•Pesch, D., McGibney, A., Sobonski, P., Rea, S., Scherer, Th., Chen, L., Engbersen, T., Mehta, D., O'Sullivan, B., Pages, E., Townley, J., Kasinathan, Dh., Torrens, J.I., Zavrel, V. & Hensen, J.L.M. (2015). The GENiC architecture for integrated data centre energy management. Proceedings 2015 IEEE/ACM 8th International Conference on Utility and Cloud Computing (UCC 2015), 7-10 December 2015, Limassol, Cyprus (pp. 540-546). Piscataway: IEEE.

The GENiC Architecture for Integrated Data Centre Energy Management

D. Pesch[1], A. McGibney[1], P. Sobonski[1], S. Rea[1], T. Scherer[2], L. Chen[2], T. Engbersen[2], D. Mehta[3], B. O'Sullivan[3], E. Pages[4], J. Townley[5], D. Kasinathan[5], J. I. Torrens[6], V. Zavrel[6], J. Hensen[6]

[1]Cork Institute of Technology, Ireland
{dirk.pesch, alan.mcgibney, piotr.sobonski, susan.rea}@cit.ie

[2]IBM Research – Zurich, Rüschlikon, Switzerland
{tsc, yic, apj}@zurich.ibm.com

[3]Insight Centre, University College Cork, Ireland
{deepak.mehta, barry.osullivan}@insight-centre.org

[4]Atos Spain, Barcelona, Spain
enric.pages@atos.net

[5]United Technologies Research Centre – Cork, Ireland
{townlejm, kasinad}@utrc.utc.com

[6]Building Physics and Services
TU Eindhoven, The Netherlands
{j.i.torrens, j.hensen}@tue.nl

***Abstract*— We present an architecture for integrated data centre energy management developed in the EC funded GENiC project. The architecture was devised to create a platform that can integrate functions for workload management, cooling, power management and control of heat recovery for future, highly efficient data centres. The architecture is based on a distributed systems approach that allows the integration of components developed by several entities through defined interfaces and data formats. We also present use cases for the architecture, a brief description of the project's prototypical implementation, evaluation metrics and some lessons learned.**

Keywords—Data centres, integrated energy management, renewable energy, optimisation

I. Introduction

Data centres have become a critical part of modern life, not just for businesses but for society at large with the huge penetration of smart phone usage, digital entertainment through media streaming, and the expected growth in the Internet of Things all relying on data centres. However, data centres are also a significant primary energy user and now consume 1.3% of worldwide electricity and with the increasing move towards cloud computing and storage, and everything as a service type computing, energy consumption is expected to grow to 8% by 2020 [1, 2]. The data centres of large cloud service providers consume many megawatts of power with corresponding annual electricity bills in the order of tens of millions of dollars, e.g. Google with over 1,120 GWh and $67M and Microsoft with over 600 GWh and $36M in 2010 [3].

On average, computing consumes 60% of total energy in data centres while cooling consumes 35%. New technologies have the potential to lead to a 40% reduction of energy consumption, but computation and cooling typically operate without joint coordination or optimisation. While server energy management can reduce energy use at CPU, rack, and overall data centre level, dynamic computation scheduling is not integrated with cooling. Data centre cooling typically operates at constant cold air temperature to protect the hottest server racks while local fans distribute the temperature across racks. However, these local server controls are typically not integrated with room cooling systems, which means that it is not possible to optimise chillers, air fans and server fans as a whole system.

The integration of renewable energy sources (RES) has received limited interest from the data centre community due to lack of interoperability of generation, storage and heat recovery and current installation and maintenance costs versus payback [4]. The intermittency of renewable energy generation is also a critical factor in an environment with very strict service level agreements and essentially 100% uptime requirements. The adoption of new technologies related to computing, cooling, generation, energy storage, and waste heat recovery individually requires sophisticated controls, but no single manufacturer provides a complete system so integration between control systems does not exist. Funded by the European Commission, the GENiC project (http://www.projectgenic.eu) develops integrated cooling and computing control strategies in conjunction with innovative power management concepts that incorporate large elements of renewable electrical power supply, energy storage, and waste heat management. The GENiC project's aim is to address this issue by developing an integrated management and control platform for data centre wide optimisation of energy consumption, reduction of carbon emissions and increased renewables usage through integrating monitoring and control of computation, data storage, cooling, local power generation, and waste heat recovery. The platform defines interfaces and common data formats, includes control and optimisation functions and decision support. We aim to achieve and demonstrate a substantial reduction in energy consumption by deploying the platform in demonstration data centres. A further premise of GENiC is that the energy consuming equipment in data centres must be supplemented with renewable energy generation and energy storage equipment and operated as a complete system to achieve an optimal energy and emissions outcome. This vision is centred on the development of a hierarchical control system to operate all of the primary data centre components in an optimal and coordinated manner. The goal is to minimize energy use through manipulation of local equipment controller set points and provision of optimised control of computing load and cooling distribution.

