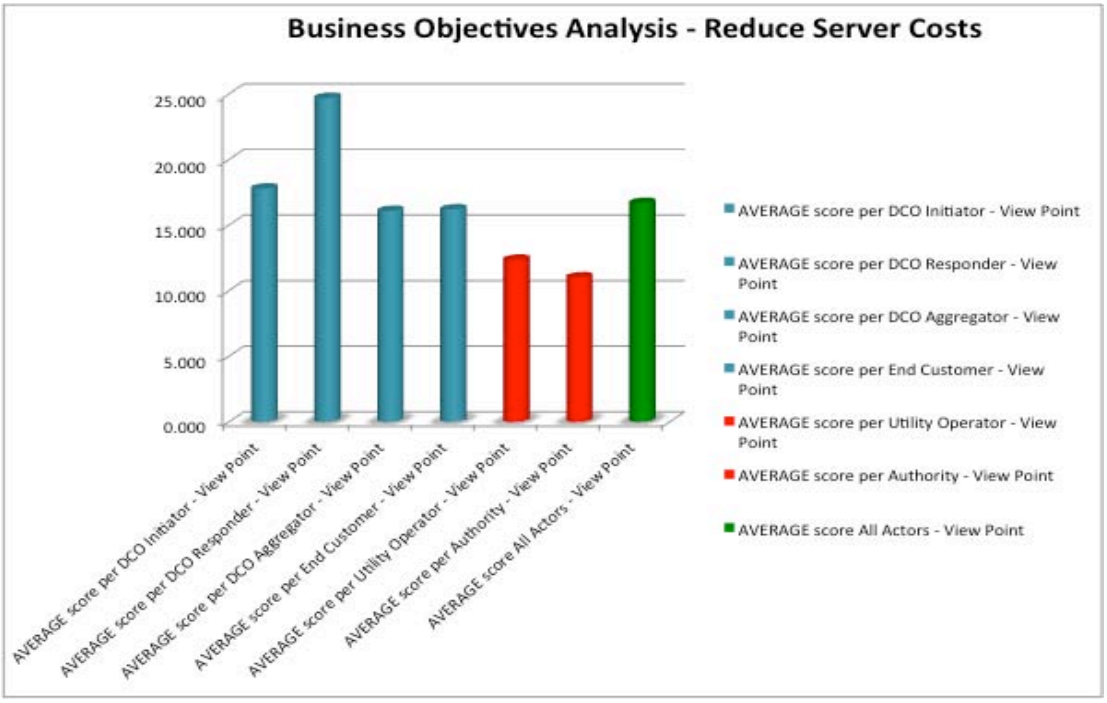

Figure 8-16: Business Objective: Reduced Server Costs

- 1 (of 6) Actor view point is emerging to contribute most to the business objective: Reduce Server Costs
  - DC Responder
- 3 (of 6) Actors view point are following as contributors to the business objective: Reduce Server Costs
  - DC Initiator
  - DC Aggregator
  - End Customer
- 3 (of 6) Actors view points have a lower than average contribution to the business objective: reduce server costs
  - Utility Operator
  - Authority

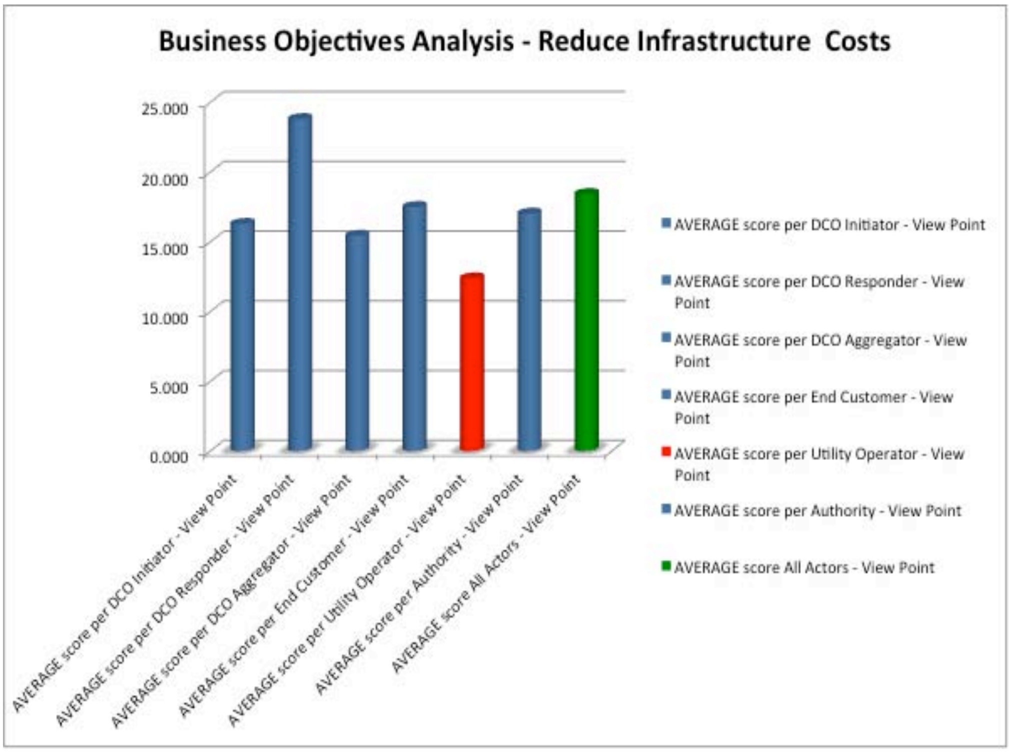

Figure 8-17: Business Objective: Reduce Infrastructure Costs

- 1 (of 6) Actor view point is emerging to contribute most to the business objective: reduce infrastructure costs
    - DC Responder
- 4 (of 6) Actors view point are following as contributors to the business objective: reduce infrastructure costs
    - DC Initiator
    - DC Aggregator
    - End Customer
    - Authority
- 1 (of 6) Actor view point have a lower than average contribution to the business objective: reduce infrastructure costs
    - Utility Operator

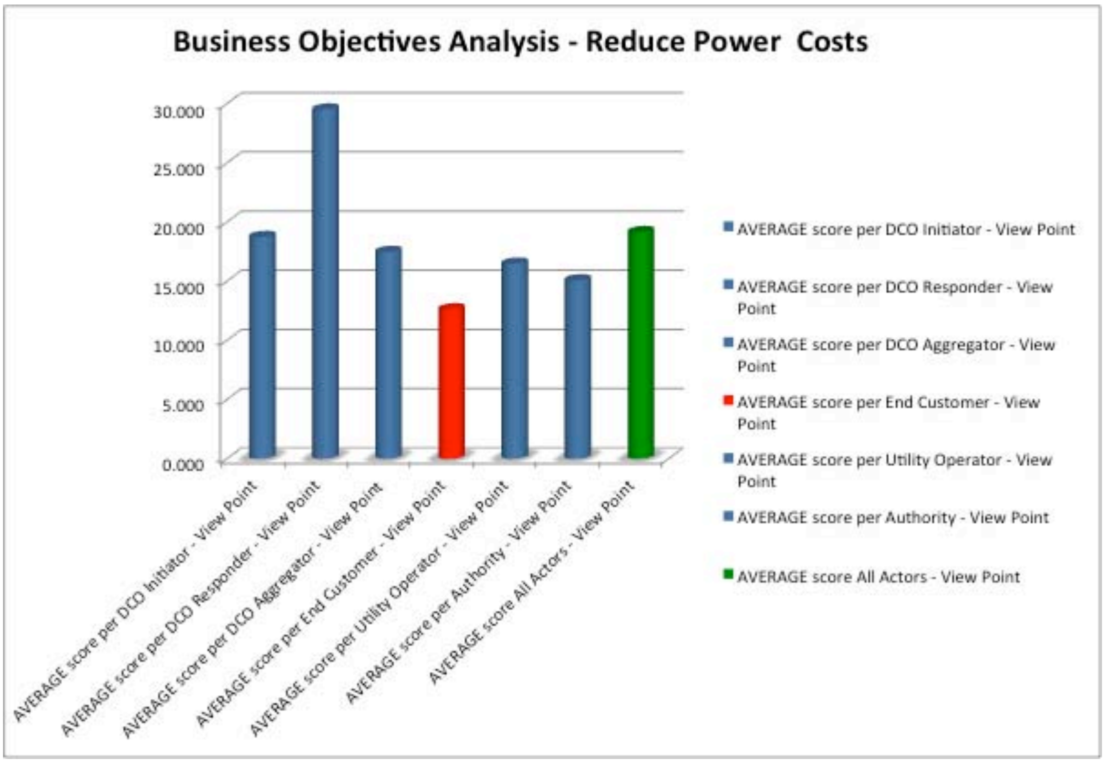

Figure 8-18: Bussiness Objectives: Reduce Power Costs

- 1 (of 6) Actor view point is emerging to contribute most to the business objective: reduce power costs
  - DC Responder
- 4 (of 6) Actors view point are following as contributors to the business objective: reduce power costs
  - DC Initiator
  - DC Aggregator
  - End Customer
  - Authority
- 1 (of 6) Actor view point have a lower than average contribution to the business objective: reduce power costs
  - Utility Operator

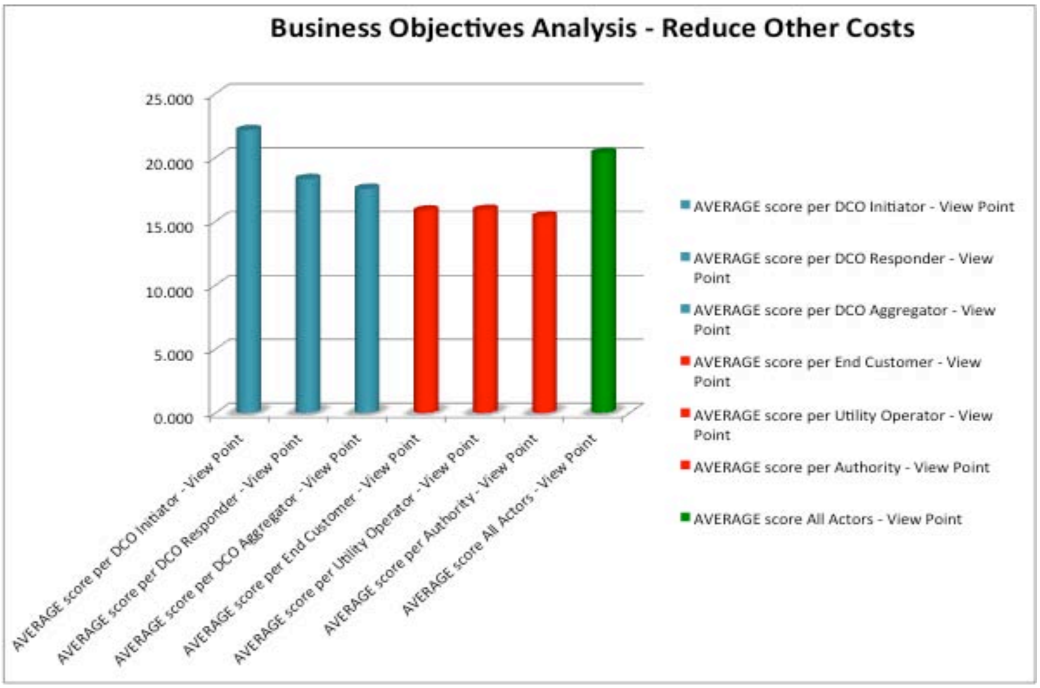

Figure 8-19: Business Objective: Reduce Other Costs

- 1 (of 6) Actor view point is emerging to contribute most to the business objective: reduce other costs
  - DC Initiator
- 2 (of 6) Actors view point are following as contributors to the business objective: reduce other costs
  - DC Responder
  - DC Aggregator
- 3 (of 6) Actors view points have a lower than average contribution to the business objective: reduce other costs
  - End Customer
  - Authority
  - Utility Operator

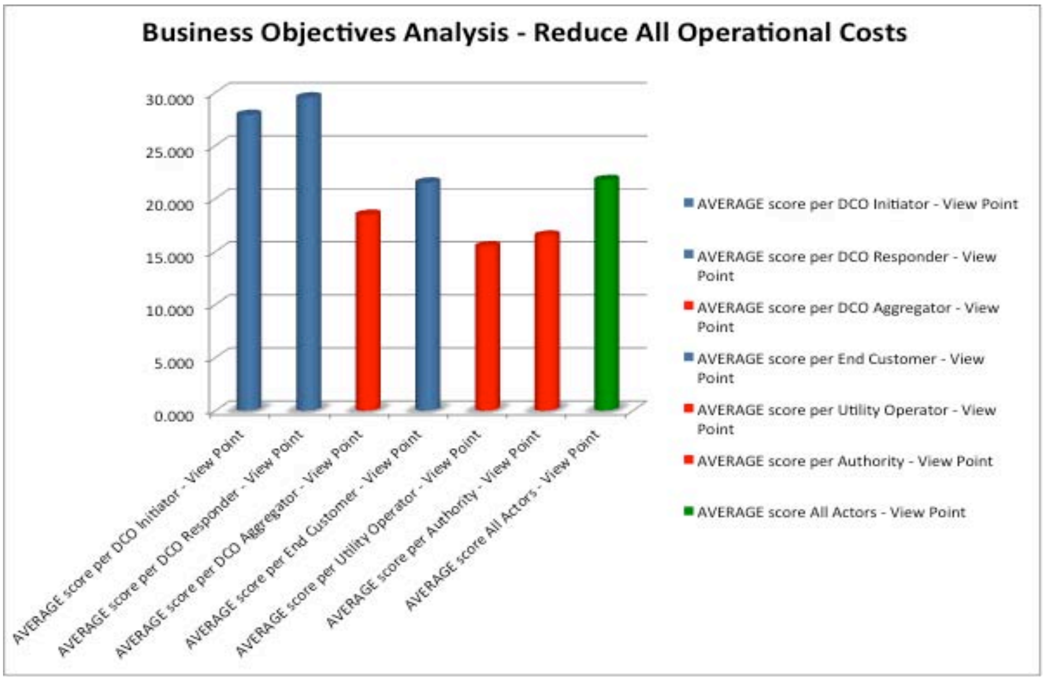

Figure 8-20: Business Objective: Reduce All Operational Costs

- 2 (of 6) Actors view point are emerging to contribute most to the business objective: reduce other costs
  - DC Initiator
  - DC Responder
- 1 (of 6) Actor view point is following as contributor to the business objective: reduce other costs
  - End Customer
- 3 (of 6) Actors view points have a lower than average contribution to the business objective: reduce other costs
  - DC Aggregator
  - Authority
  - Utility Operator

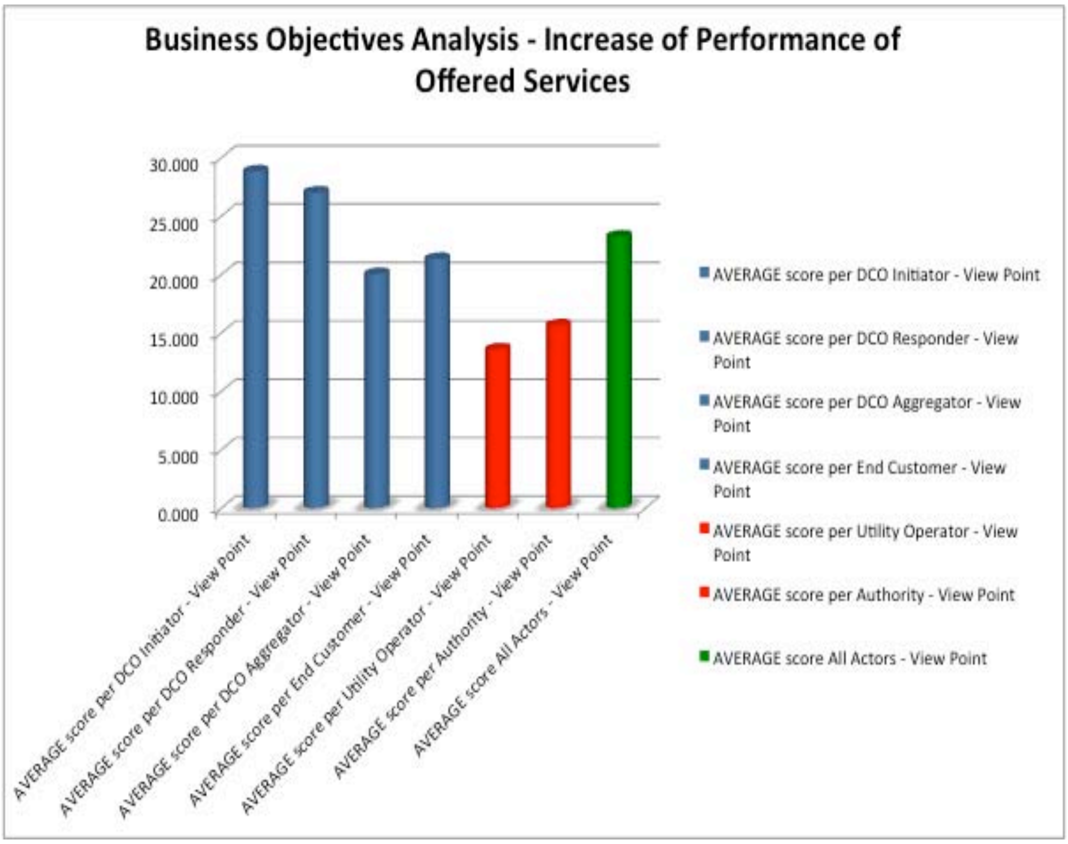

Figure 8-21 Bussiness Objective: Increase of Performance of Offered Service

- 2 (of 6) Actors view point are emerging to contribute most to the business objective: Increase Performance of Offered Services
  - DC Initiator
  - DC Responder
- 2 (of 6) Actors view point are following as contributors to the business objective: Increase Performance of Offered Services
  - DC Aggregator
  - End Customer
- 2 (of 6) Actors view points have a lower than average contribution to the business objective: Increase Performance of Offered Services
  - Authority
  - Utility Operator

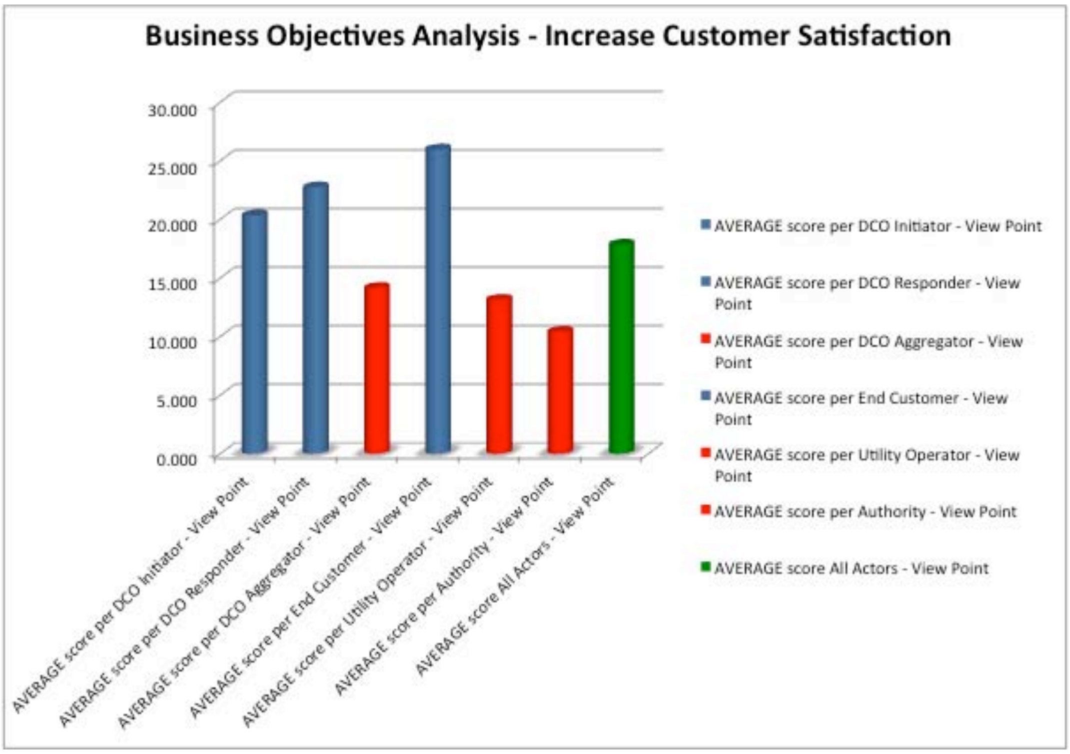

Figure 8-22: Business Objective: Increase Customer Satisfaction

- 3 (of 6) Actors view point are emerging to contribute most to the business objective: Increase Customer Satisfaction
  - DC Initiator
  - DC Responder
  - End Customer
- 3 (of 6) Actors view points have a lower than average contribution to the business objective: Increase Customer Satisfaction
  - Authority
  - Utility Operator

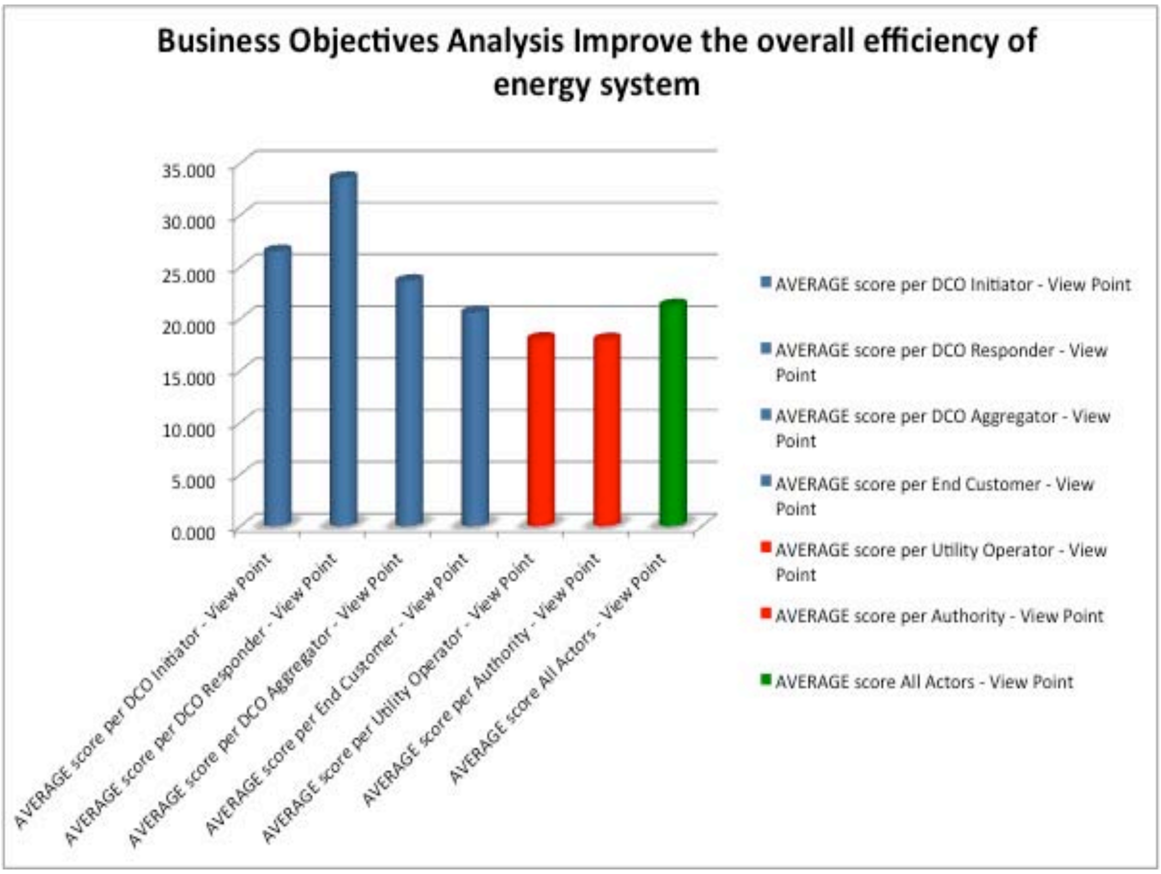

Figure 8-23: Businress Objective: Improve overall efficiency of energy systems

- 3 (of 6) Actors view point are emerging to contribute most to the business objective: Improve Overall Efficiency of Energy Systems
    - DC Initiator
    - DC Responder
    - DC Aggregator
- 1 (of 6) Actor view point is following as contributors to the business objective: Improve Overall Efficiency of Energy Systems
    - End Customer
- 2 (of 6) Actors view points have a lower than average contribution to the business objective: Improve Overall Efficiency of Energy Systems
    - Authority
    - Utility Operator