In this paper we present the overall GENiC system architecture in order to address some of the challenges associated with an integrated approach to data centre management and discuss the first prototype implementation of the proposed architecture. The challenges are outlined in Section II, the architecture and associated use cases are presented in Section III and the prototype implementation is discussed in Section IV. Finally the paper draws conclusions on the experience we have gained during the development and implementation of the GENiC architecture.

II. Challenges in Data Centre Energy Management

Data centres have evolved into critical information technology (IT) infrastructure and much of today's IT services, both for businesses and consumers, depend on their operation. As pointed out, they consume an increasing amount of energy and contribute significantly to CO_2 emissions. However, opportunities exist to enhance the energy and power management of data centres in conjunction with renewable energy generation and integration with their surrounding infrastructure. Work has been done on powering of data centres by renewable energy [5 and references therein], but this has not been fully integrated into a complete energy management system considering coordinated workload management, cooling, powering, and heat recovery management. While much work has focused on integrated energy management for data centres [6, 7], there is still a lack of an overall consideration of energy usage and powering with the recovery of waste heat as part of an overall thermal management approach. In order to bring the elements of workload management, cooling, powering and heat recovery together in such a way that it will be possible to achieve a high level of renewable energy powering of data centres, a comprehensive integrated energy management system is needed. The challenges that such a system needs to address are

- Comprehensive, scalable integration of workload management with cooling approaches.
- Effective power management with a high level of renewable energy supply integration while meeting service level agreements. For example, how to deliver on such agreements while trying to execute workload during time windows where power supply is cheaper.
- Strategies for waste heat recovery in conjunction with the heating needs of surrounding areas.
- Design and decision support tools for data centres' energy management. Effective monitoring and fault management.

III. GENiC Architecture

To address the challenges outlined above, the GENiC project has developed a high level architecture for an integrated design, management and control platform, targeting data centre wide optimisation of energy consumption by encapsulating monitoring and control of IT workload, data centre cooling, local power generation, energy storage, and waste heat recovery. In the following, a functional specification of the GENiC architecture is presented and an overview of the integration framework is provided. The applicability of the proposed functional architecture is illustrated by two use cases. More detail can be found in [8, 9].

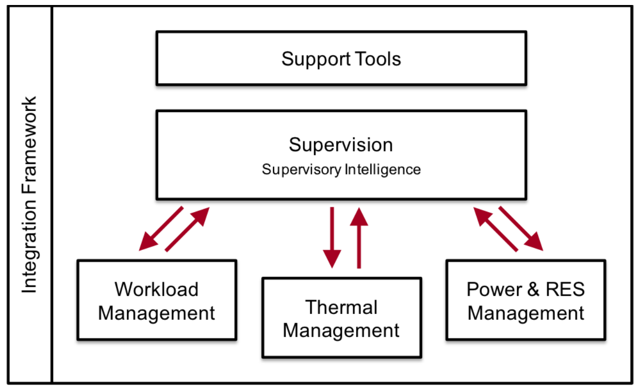

Fig. 1. High level overview of the GENiC architecture (from [9])

A. GENIC Functional Architecture

The GENiC system integrates workload management, thermal management and power management by using a hierarchical control concept to coordinate the management sub-systems in an optimal manner with respect to the cost of energy consumption and environmental impact, and cost policies. Fig. 1 provides a high level overview of the proposed GENiC system architecture, which consists of six functional groups known as GENiC Component Groups (GCGs):

- The **Workload Management GCG** is responsible for monitoring, analyzing, predicting, allocating, and actuating IT workload within the data centre.
- The **Thermal Management GCG** is responsible for monitoring the thermal environment and cooling systems in the data centre, predicting temperature profiles and cooling demand, and optimally coordinating and actuating the cooling systems.
- The **Power & RES Management GCG** is responsible for monitoring and predicting power supply and demand, and for actuating the on-site power supply of the data centre.
- The **Supervision GCG** includes the supervisory intelligence which provides optimal IT power demand, power supply, and thermal policies to the individual sub-systems based on monitoring data, predicted systems states, and actuation feedback.
- The **Support Tools GCG** includes a number of tools that provide decision support for data centre planners, system integrators, and data centre operators.
- The **Integration Framework GCG** provides the communication infrastructure and data formats that are used for interactions between all components of the GENiC system.

Each GCG is composed of a number of functional components known as GENiC Components (GCs). The individual GCs are shown in Fig. 2. The core function of the GENiC system for continuous holistic data centre optimisation can be divided into four basic steps:

1. **Monitoring** components within the management GCGs collect data about IT workload, thermal environment, cooling systems, power demand and on-site power supply.
2. **Prediction** components within the management GCGs update their internal models and estimate future system states based on the collected monitoring data.