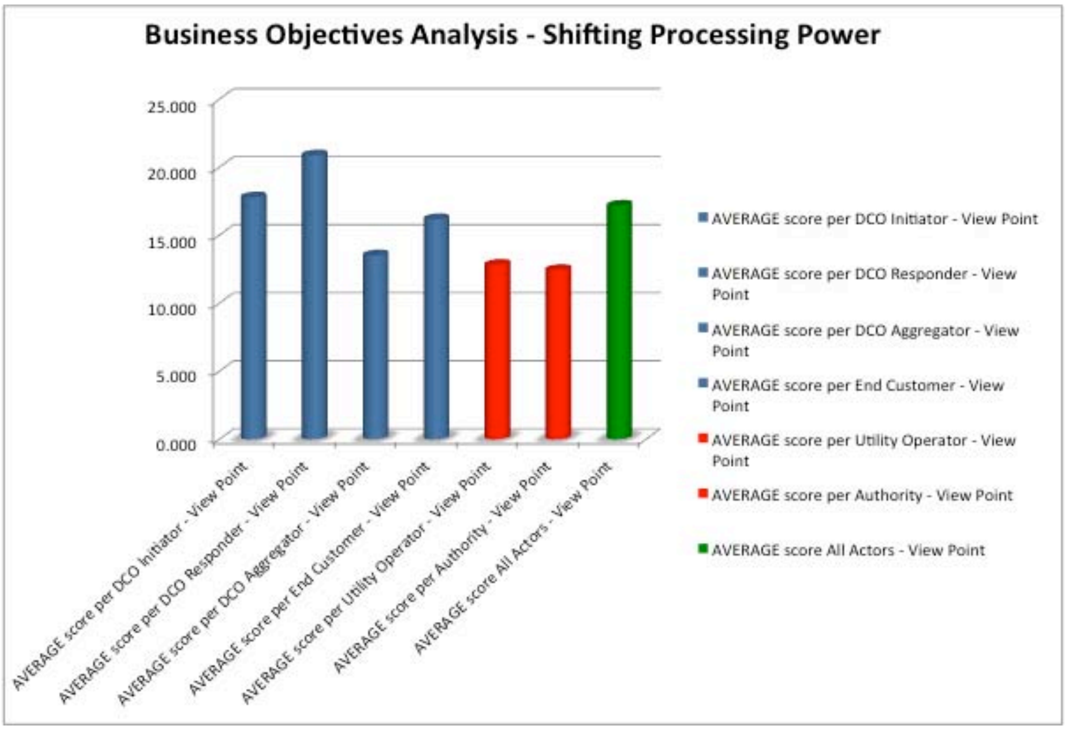

Figure 8-24: Business Objective: Shifting Processing Power

- 2 (of 6) Actors view point are emerging to contribute most to the business objective: Shifting Processing Power
  - DC Initiator
  - DC Responder
- 2 (of 6) Actors view points are following as contributors to the business objective: Shifting Processing Power
  - Aggregator
  - End Customer
- 2 (of 6) Actors view points have a lower than average contribution to the business objective: Shifting Processing Power
  - Utility Operator
  - Authority

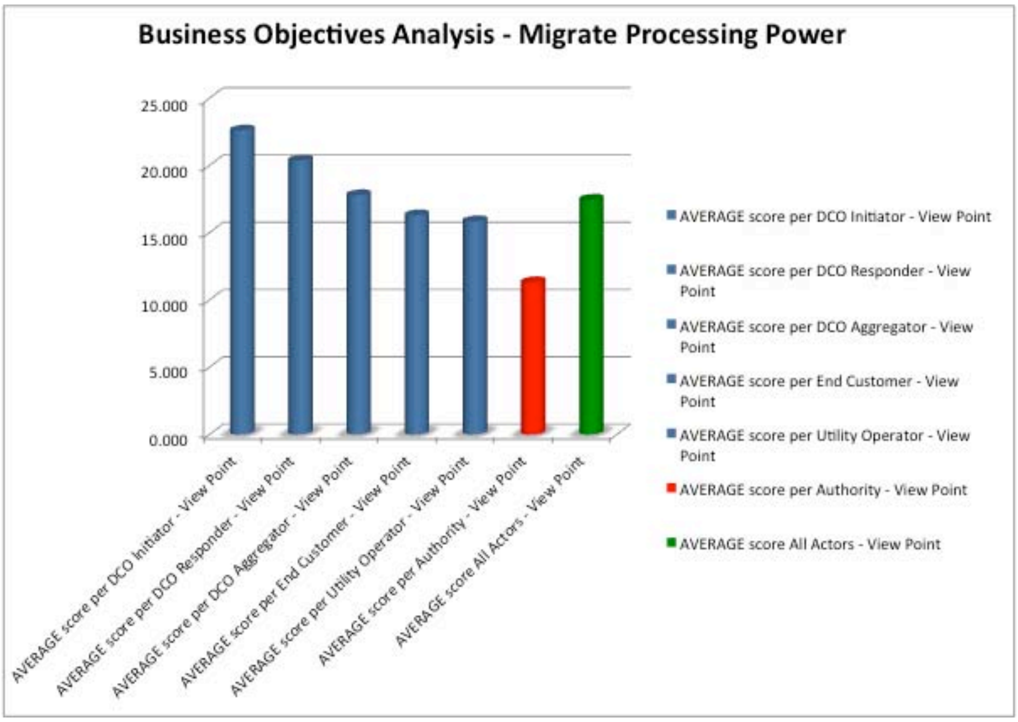

Figure 8-25: Business Objective: Migrate Processing Power

- 2 (of 6) Actors view point are emerging to contribute most to the business objective: Migrate Processing Power
    - DC Initiator
    - DC Responder
- 3 (of 6) Actors view points are following as contributors to the business objective: Migrate Processing Power
    - Aggregator
    - End Customer
    - Utility Operator
- 1 (of 6) Actor view point has a lower than average contribution to the business objective: Migrate Processing Power
    - Authority

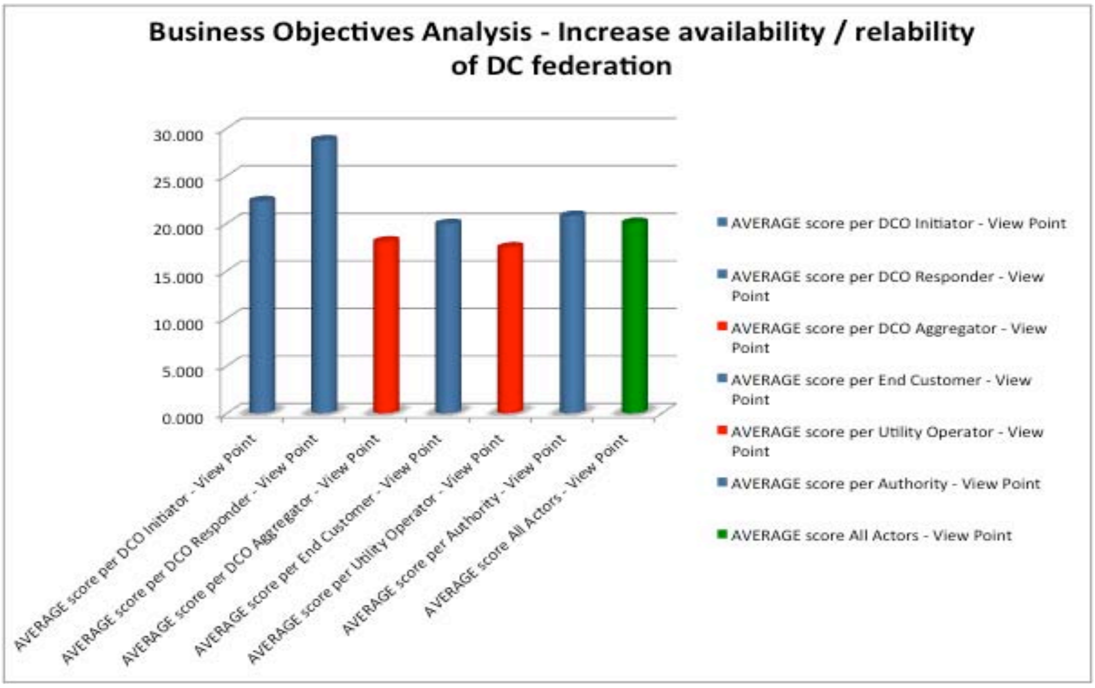

Figure 8-26: Business Objective: Increase availability / relability of DC federation

- 2 (of 6) Actors view point are emerging to contribute most to the business objective: Increase availability / relability of DC federation
    - DC Initiator
    - DC Responder
- 2 (of 6) Actors view points are following as contributors to the business objective: : Increase availability / relability of DC federation
    - End Customer
    - Authority
- 2 (of 6) Actor view point has a lower than average contribution to the business objective: : Increase availability / relability of DC federation
    - Aggregator
    - Utility Operator

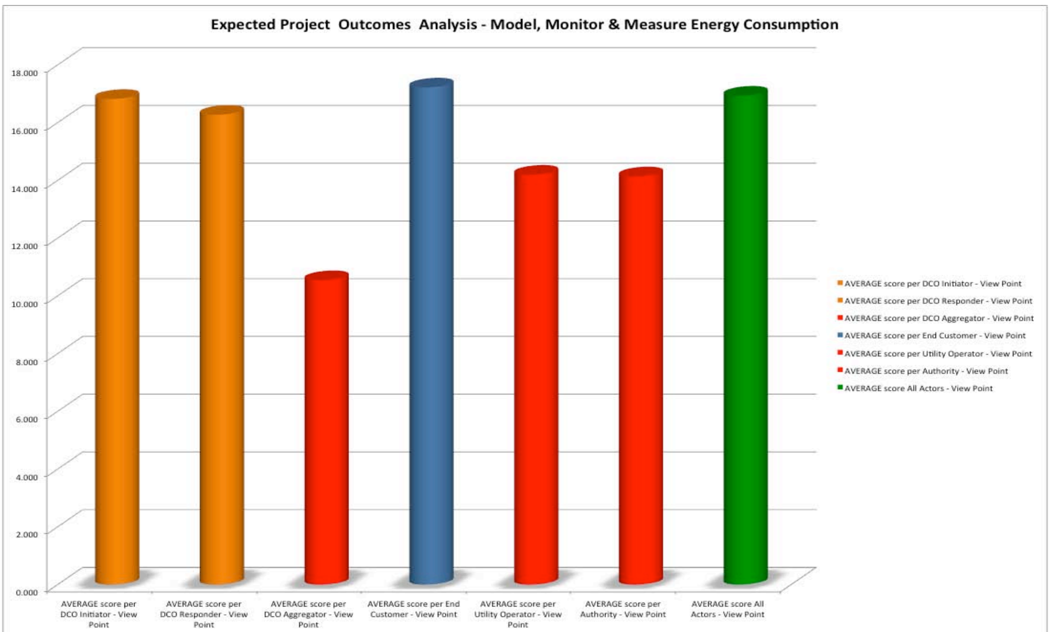

Figure 8-27: Expected Project Outcome: Model, Monitor & Measure Energy Consumption

- 1 (of 6) Actor view point are emerging to contribute most to the Project Outcome: Model, Monitor & Measure Energy Consumption
    - End Customer
- 2 (of 6) Actors view points are following as contributors to the Project Outcome: Model, Monitor & Measure Energy Consumption
    - DC Initiator
    - DC Responder
- 3 (of 6) Actor view point has a lower than average contribution to the Project Outcome: Model, Monitor & Measure Energy Consumption
    - Aggregator
    - Utility Operator
    - Authority

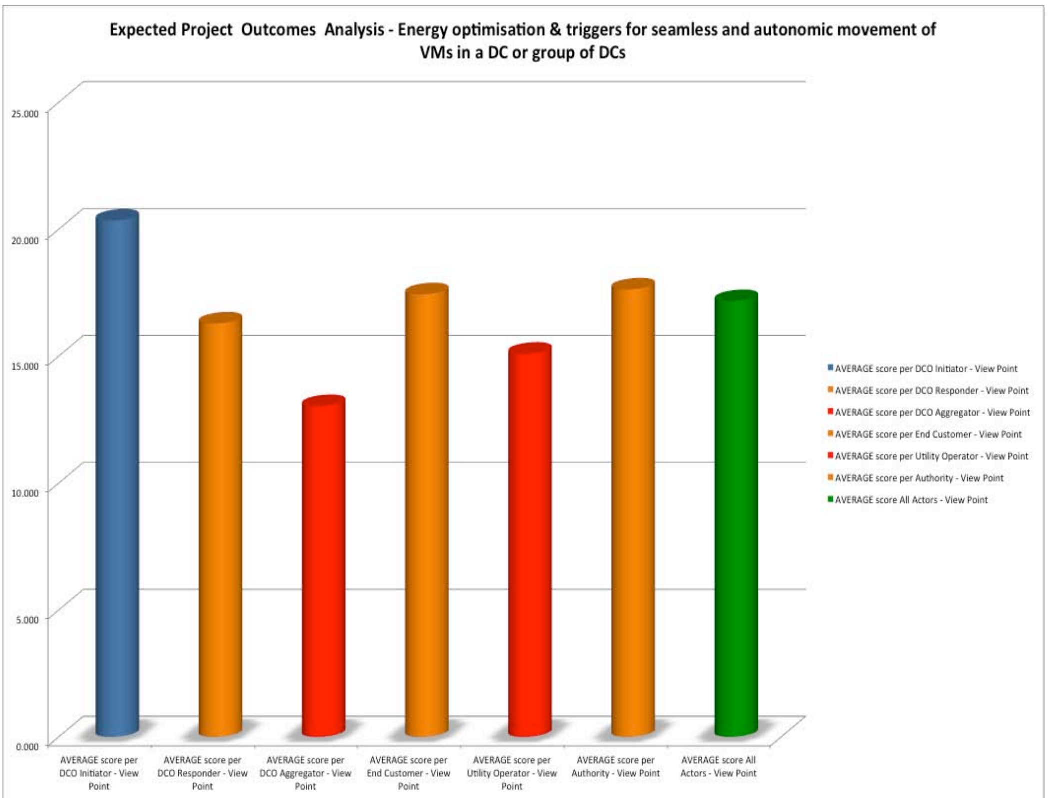

Figure 8-28: Expected Project Outcome: Energy consumption & triggers for seamless and autonomic movement of VMs in a DC or group of DCs

- 1 (of 6) Actor view point are emerging to contribute most to the Project Outcome: Energy consumption & triggers for seamless and autonomic movement of VMs in a DC or group of DCs
    - DC Initiator
- 3 (of 6) Actors view points are following as contributors to the Project Outcome: Energy consumption & triggers for seamless and autonomic movement of VMs in a DC or group of DCs
    - DC Responder
    - End Customer
    - Authority
- 3 (of 6) Actor view point has a lower than average contribution to the Project Outcome: Energy consumption & triggers for seamless and autonomic movement of VMs in a DC or group of DCs
    - Aggregator
    - Utility Operator

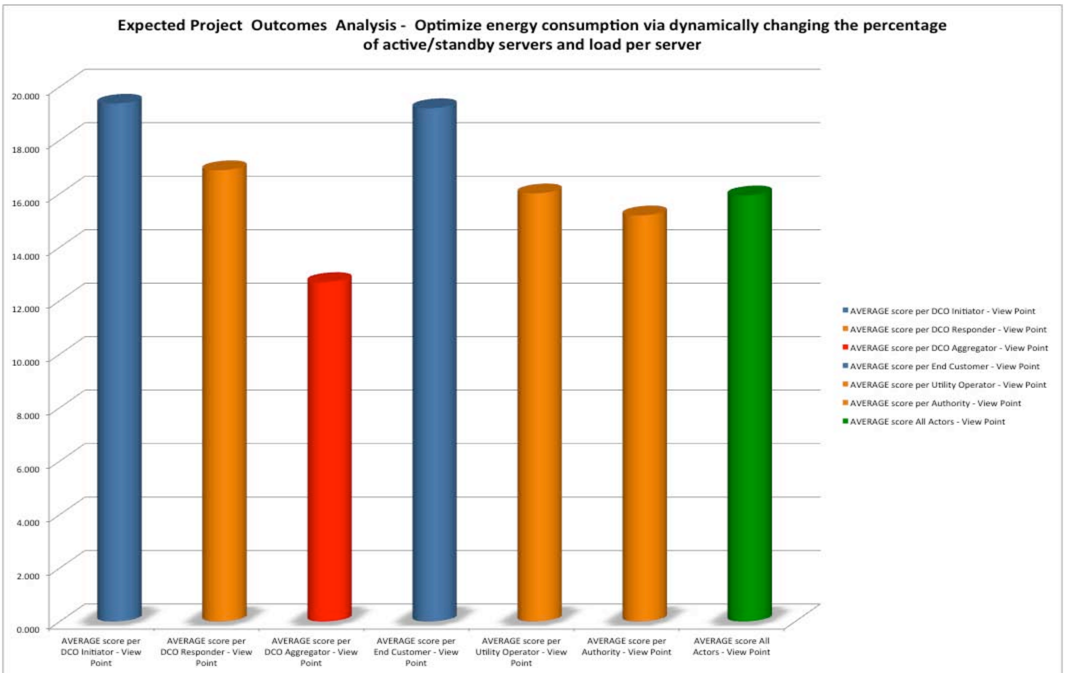

Figure 8-29: Expected Project Outcome: Optimise energy consumption via dynamically changing the percentage of active/standby server and load per server

- 2 (of 6) Actors view points are emerging to contribute most to the Project Outcome: Optimise energy consumption via dynamically changing the percentage of active/standby server and load per server
    - DC Initiator
    - End Customer
- 3 (of 6) Actors view points are following as contributors to the Project Outcome: Optimise energy consumption via dynamically changing the percentage of active/standby server and load per server
    - DC Responder
    - Utility Operator
    - Authority
- 1 (of 6) Actor view point has a lower than average contribution to the Project Outcome: Optimise energy consumption via dynamically changing the percentage of active/standby server and load per server
    - Aggregator

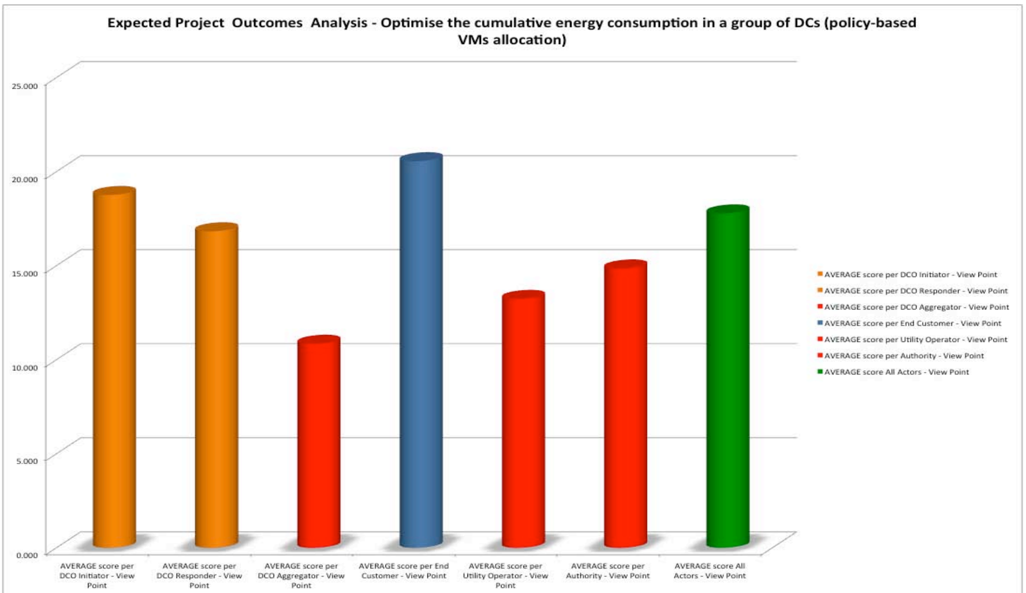

Figure 8-30: Expected Project Outcome: Optimise the cumulative consumption in a group of DCs (policy-based VMs allocation)

- 1 (of 6) Actor view point is emerging to contribute most to the Project Outcome: Optimise the cumulative consumption in a group of DCs (policy-based VMs allocation)
    - End Customer
- 2 (of 6) Actors view points are following as contributors to the Project Outcome: Optimise the cumulative consumption in a group of DCs (policy-based VMs allocation)
    - DC Initiator
    - DC Responder
- 3 (of 6) Actor view point has a lower than average contribution to the Project Outcome: Optimise the cumulative consumption in a group of DCs (policy-based VMs allocation)
    - Aggregator
    - Utility Operator
    - Authority

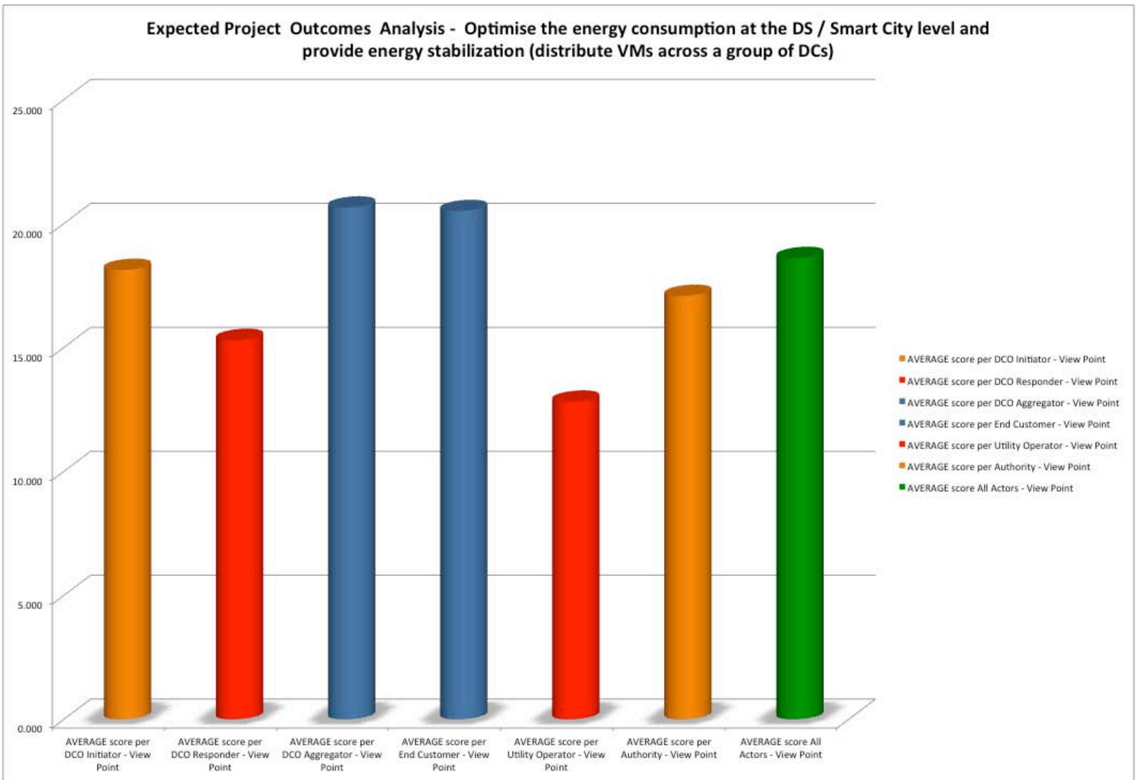

Figure 8-31: Expected Project Outcome: Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs)

- 1 (of 6) Actor view point is emerging to contribute most to the Project Outcome: Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs)
  - Aggregator
  - End Customer
- 2 (of 6) Actors view points are following as contributors to the Project Outcome: Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs)
  - DC Initiator
  - Authority
- 2 (of 6) Actor view point has a lower than average contribution to the Project Outcome: Optimise the energy consumption at the DS / Smart City level and provide energy stabilization (distribute VMs across a group of DCs)
  - DC Responder
  - Utility Operator

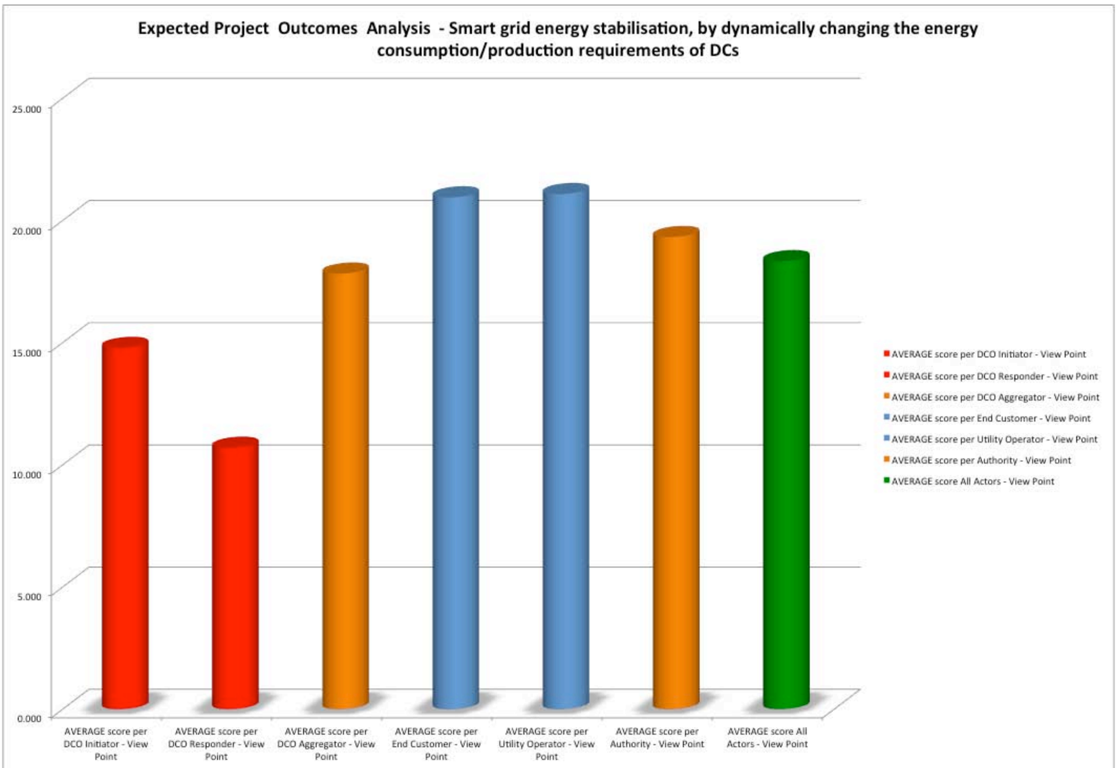

Figure 8-32: Expected Project Outcome: Smart grid energy stabilization (dynamically changing the energy consumption / production requirements of DCs)

- 2 (of 6) Actors view points are emerging to contribute most to the Project Outcome: Smart grid energy stabilization (dynamically changing the energy consumption / production requirements of DCs)
  - End Customer
  - Utility Operator
- 2 (of 6) Actors view points are following as contributors to the Project Outcome: Smart grid energy stabilization (dynamically changing the energy consumption / production requirements of DCs)
  - Aggregator
  - Authority
- 2 (of 6) Actor view point has a lower than average contribution to the Project Outcome: Smart grid energy stabilization (dynamically changing the energy consumption / production requirements of DCs)
  - DC Initiator
  - DC Responder

9. Appendix II – Requirements List

Definitions of Requirements

A requirement can be defined as a statement that identifies a capability or function that is needed by a system in order to satisfy its customer's needs.

In the context of DOLFIN, a functional requirement is meant to describe a function or a feature of a DC system, or its components, capable of solving a certain energy optimisation problem or replying to a certain Energy consumption need/request from key actors. It presents a complete description of how a specific system will function as far as energy management is concerned, capturing every aspect of how it should work before it is built, including information handling, computation handling, storage handling and connectivity handling.

Non-functional requirements are meant as implementation attributes and artefacts that a specific DC system must have. Examples of categories of non-functional requirements are:

- Usability: it describes the ease with which a system performing certain functions or features can be adopted and used.
- Reliability: it describes the degree to which a system must work. Specifications for reliability typically refer to stability, availability, accuracy, and maximum acceptable bugs.
- Performance: it describes the degree of performances of the system (according to certain predefined metrics, e.g. convergence time).
- Supportability: it refers to a system's ability to be easily modified or maintained to accommodate usage in typical situations and change scenarios. For instance, how easy should it be to add new blocks and/or subsystems to the support framework.
- Security: it refers to the ability to prevent and/or forbid access to a system by unauthorized parties.
- Safety: It refers to conditions of being protected against different types and the consequences of failure, error harm or any other event, which could be considered non-desirable.
- Resilience: it refers to the ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operation
- Compliance: it refers to the conformance to a rule, such as a specification, policy, standard or regulatory.
- Extensibility: it refers to the ability to extend a system and the level of effort and complexity required to realize an extension. Extensions can be through the addition of new functionality, new characteristics or through modification of existing functionality/characteristics, while minimizing impact to existing system functions.
- Interoperability: it refers to the ability of diverse systems and subsystems to work together (inter-operate).
- Operability: it refers to the ability to keep a system in a safe and reliable functioning condition, according to pre-defined operational requirements.
- Privacy: it refers the ability of system or actor to seclude itself or information about itself and thereby reveal itself selectively.

- Scalability: it refers to the ability of a system to handle growing amounts of work or usage in a graceful manner and its ability to be enlarged to accommodate that growth.

DOLFIN Requirements

For the priority use cases identified in the deliverable D2.1 (e.g. UC 5.2, UC 5.2.1, UC 5.2.2, UC 5.2.3, UC 5.3, UC 5.3.1, UC 5.3.2, UC 5.3.3, UC 5.4, UC 5.4.1, UC 5.4.2, UC 5.4.3, UC 5.5, UC 5.5.1, UC 5.5.2, UC 5.5.3, UC 5.6, UC 5.6.1, UC 5.6.2, UC 5.6.3, UC 5.7, UC 5.7.1, UC 5.7.2, UC 5.7.3, UC 5.8) 478 row functional requirements were identified in terms if interactions between key actors, functional role and links with the relevant components. The row functional requirements were also assign a specific name and identity (i.e. Rx). They were compiled in a requirement database at **http://claydesk2.ee.ucl.ac.uk:8080/DOLFIN/requirement/index.**

A similarity and consolidation exercise was performed on the 478 row functional requirements with the view of generating higher level aggregation of requirements with the following characteristics:

- Consistent and atomic: a requirement addresses only one function or feature. The requirement is atomic, i.e., it does not contain conjunctions.
- Complete: a requirement is fully defined in one place with no missing information.
- Dependable: a requirement does not contradict any other requirements and is fully consistent with all relevant references.
- Current: a requirement has not been made obsolete.
- Feasible: a requirement can be implemented and supported by the enabling technology
- Verifiable: the implementation of a design goal/objective can be determined through one of five possible methods: inspection, demonstration, test, trial or analysis

As such 39 aggregated requirements were identified, each requirement identified with an identity (i.e. Qx), name, description, technical actors involved, roles, use cases links and links with the row functional requirements. Theses aggregated requirements were also compiled in the requirement database at http://claydesk2.ee.ucl.ac.uk:8080/DOLFIN/requirement/index which could by grouped on different dimensions. The aggregated requirements were ranked for importance and impact on the business goals and DOLFIN project expected outcome using the Quality Function Deployment (QFD) analysis in Chapter 1. QFD is a quantitative technique used to rank the multi-dimensional options of an option set.

The Aggregated Requirements are provided next.

Aggregated Requirement	Aggregated Requirement - Description	Technical Actors	Roles	Row Requirements	UseCases
Identifier:	**Q50**				
Name:	**Ask for PostponeExecutionPolicy**				
Raw Entries:	3				
Description:	The Energy Eff. Policy Maker and Actuator should be able issue a request of the Workload and VM Manager to shift computing load (postpone it for later) when such a need arises, e.g. when this load shifting could result in substantial cost reduction due to electricity price or it is considered to be profitable in the case of an end-user SLA renegotiation process, or such a request arrives DCO / Smart City operator.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	DCO (Initiator)	DOLFIN	Interacting	Ask for PostponeExecutionPolicy	UC 5.3.1
					UC 5.5.1
					UC 5.8

R199	DCO (Initiator)	DOLFIN	Interacting	Ask for PostponeExecutionPolicy	UC 5.3.1
R211	DCO (Initiator)	DOLFIN	Interacting	Ask for PostponeExecutionPolicy	UC 5.5.1
R228	DCO (Initiator)	DOLFIN	Interacting	Ask for PostponeExecutionPolicy	UC 5.8

Identifier:	**Q51**				
Name:	**Ask for RenegotiatioPolicy**				
Raw Entries:	1				
Description:	The Energy Eff. Policy Maker and Actuator should be able issue a request of the Renegotiation Policy (e.g. renegotiation of Workload and VM Manager to shift computing load (postpone it for later) when such a need arises, when this load shifting could result in substantial cost reduction due to electricity price or it is considered to be profitable in the case of an end-user SLA renegotiation process, or such a request arrives from DCO / Smart City operator.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DOLFIN	Interacting	Ask for RenegotiatioPolicy	UC 5.2.1
Raw List					

R183	DCO (Initiator)	DOLFIN	Interacting	Ask for RenegotiatioPolicy	UC 5.2.1

Identifier:	**Q52**				
Name:	**Ask for VMParamsAdjustment**				
Raw Entries:	3				
Description:	In case a workload redistribution is necessitated, the Energy Eff. Policy Maker and Actuator should be able to ask from the Workload and VM Manager to adjust the VM characteristics of a certain set of VMs in order to meet the new DC configuration scheme, as generated by the Energy Eff. Policy Maker and Actuator.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	DCO (Initiator)	DOLFIN	Interacting	Ask for VMParamsAdjustment	UC 5.3.1
					UC 5.5.1
					UC 5.8
Raw List					
R200	DCO (Initiator)	DOLFIN	Interacting	Ask for VMParamsAdjustment	UC 5.3.1
R212	DCO (Initiator)	DOLFIN	Interacting	Ask for VMParamsAdjustment	UC 5.5.1

R229	DCO (Initiator)	DOLFIN	Interacting	Ask for VMParamsAdjustment	UC 5.8

Identifier:	**Q37**				
Name:	**DC2BInterface MgmntSubOptimalDetection**				
Raw Entries:	4				
Description:	The management subsystem shall detect DC suboptimal states (i.e. high KPI) by utilizing current DC metrics and defined thresholds				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	4 useCases
	DCO (Initiator)	DOLFIN	Interacting	DC2BInterface MgmntSubOptimalDetection	UC 5.2.1
					UC 5.3.1
					UC 5.5.1
					UC 5.8
Raw List					
R182	DCO (Initiator)	DOLFIN	Interacting	DC2BInterface MgmntSubOptimalDetection	UC 5.2.1
R198	DCO (Initiator)	DOLFIN	Interacting	DC2BInterface MgmntSubOptimalDetection	UC 5.3.1

R210	DCO (Initiator)	DOLFIN	Interacting	DC2BInterface MgmntSubOptimalDetection	UC 5.5.1
R227	DCO (Initiator)	DOLFIN	Interacting	DC2BInterface MgmntSubOptimalDetection	UC 5.8

Identifier:	**Q15**				
Name:	**DC2DCControlInterface**				
Raw Entries:	9				
Description:	The DC2DC subsystem shall implements CONTROL interfaces to coordinate the VM migration between DCs				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	5 useCases
	DCO (Initiator)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.4.2
	DCO (Responder)				UC 5.4.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
Raw List					
R19	DCO (Initiator)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.6.2
R29	DCO (Initiator)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.6.3

R46	DCO (Initiator)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.4.2
R57	DCO (Initiator)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.4.3
R234	DCO (Responder)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.6.1
R242	DCO (Responder)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.6.2
R250	DCO (Responder)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.6.3
R288	DCO (Responder)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.4.2
R296	DCO (Responder)	DCO Manager	Interacting	DC2DCControlInterface	UC 5.4.3

Identifier:	**Q16**				
Name:	**DC2DCDataInterface**				
Raw Entries:	9				
Description:	[DUPLICATE of DC2DCInterface] The DC2DC subsystem shall implements DATA interfaces to perform the VM migration between DCs				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	5 useCases
	DCO (Initiator)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.4.2
	DCO (Responder)				UC 5.4.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3

Raw List					
R20	DCO (Initiator)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.6.2
R30	DCO (Initiator)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.6.3
R47	DCO (Initiator)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.4.2
R58	DCO (Initiator)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.4.3
R235	DCO (Responder)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.6.1
R243	DCO (Responder)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.6.2
R251	DCO (Responder)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.6.3
R289	DCO (Responder)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.4.2
R297	DCO (Responder)	DCO Manager	Interacting	DC2DCDataInterface	UC 5.4.3

Identifier:	**Q14**				
Name:	**DC2DCInterface**				
Raw Entries:	40				
Description:	The system shall be able to interface with other federated DCs				
Aggregated List					
	2 keyActors	2 techActors	1 roles	1 requirements	16 useCases
	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.2
	DCO (Responder)	DOLFIN			UC 5.2.1
					UC 5.2.2

					UC 5.2.3
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.4.2
					UC 5.4.3
					UC 5.5.2
					UC 5.5.3
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
Raw List					
R2	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.2
R71	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.3.2
R80	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.3.3
R92	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.5.2
R102	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.5.3
R117	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.7.1
R126	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.7.2
R135	DCO (Initiator)	DCO Manager	Interacting	DC2DCInterface	UC 5.7.3
R171	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.7.1

R174	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.7.2
R177	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.7.3
R186	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.2.2
R189	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.2.3
R202	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.3.2
R207	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.3.3
R216	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.5.2
R222	DCO (Initiator)	DOLFIN	Interacting	DC2DCInterface	UC 5.5.3
R258	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.7.1
R263	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.7.2
R268	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.7.3
R273	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.2.1
R278	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.2.2
R283	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.2.3
R304	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.3.1
R308	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.3.2
R313	DCO (Responder)	DCO Manager	Interacting	DC2DCInterface	UC 5.3.3
R377	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.6.2
R380	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.6.3
R383	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.7.1
R386	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.7.2
R389	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.7.3
R392	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.2.1

R395	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.2.2
R398	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.2.3
R401	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.4.2
R404	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.4.3
R407	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.3.2
R409	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.3.3
R413	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.5.2
R418	DCO (Responder)	DOLFIN	Interacting	DC2DCInterface	UC 5.5.3

Identifier:	**Q53**				
Name:	**DC2SmartGridInterface**				
Raw Entries:	5				
Description:	In order for the Smart Grid Controller to be able to communicate with both the Smart Grid/EPs and the DOLFIN-enabled DCs, a proper interface serving the relevant requests/responses of such communication is needed.				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	DCO (Initiator)	DCO Manager	Interacting	DC2SmartGridInterface	UC 5.4.1
	DCO (Responder)				UC 5.4.2
					UC 5.4.3
Raw List					
R42	DCO (Initiator)	DCO Manager	Interacting	DC2SmartGridInterface	UC 5.4.1
R52	DCO (Initiator)	DCO	Interacting	DC2SmartGridInterface	UC 5.4.2

		Manager			
R63	DCO (Initiator)	DCO Manager	Interacting	DC2SmartGridInterface	UC 5.4.3
R294	DCO (Responder)	DCO Manager	Interacting	DC2SmartGridInterface	UC 5.4.2
R302	DCO (Responder)	DCO Manager	Interacting	DC2SmartGridInterface	UC 5.4.3

Identifier:	**Q54**				
Name:	**DCHeatControl**				
Raw Entries:	1				
Description:	In case the Smart Grid Controller retrieves a demand for a heat exchange state change from the DC part, the Energy Eff. Policy Maker and Actuator should be able to control the heat exchange towards the Smart City (turn it off or on).				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Manager	Interacting	DCHeatControl	UC 5.8
Raw List					
R112	DCO (Initiator)	DCO Manager	Interacting	DCHeatControl	UC 5.8

Identifier:	**Q55**				
Name:	**DCHVACMngmt**				
Raw Entries:	1				

Description:	[DUPLICATE of HVACPowerControl]				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Manager	Interacting	DCHVACMngmt	UC 5.8
Raw List					
R111	DCO (Initiator)	DCO Manager	Interacting	DCHVACMngmt	UC 5.8

Identifier:	**Q13**				
Name:	**DCManagement**				
Raw Entries:	4				
Description:	The system shall include a management entity for orchestarting the whole process				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	4 useCases
	DCO (Initiator)	DCO Manager	Interacting	DCManagement	UC 5.2
					UC 5.3.1
					UC 5.5.1
					UC 5.8
Raw List					
R1	DCO (Initiator)	DCO Manager	Interacting	DCManagement	UC 5.2

R64	DCO (Initiator)	DCO Manager	Interacting	DCManagement	UC 5.3.1
R84	DCO (Initiator)	DCO Manager	Interacting	DCManagement	UC 5.5.1
R106	DCO (Initiator)	DCO Manager	Interacting	DCManagement	UC 5.8

Identifier:	**Q1**				
Name:	**DCMonitoring**				
Raw Entries:	41				
Description:	The system shall monitor the DC infrastructure: IT devices, power units, other DC facilities that might be involved in the process of energy optimization.				
Aggregated List					
	2 keyActors	2 techActors	4 roles	1 requirements	19 useCases
	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.2
	DCO (Responder)	DCO Monitor	Input		UC 5.2.1
			Supp Service		UC 5.2.2
			Service		UC 5.2.3
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.4.1
					UC 5.4.2
					UC 5.4.3

					UC 5.5.2
					UC 5.5.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R10	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.6.1
R18	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.6.2
R28	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.6.3
R37	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.4.1
R45	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.4.2
R56	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.4.3
R70	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.3.2
R79	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.3.3
R91	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.5.2
R101	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.5.3
R116	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.7.1

R125	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.7.2
R134	DCO (Initiator)	DCO Manager	Interacting	DCMonitoring	UC 5.7.3
R145	DCO (Initiator)	DCO Monitor	Input	DCMonitoring	UC 5.2
R164	DCO (Initiator)	DCO Monitor	Supp Service	DCMonitoring	UC 5.8
R233	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.6.1
R241	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.6.2
R249	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.6.3
R257	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.7.1
R262	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.7.2
R267	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.7.3
R272	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.2.1
R277	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.2.2
R282	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.2.3
R287	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.4.2
R295	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.4.3
R303	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.3.1
R307	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.3.2
R312	DCO (Responder)	DCO Manager	Interacting	DCMonitoring	UC 5.3.3

R316	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.6.2
R322	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.6.3
R328	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.7.2
R332	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.7.3
R336	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.2.1
R340	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.2.2
R344	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.2.3
R348	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.4.2
R354	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.4.3
R360	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.3.2
R364	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.3.3
R368	DCO (Responder)	DCO Monitor	Service	DCMonitoring	UC 5.5.2

Identifier:	**Q20**				
Name:	**DCVMHandling**				
Raw Entries:	27				
Description:	The PolicyEnforcement subsystem shall be able to manipulate VMs (migrate them from server to server, from DC to DC, shutdown, etc.)				
Aggregated List					

	2 keyActors	1 techActors	1 roles	1 requirements	17 useCases
	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.2.1
	DCO (Responder)				UC 5.2.2
					UC 5.2.3
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.4.1
					UC 5.4.2
					UC 5.4.3
					UC 5.5.2
					UC 5.5.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
Raw List					
R11	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.6.1
R21	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.6.2
R31	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.6.3
R38	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.4.1
R48	DCO (Initiator)	DCO	Interacting	DCVMHandling	UC 5.4.2

		Manager			
R59	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.4.3
R72	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.3.2
R81	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.3.3
R93	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.5.2
R103	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.5.3
R118	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.7.1
R127	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.7.2
R136	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling	UC 5.7.3
R236	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.6.1
R244	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.6.2
R252	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.6.3
R259	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.7.1
R264	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.7.2
R269	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.7.3
R274	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.2.1
R279	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.2.2
R284	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.2.3

R290	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.4.2
R298	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.4.3
R305	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.3.1
R309	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.3.2
R314	DCO (Responder)	DCO Manager	Interacting	DCVMHandling	UC 5.3.3

Identifier:	**Q56**				
Name:	**DCVMHandling (VM params)**				
Raw Entries:	3				
Description:	In accordance with Ask for VMParamsAdjustment" this requirement implies that the Workload and VM Manager should be able to fulfills a request for changing the parameters of a certain set of VMs"				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling (VM params)	UC 5.3.1
					UC 5.5.1
					UC 5.8
Raw List					
R67	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling (VM params)	UC 5.3.1

R87	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling (VM params)	UC 5.5.1
R109	DCO (Initiator)	DCO Manager	Interacting	DCVMHandling (VM params)	UC 5.8
Identifier:	**Q25**				
Name:	**Energy Provider Responses Mngmnt**				
Raw Entries:	10				
Description:	The EPs shall be able to receive feedback from the consumers				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	10 useCases
	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.2
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.5.1
					UC 5.5.2
					UC 5.5.3
					UC 5.6
					UC 5.7
					UC 5.8
Raw List					
R426	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.6

R431	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.7
R436	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.2
R441	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.3.1
R445	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.3.2
R449	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.3.3
R454	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.5.1
R458	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.5.2
R462	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.5.3
R467	AGGREGATOR	EPS	Interacting	Energy Provider Responses Mngmnt	UC 5.8

Identifier:	**Q62**				
Name:	**Energy Provider Cost profile**				
Raw Entries:	3				
Description:	The Aggregators should be able to request for the tariff schemes of a certain set of EPs, through the Smart Grid Controller.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	3 useCases

	AGGREGATOR	EPS	Interacting	Energy Provider Cost profile	UC 5.3.1
					UC 5.5.1
					UC 5.8
Raw List					
R439	AGGREGATOR	EPS	Interacting	Energy Provider Cost profile	UC 5.3.1
R452	AGGREGATOR	EPS	Interacting	Energy Provider Cost profile	UC 5.5.1
R465	AGGREGATOR	EPS	Interacting	Energy Provider Cost profile	UC 5.8

Identifier:	**Q23**				
Name:	**Energy Provider Emission rate**				
Raw Entries:	3				
Description:	The EPs shall be able to produce energy production statistics and forecasts				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	AGGREGATOR	EPS	Interacting	Energy Provider Emission rate	UC 5.2
					UC 5.6
					UC 5.7
Raw List					

R424	AGGREGATOR	EPS	Interacting	Energy Provider Emission rate	UC 5.6
R429	AGGREGATOR	EPS	Interacting	Energy Provider Emission rate	UC 5.7
R434	AGGREGATOR	EPS	Interacting	Energy Provider Emission rate	UC 5.2

Identifier:	**Q24**				
Name:	**Energy Provider Requests Mngmnt**				
Raw Entries:	10				
Description:	The EPs shall be able to forward requests and to consumers about status changes of the energy provisioning				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	10 useCases
	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.2
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.5.1
					UC 5.5.2
					UC 5.5.3
					UC 5.6
					UC 5.7
					UC 5.8

Raw List					
R425	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.6
R430	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.7
R435	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.2
R440	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.3.1
R444	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.3.2
R448	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.3.3
R453	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.5.1
R457	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.5.2
R461	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.5.3
R466	AGGREGATOR	EPS	Interacting	Energy Provider Requests Mngmnt	UC 5.8

Identifier:	**Q57**				
Name:	**HandleExceptions**				
Raw Entries:	3				
Description:	When a customer denies a SLA renegotiation proposal or a synergetic DC denies cooperation, the DCO should be able to handle the rejection and consider alternative self-optimization				

	actions				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DOLFIN	Interacting	HandleExceptions	UC 5.3.2
	DCO (Responder)				UC 5.5.2
Raw List					
R205	DCO (Initiator)	DOLFIN	Interacting	HandleExceptions	UC 5.3.2
R219	DCO (Initiator)	DOLFIN	Interacting	HandleExceptions	UC 5.5.2
R415	DCO (Responder)	DOLFIN	Interacting	HandleExceptions	UC 5.5.2

Identifier:	**Q41**				
Name:	**HVACParamsInterface**				
Raw Entries:	1				
Description:	This interface provides information to the monitor subsystem concerning power dissipation and status of HVAC equipment.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Monitor	Input	HVACParamsInterface	UC 5.6

Raw List					
R141	DCO (Initiator)	DCO Monitor	Input	HVACParamsInterface	UC 5.6

Identifier:	**Q64**				
Name:	**HVACPowerContro**				
Raw Entries:	1				
Description:	Should be deleted.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Manager	Interacting	HVACPowerContro	UC 5.4.3
Raw List					
R62	DCO (Initiator)	DCO Manager	Interacting	HVACPowerContro	UC 5.4.3

Identifier:	**Q19**				
Name:	**HVACPowerControl**				
Raw Entries:	13				
Description:	The PolicyEnforcement subsystem shall be able to manage noni CT devices power dissipation (i.e. control/shutdown HVAC equipment)				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	9 useCases
	DCO (Initiator)	DCO	Interacting	HVACPowerControl	UC 5.4.1

		Manager			
	DCO (Responder)				UC 5.4.2
					UC 5.4.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
Raw List					
R14	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.6.1
R24	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.6.2
R34	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.6.3
R41	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.4.1
R51	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.4.2
R121	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.7.1
R130	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.7.2
R139	DCO (Initiator)	DCO Manager	Interacting	HVACPowerControl	UC 5.7.3
R239	DCO (Responder)	DCO Manager	Interacting	HVACPowerControl	UC 5.6.1
R247	DCO (Responder)	DCO Manager	Interacting	HVACPowerControl	UC 5.6.2
R255	DCO (Responder)	DCO Manager	Interacting	HVACPowerControl	UC 5.6.3

R293	DCO (Responder)	DCO Manager	Interacting	HVACPowerControl	UC 5.4.2
R301	DCO (Responder)	DCO Manager	Interacting	HVACPowerControl	UC 5.4.3

Identifier:	**Q42**				
Name:	**ICTParamsInterface**				
Raw Entries:	1				
Description:	This interface provides information to the monitor subsystem concerning power dissipation and status of IT equipment.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Monitor	Input	ICTParamsInterface	UC 5.6
Raw List					
R140	DCO (Initiator)	DCO Monitor	Input	ICTParamsInterface	UC 5.6

Identifier:	**Q18**				
Name:	**ICTPowerControl**				
Raw Entries:	18				
Description:	The PolicyEnforcement subsystem shall be able to manage ICT devices power dissipation (i.e. DVFS, ACPI, etc.)				
Aggregated List					

	2 keyActors	1 techActors	1 roles	1 requirements	13 useCases
	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.3.2
	DCO (Responder)				UC 5.3.3
					UC 5.4.1
					UC 5.4.2
					UC 5.4.3
					UC 5.5.2
					UC 5.5.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
Raw List					
R13	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.6.1
R23	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.6.2
R33	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.6.3
R40	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.4.1
R50	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.4.2
R61	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.4.3
R74	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.3.2

R83	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.3.3
R95	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.5.2
R105	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.5.3
R120	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.7.1
R129	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.7.2
R138	DCO (Initiator)	DCO Manager	Interacting	ICTPowerControl	UC 5.7.3
R238	DCO (Responder)	DCO Manager	Interacting	ICTPowerControl	UC 5.6.1
R246	DCO (Responder)	DCO Manager	Interacting	ICTPowerControl	UC 5.6.2
R254	DCO (Responder)	DCO Manager	Interacting	ICTPowerControl	UC 5.6.3
R292	DCO (Responder)	DCO Manager	Interacting	ICTPowerControl	UC 5.4.2
R300	DCO (Responder)	DCO Manager	Interacting	ICTPowerControl	UC 5.4.3

Identifier:	**Q58**				
Name:	**LogRenegotiationSteps**				
Raw Entries:	2				
Description:	The SLA Renegotiation Controller should be able to persist the SLA renegotiation steps upon happening {Siemens} I confirm the explanation given, this logs should support development for Monitoring, Analytics, Audit and Security.				

Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DCO Manager	Interacting	LogRenegotiationSteps	UC 5.2
					UC 5.5.2
Raw List					
R6	DCO (Initiator)	DCO Manager	Interacting	LogRenegotiationSteps	UC 5.2
R96	DCO (Initiator)	DCO Manager	Interacting	LogRenegotiationSteps	UC 5.5.2

Identifier:	**Q69**				
Name:	**LogReplanningSteps**				
Raw Entries:	3				
Description:	This logs should support development for Monitoring, Analytics, Audit and Security. It can be implemented as a utility package for Management of Replanting of energy scheme.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	DCO (Initiator)	DCO Manager	Interacting	LogReplanningSteps	UC 5.3.1
					UC 5.5.1
					UC 5.8

Raw List					
R68	DCO (Initiator)	DCO Manager	Interacting	LogReplanningSteps	UC 5.3.1
R88	DCO (Initiator)	DCO Manager	Interacting	LogReplanningSteps	UC 5.5.1
R110	DCO (Initiator)	DCO Manager	Interacting	LogReplanningSteps	UC 5.8

Identifier:	**Q63**				
Name:	**MgmntCostDetection**				
Raw Entries:	3				
Description:	The Energy Eff. Policy Maker and Actuator should be able to detect situations when the operating costs (due to energy consumption) of the DC are higher than minimum and proceed to the necessary actions to overturn this situations. [Siemens] I agree with given definition.				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DCO Manager	Interacting	MgmntCostDetection	UC 5.3.3
	DCO (Responder)				UC 5.5.3
Raw List					
R75	DCO (Initiator)	DCO Manager	Interacting	MgmntCostDetection	UC 5.3.3
R97	DCO (Initiator)	DCO Manager	Interacting	MgmntCostDetection	UC 5.5.3

R311	DCO (Responder)	DCO Manager	Interacting	MgmntCostDetection	UC 5.3.3

Identifier:	**Q70**				
Name:	**MgmntCostsOpportunity**				
Raw Entries:	4				
Description:	Represents an Analytics functional package used my Management Interfaces to evaluate the impact of potential actions taking in account history logs and current state. Main result is a decision support regarding the opportunity to search and/or accept a different cost scheme and it's time based impact on managed DC.				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	3 useCases
	DCO (Initiator)	DOLFIN	Interacting	MgmntCostsOpportunity	UC 5.3.2
	DCO (Responder)				UC 5.5.2
					UC 5.5.3
Raw List					
R203	DCO (Initiator)	DOLFIN	Interacting	MgmntCostsOpportunity	UC 5.3.2
R217	DCO (Initiator)	DOLFIN	Interacting	MgmntCostsOpportunity	UC 5.5.2
R225	DCO (Initiator)	DOLFIN	Interacting	MgmntCostsOpportunity	UC 5.5.3
R421	DCO (Responder)	DOLFIN	Interacting	MgmntCostsOpportunity	UC 5.5.3

Identifier:	**Q22**				
Name:	**MgmntDOLFINResponses**				

Raw Entries:	10				
Description:	The system shall be able to send feedback and notifications to the EPs				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	10 useCases
	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.2
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.5.1
					UC 5.5.2
					UC 5.5.3
					UC 5.6
					UC 5.7
					UC 5.8
Raw List					
R423	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.6
R428	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.7
R433	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.2
R438	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.3.1
R443	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.3.2
R447	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.3.3
R451	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.5.1
R456	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.5.2
R460	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.5.3
R464	AGGREGATOR	Broker	Service	MgmntDOLFINResponses	UC 5.8

Identifier:	**Q21**				
Name:	**MgmntEPsRequests**				
Raw Entries:	10				
Description:	The system shall handle incoming EP requests and notifications about changes and trends in the energy provision service				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	10 useCases
	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.2
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.5.1
					UC 5.5.2
					UC 5.5.3
					UC 5.6
					UC 5.7
					UC 5.8
Raw List					
R422	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.6
R427	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.7
R432	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.2
R437	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.3.1
R442	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.3.2
R446	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.3.3

R450	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.5.1
R455	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.5.2
R459	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.5.3
R463	AGGREGATOR	Broker	Service	MgmntEPsRequests	UC 5.8

Identifier:	**Q5**				
Name:	**MgmntPolicyDecision**				
Raw Entries:	45				
Description:	The management subsystem shall automatically identify the actions that could be taken to optimize the DC energy state.				
Aggregated List					
	2 keyActors	2 techActors	1 roles	1 requirements	20 useCases
	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.2.1
	DCO (Responder)	DOLFIN			UC 5.2.2
					UC 5.2.3
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.4.1
					UC 5.4.2
					UC 5.4.3
					UC 5.5.1
					UC 5.5.2
					UC 5.5.3

					UC 5.6
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R8	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.6.1
R16	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.6.2
R26	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.6.3
R36	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.4.1
R44	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.4.2
R54	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.4.3
R69	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.3.2
R77	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.3.3
R90	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.5.2
R99	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.5.3
R114	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.7.1
R123	DCO (Initiator)	DCO	Interacting	MgmntPolicyDecision	UC 5.7.2

		Manager			
R132	DCO (Initiator)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.7.3
R168	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.6
R170	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.7.1
R173	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.7.2
R176	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.7.3
R181	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.2.1
R185	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.2.2
R188	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.2.3
R191	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.4.1
R193	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.4.2
R195	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.4.3
R197	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.3.1
R201	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.3.2
R206	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.3.3
R209	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.5.1
R215	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.5.2
R221	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.5.3
R226	DCO (Initiator)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.8
R231	DCO (Responder)	DCO Manager	Interacting	MgmntPolicyDecision	UC 5.6.1
R376	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.6.2
R379	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.6.3
R382	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.7.1
R385	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.7.2
R388	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.7.3
R391	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.2.1
R394	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.2.2

R397	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.2.3
R400	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.4.2
R403	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.4.3
R406	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.3.2
R408	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.3.3
R412	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.5.2
R417	DCO (Responder)	DOLFIN	Interacting	MgmntPolicyDecision	UC 5.5.3

Identifier:	**Q6**				
Name:	**MgmntSLARenegotiation**				
Raw Entries:	18				
Description:	The management subsystem shall be able to (automatically) request the SLA renegotiation to better adapt the SLA KPI to current optimal assets.				
Aggregated List					
	5 keyActors	4 techActors	1 roles	1 requirements	15 useCases
	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.2.1
	DCO (Responder)	DOLFIN			UC 5.2.3
	END CUSTOMERS	Customer			UC 5.3.1
	UTIL OPERATOR	Provider			UC 5.3.3
	Utilities				UC 5.4.3
					UC 5.5.2
					UC 5.5.3
					UC 5.6
					UC 5.6.1
					UC 5.6.2

					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R9	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.6.1
R17	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.6.2
R27	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.6.3
R55	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.4.3
R78	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.3.3
R100	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.5.3
R115	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.7.1
R124	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.7.2
R133	DCO (Initiator)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.7.3
R180	DCO (Initiator)	DOLFIN	Interacting	MgmntSLARenegotiation	UC 5.2.1
R214	DCO (Initiator)	DOLFIN	Interacting	MgmntSLARenegotiation	UC 5.5.2
R232	DCO (Responder)	DCO Manager	Interacting	MgmntSLARenegotiation	UC 5.6.1
R469	END CUSTOMERS	Customer	Interacting	MgmntSLARenegotiation	UC 5.6
R471	END CUSTOMERS	Customer	Interacting	MgmntSLARenegotiation	UC 5.2.1
R473	END CUSTOMERS	Customer	Interacting	MgmntSLARenegotiation	UC 5.2.3
R475	END CUSTOMERS	Customer	Interacting	MgmntSLARenegotiation	UC 5.4.3

R477	UTIL OPERATOR	Provider	Interacting	MgmntSLARenegotiation	UC 5.3.1
R479	Utilities	Customer	Interacting	MgmntSLARenegotiation	UC 5.8

Identifier:	**Q4**				
Name:	**MgmntSubOptimalDetection**				
Raw Entries:	39				
Description:	The management subsystem shall detect DC suboptimal states (i.e. high KPI) by utilizing current DC metrics and defined thresholds				
Aggregated List					
	2 keyActors	2 techActors	1 roles	1 requirements	14 useCases
	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.2.1
	DCO (Responder)	DOLFIN			UC 5.2.2
					UC 5.2.3
					UC 5.3.3
					UC 5.4.1
					UC 5.4.2
					UC 5.4.3
					UC 5.6
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1
					UC 5.7.2
					UC 5.7.3

Raw List					
R7	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.6.1
R15	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.6.2
R25	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.6.3
R35	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.4.1
R43	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.4.2
R53	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.4.3
R113	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.7.1
R122	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.7.1
R131	DCO (Initiator)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.7.3
R169	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.6
R172	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.7.1
R175	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.7.2
R178	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.7.3

R187	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.2.2
R190	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.2.3
R192	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.4.1
R194	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.4.2
R196	DCO (Initiator)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.4.3
R230	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.6.1
R240	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.6.2
R248	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.6.3
R256	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.7.1
R261	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.7.2
R266	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.7.3
R271	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.2.1
R276	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.2.2
R281	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.2.3

R286	DCO (Responder)	DCO Manager	Interacting	MgmntSubOptimalDetection	UC 5.4.2
R378	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.6.2
R381	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.6.3
R384	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.7.1
R387	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.7.2
R390	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.7.3
R393	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.2.1
R396	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.2.2
R399	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.2.3
R402	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.4.2
R405	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.4.3
R410	DCO (Responder)	DOLFIN	Interacting	MgmntSubOptimalDetection	UC 5.3.3

Identifier:	**Q71**				
Name:	**MgmntTransfer**				
Raw Entries:	3				

Description:	Functional package responsible for technology independent implementation of DC to DC content transfer.				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DOLFIN	Interacting	MgmntTransfer	UC 5.3.2
	DCO (Responder)				UC 5.5.2
Raw List					
R204	DCO (Initiator)	DOLFIN	Interacting	MgmntTransfer	UC 5.3.2
R218	DCO (Initiator)	DOLFIN	Interacting	MgmntTransfer	UC 5.5.2
R414	DCO (Responder)	DOLFIN	Interacting	MgmntTransfer	UC 5.5.2

Identifier:	**Q68**				
Name:	**MgmtWorkload**				
Raw Entries:	2				
Description:	A functional subinterface of DCVMHandling.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DCO Manager	Interacting	MgmtWorkload	UC 5.3.3
					UC 5.5.3

Raw List					
R76	DCO (Initiator)	DCO Manager	Interacting	MgmtWorkload	UC 5.3.3
R98	DCO (Initiator)	DCO Manager	Interacting	MgmtWorkload	UC 5.5.3

Identifier:	**Q72**				
Name:	**MonitorCostDataRequests**				
Raw Entries:	1				
Description:	Listener of Cost Information requests, format and distribute to subscribed interfaces in published formats.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	UTIL OPERATOR	Provider	Interacting	MonitorCostDataRequests	UC 5.3.1
Raw List					
R476	UTIL OPERATOR	Provider	Interacting	MonitorCostDataRequests	UC 5.3.1

Identifier:	**Q9**				
Name:	**MonitorInterface**				
Raw Entries:	19				

Description:	The monitoring subsystem shall be able to present metrics-information to the management subsystem This interface is considered to expose in a normalized form a number of metrics and their frequency. It could receive the metrics to retrieve, their frequency and response format (e.g. publish to a message broker).				
Aggregated List					
	2 keyActors	1 techActors	3 roles	1 requirements	19 useCases
	DCO (Initiator)	DCO Monitor	Output	MonitorInterface	UC 5.2
	DCO (Responder)		Service		UC 5.2.1
			Supp Service		UC 5.2.2
					UC 5.2.3
					UC 5.3
					UC 5.3.2
					UC 5.3.3
					UC 5.4
					UC 5.4.2
					UC 5.4.3
					UC 5.5
					UC 5.5.2
					UC 5.5.3
					UC 5.6
					UC 5.6.2
					UC 5.6.3

					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R144	DCO (Initiator)	DCO Monitor	Output	MonitorInterface	UC 5.6
R148	DCO (Initiator)	DCO Monitor	Service	MonitorInterface	UC 5.2
R152	DCO (Initiator)	DCO Monitor	Supp Service	MonitorInterface	UC 5.4
R158	DCO (Initiator)	DCO Monitor	Supp Service	MonitorInterface	UC 5.3
R162	DCO (Initiator)	DCO Monitor	Supp Service	MonitorInterface	UC 5.5
R166	DCO (Initiator)	DCO Monitor	Supp Service	MonitorInterface	UC 5.8
R318	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.6.2
R324	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.6.3
R330	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.7.2
R334	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.7.3
R338	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.2.1
R342	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.2.2
R346	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.2.3
R350	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.4.2
R356	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.4.3

R362	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.3.2
R366	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.3.3
R370	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.5.2
R374	DCO (Responder)	DCO Monitor	Service	MonitorInterface	UC 5.5.3

Identifier:	**Q12**				
Name:	**MonitorOpenDataFmt**				
Raw Entries:	5				
Description:	The monitoring subsystem should utilize open data formats is available				
Aggregated List					
	2 keyActors	1 techActors	2 roles	1 requirements	5 useCases
	DCO (Initiator)	DCO Monitor	Supp Service	MonitorOpenDataFmt	UC 5.4
	DCO (Responder)		Service		UC 5.4.2
					UC 5.4.3
					UC 5.6.2
					UC 5.6.3
Raw List					
R154	DCO (Initiator)	DCO Monitor	Supp Service	MonitorOpenDataFmt	UC 5.4
R320	DCO (Responder)	DCO Monitor	Service	MonitorOpenDataFmt	UC 5.6.2
R326	DCO (Responder)	DCO Monitor	Service	MonitorOpenDataFmt	UC 5.6.3

R352	DCO (Responder)	DCO Monitor	Service	MonitorOpenDataFmt	UC 5.4.2
R358	DCO (Responder)	DCO Monitor	Service	MonitorOpenDataFmt	UC 5.4.3

Identifier:	**Q7**				
Name:	**MonitorParams**				
Raw Entries:	18				
Description:	The monitoring subsystem shall always monitor DC parameters. This requirements identify the necessity to make a formal collection of all parameters that have to be monitored to build KPI and other triggers levels used to evaluate the DC performance and schedule the appropriate policies' actions. The list of params to be published should be configurable in terms of frequency of measurements and way to deliver. Also, the access to this functions should be controlled.				
Aggregated List					
	2 keyActors	1 techActors	3 roles	1 requirements	18 useCases
	DCO (Initiator)	DCO Monitor	Input	MonitorParams	UC 5.2
	DCO (Responder)		Supp Service		UC 5.2.1
			Service		UC 5.2.2
					UC 5.2.3

					UC 5.3
					UC 5.3.2
					UC 5.3.3
					UC 5.4
					UC 5.4.2
					UC 5.4.3
					UC 5.5
					UC 5.5.2
					UC 5.5.3
					UC 5.6.2
					UC 5.6.3
					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R146	DCO (Initiator)	DCO Monitor	Input	MonitorParams	UC 5.2
R150	DCO (Initiator)	DCO Monitor	Supp Service	MonitorParams	UC 5.4
R156	DCO (Initiator)	DCO Monitor	Supp Service	MonitorParams	UC 5.3
R160	DCO (Initiator)	DCO Monitor	Supp Service	MonitorParams	UC 5.5
R165	DCO (Initiator)	DCO Monitor	Supp Service	MonitorParams	UC 5.8
R317	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.6.2
R323	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.6.3
R329	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.7.2

R333	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.7.3
R337	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.2.1
R341	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.2.2
R345	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.2.3
R349	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.4.2
R355	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.4.3
R361	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.3.2
R365	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.3.3
R369	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.5.2
R372	DCO (Responder)	DCO Monitor	Service	MonitorParams	UC 5.5.3

Identifier:	**Q11**				
Name:	**MonitorResilience**				
Raw Entries:	5				
Description:	The monitoring subsystem should be as much resilient as possible (i.e. implement heartbeat, watchdog functions)				
Aggregated List					
	2 keyActors	1 techActors	2 roles	1 requirements	5 useCases
	DCO (Initiator)	DCO Monitor	Supp Service	MonitorResilience	UC 5.4

	DCO (Responder)		Service		UC 5.4.2
					UC 5.4.3
					UC 5.6.2
					UC 5.6.3
Raw List					
R155	DCO (Initiator)	DCO Monitor	Supp Service	MonitorResilience	UC 5.4
R321	DCO (Responder)	DCO Monitor	Service	MonitorResilience	UC 5.6.2
R327	DCO (Responder)	DCO Monitor	Service	MonitorResilience	UC 5.6.3
R353	DCO (Responder)	DCO Monitor	Service	MonitorResilience	UC 5.4.2
R359	DCO (Responder)	DCO Monitor	Service	MonitorResilience	UC 5.4.3

Identifier:	**Q66**				
Name:	**MonitorSLA Requests**				
Raw Entries:	7				
Description:	The SLA Renegotiation Controller should keep track of the SLA Renegotiation requests for historical reasons.				
Aggregated List					
	3 keyActors	2 techActors	1 roles	1 requirements	5 useCases
	DCO (Initiator)	DOLFIN	Interacting	MonitorSLA Requests	UC 5.2.1
	DCO (Responder)	Customer			UC 5.2.3

	END CUSTOMERS				UC 5.4.3
					UC 5.5.2
					UC 5.6
Raw List					
R179	DCO (Initiator)	DOLFIN	Interacting	MonitorSLA Requests	UC 5.2.1
R213	DCO (Initiator)	DOLFIN	Interacting	MonitorSLA Requests	UC 5.5.2
R411	DCO (Responder)	DOLFIN	Interacting	MonitorSLA Requests	UC 5.5.2
R468	END CUSTOMERS	Customer	Interacting	MonitorSLA Requests	UC 5.6
R470	END CUSTOMERS	Customer	Interacting	MonitorSLA Requests	UC 5.2.1
R472	END CUSTOMERS	Customer	Interacting	MonitorSLA Requests	UC 5.2.3
R474	END CUSTOMERS	Customer	Interacting	MonitorSLA Requests	UC 5.4.3

Identifier:	**Q8**				
Name:	**MonitorTranslate**				
Raw Entries:	6				
Description:	The monitoring subsystem shall always translate params to metrics (KPI)				
Aggregated List					
	2 keyActors	1 techActors	2 roles	1 requirements	6 useCases
	DCO (Initiator)	DCO Monitor	Service	MonitorTranslate	UC 5.2
	DCO (Responder)		Supp Service		UC 5.3
					UC 5.4
					UC 5.5
					UC 5.5.3
					UC 5.6

Raw List					
R142	DCO (Initiator)	DCO Monitor	Service	MonitorTranslate	UC 5.6
R147	DCO (Initiator)	DCO Monitor	Service	MonitorTranslate	UC 5.2
R151	DCO (Initiator)	DCO Monitor	Supp Service	MonitorTranslate	UC 5.4
R157	DCO (Initiator)	DCO Monitor	Supp Service	MonitorTranslate	UC 5.3
R161	DCO (Initiator)	DCO Monitor	Supp Service	MonitorTranslate	UC 5.5
R373	DCO (Responder)	DCO Monitor	Service	MonitorTranslate	UC 5.5.3

Identifier:	**Q10**				
Name:	**MonitorVariousData**				
Raw Entries:	18				
Description:	The monitoring subsystem shall handle many kind of data and with different nature (i.e. performance information, energy consumption, power dissipation, VM status and deployment, semi-static parameters, etc) This function will complement in terms of Data the other peer Monitoring interfaces.				
Aggregated List					
	2 keyActors	1 techActors	3 roles	1 requirements	18 useCases
	DCO (Initiator)	DCO Monitor	Output	MonitorVariousData	UC 5.2

	DCO (Responder)		Supp Service		UC 5.2.1
			Service		UC 5.2.2
					UC 5.2.3
					UC 5.3
					UC 5.3.2
					UC 5.3.3
					UC 5.4
					UC 5.4.2
					UC 5.4.3
					UC 5.5
					UC 5.5.2
					UC 5.5.3
					UC 5.6.2
					UC 5.6.3
					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R149	DCO (Initiator)	DCO Monitor	Output	MonitorVariousData	UC 5.2
R153	DCO (Initiator)	DCO Monitor	Supp Service	MonitorVariousData	UC 5.4
R159	DCO (Initiator)	DCO Monitor	Supp Service	MonitorVariousData	UC 5.3
R163	DCO (Initiator)	DCO Monitor	Supp Service	MonitorVariousData	UC 5.5
R167	DCO (Initiator)	DCO Monitor	Supp Service	MonitorVariousData	UC 5.8
R319	DCO (Responder)	DCO	Service	MonitorVariousData	UC 5.6.2

		Monitor			
R325	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.6.3
R331	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.7.2
R335	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.7.3
R339	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.2.1
R343	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.2.2
R347	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.2.3
R351	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.4.2
R357	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.4.3
R363	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.3.2
R367	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.3.3
R371	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.5.2
R375	DCO (Responder)	DCO Monitor	Service	MonitorVariousData	UC 5.5.3

Identifier:	**Q67**				
Name:	**NegotiateSLAUpdate**				
Raw Entries:	2				

Description:	When required to do so (e.g. by the MgmntSLARenegotiation registration), the SLA Renegotiation Controller should be able to actually perform the renegotiation. [Siemens] Agree with explanation, use case owners should comment on this.				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DOLFIN	Interacting	NegotiateSLAUpdate	UC 5.5.3
	DCO (Responder)				
Raw List					
R224	DCO (Initiator)	DOLFIN	Interacting	NegotiateSLAUpdate	UC 5.5.3
R420	DCO (Responder)	DOLFIN	Interacting	NegotiateSLAUpdate	UC 5.5.3

Identifier:	**Q65**				
Name:	**NegotiateSLAUpdates**				
Raw Entries:	2				
Description:	[DUPLICATE of NegotiateSLAUpdate]				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DOLFIN	Interacting	NegotiateSLAUpdates	UC 5.5.2
	DCO (Responder)				UC 5.5.3

Raw List					
R220	DCO (Initiator)	DOLFIN	Interacting	NegotiateSLAUpdates	UC 5.5.3
R416	DCO (Responder)	DOLFIN	Interacting	NegotiateSLAUpdates	UC 5.5.2

Identifier:	**Q73**				
Name:	**NegotiateVMTransfer**				
Raw Entries:	3				
Description:	Since DC to DC transfer may involve a multi-part balancing the way of what and how to transfer should be negotiated at VM level.				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	2 useCases
	DCO (Initiator)	DOLFIN	Interacting	NegotiateVMTransfer	UC 5.3.3
	DCO (Responder)				UC 5.5.3
Raw List					
R208	DCO (Initiator)	DOLFIN	Interacting	NegotiateVMTransfer	UC 5.3.3
R223	DCO (Initiator)	DOLFIN	Interacting	NegotiateVMTransfer	UC 5.5.3
R419	DCO (Responder)	DOLFIN	Interacting	NegotiateVMTransfer	UC 5.5.3

Identifier:	**Q17**				
Name:	**PolicyEnfoercementInterface**				
Raw Entries:	31				

Description:	The system shall be able to enforce different power dissipation policies (i.e. using DVFS and/or ACPI functions, control HVAC equipment and also through VM migration). The PolicyEnforcement subsystem shall utilize an interface to communicate with the management subsystem that should be common across all platforms				
Aggregated List					
	2 keyActors	1 techActors	1 roles	1 requirements	20 useCases
	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.2
	DCO (Responder)				UC 5.2.1
					UC 5.2.2
					UC 5.2.3
					UC 5.3.1
					UC 5.3.2
					UC 5.3.3
					UC 5.4.1
					UC 5.4.2
					UC 5.4.3
					UC 5.5.1
					UC 5.5.2
					UC 5.5.3
					UC 5.6.1
					UC 5.6.2
					UC 5.6.3
					UC 5.7.1

					UC 5.7.2
					UC 5.7.3
					UC 5.8
Raw List					
R4	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.2
R12	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.6.1
R22	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.6.2
R32	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.6.3
R39	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.4.1
R49	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.4.2
R60	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.4.3
R66	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.3.1
R73	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.3.2
R82	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.3.3
R86	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.5.1
R94	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.5.2

R104	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.5.3
R108	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.8
R119	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.7.1
R128	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.7.2
R137	DCO (Initiator)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.7.3
R237	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.6.1
R245	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.6.2
R253	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.6.3
R260	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.7.1
R265	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.7.2
R270	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.7.3
R275	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.2.1
R280	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.2.2
R285	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.2.3

R291	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.4.2
R299	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.4.3
R306	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.3.1
R310	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.3.2
R315	DCO (Responder)	DCO Manager	Interacting	PolicyEnfoercementInterface	UC 5.3.3

Identifier:	**Q61**				
Name:	**Provide Public Policies**				
Raw Entries:	1				
Description:	The Authorities should be able to provide the set of acceptable policies that rule the interaction among customers and DCs				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	Authorities	Customer	Interacting	Provide Public Policies	UC 5.3.2
Raw List					
R478	Authorities	Customer	Interacting	Provide Public Policies	UC 5.3.2

Identifier:	**Q60**				
Name:	**ReceiveRenegotiationRequest**				

Raw Entries:	1				
Description:	Maybe add it in more UCs] The SLA Renegotiation Controller should be able to receive and handle a SLA renegotiation request from the Energy Efficient Policy Maker and Actuator.				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Manager	Interacting	ReceiveRenegotiationRequest	UC 5.2
Raw List					
R5	DCO (Initiator)	DCO Manager	Interacting	ReceiveRenegotiationRequest	UC 5.2

Identifier:	**Q3**				
Name:	**StoreParams**				
Raw Entries:	1				
Description:	Store performance and power-dissipation data (historical data)				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	1 useCases
	DCO (Initiator)	DCO Monitor	Service	StoreParams	UC 5.6
Raw List					
R143	DCO (Initiator)	DCO	Service	StoreParams	UC 5.6

		Monitor			

Identifier:	**Q59**				
Name:	**VMHandling**				
Raw Entries:	4				
Description:	[DUPLICATE of DCVMHandling]				
Aggregated List					
	1 keyActors	1 techActors	1 roles	1 requirements	4 useCases
	DCO (Initiator)	DCO Manager	Interacting	VMHandling	UC 5.2
					UC 5.3.1
					UC 5.5.1
					UC 5.8
Raw List					
R3	DCO (Initiator)	DCO Manager	Interacting	VMHandling	UC 5.2
R65	DCO (Initiator)	DCO Manager	Interacting	VMHandling	UC 5.3.1
R85	DCO (Initiator)	DCO Manager	Interacting	VMHandling	UC 5.5.1
R107	DCO (Initiator)	DCO Manager	Interacting	VMHandling	UC 5.8