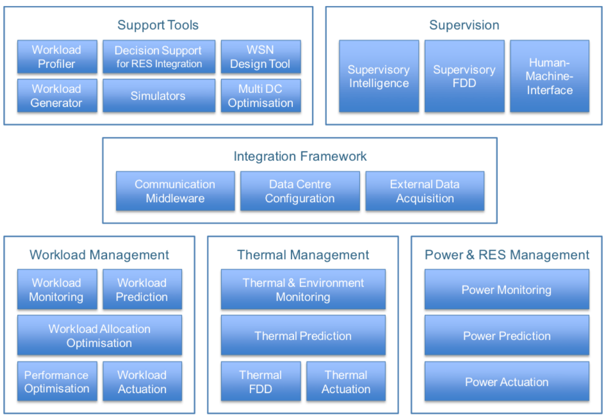

Fig. 2. Components of the GENiC functional architecture (from [9])

3. **Optimisation** components determine optimal policies based on the collected monitoring data and calculated prediction data. These policies are provided to the management GCGs.
4. **Actuation** components within the individual management GCGs implement the policies provided by the optimisation components in the data centre and at the renewable energy sources facilities.

In the following, we take a closer look at the GCGs and their individual components:

Workload Management GCG: The primary objective of this GCG is to allocate virtual machines (VMs) to physical machines (PMs) such that service level objectives (SLOs) are satisfied with low operational cost. Monitoring data from the IT resources deployed within the data centre is collected by the Workload Monitoring GC. The Workload Prediction GC uses this information to provide short- and long-term predictions about the resource utilization. The allocation and migration of VMs to PMs is determined by the Workload Allocation Optimisation GC, which solves a constrained optimisation problem, taking the predicted workload as well as constraints provided by the Supervisory Intelligence GC, Thermal Prediction and Performance Optimisation GC into consideration. The Performance Optimisation GC defines colocation and anti-colocation constraints for individual VMs and modifies the individual VMs' priorities to fulfil application specific SLOs. The VM allocation plan is finally applied by the Workload Actuation GC, which provides an interface to the data centre specific virtualization platform.

Thermal Management GCG: The Thermal & Environment Monitoring GC integrates monitoring of cooling systems and wireless sensor network infrastructure for collecting temperature and other environmental data in the data centre room. The collected data is used by the Thermal Prediction GC to provide short-term and long-term predictions to support supervisory control decisions, thermal actuation and workload allocation. Long-term predictions

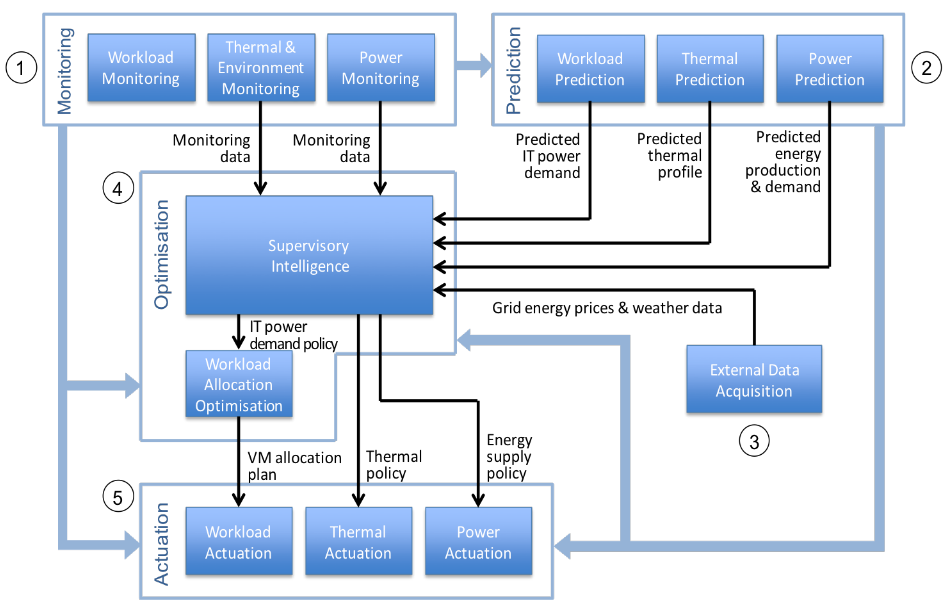

Fig. 3. Energy management use case (simplified) – from [9]

obtained with mathematical models are used for making decisions at the supervisory level. Short-term thermal predictions based on a discrete time mathematical model are required by the Thermal Actuation GC along with real-time sensor measurements to determine optimal set points for the cooling system in order to achieve the targets set by the Supervisory Intelligence GC. These short-term thermal predictions are also necessary input to the Workload Allocation GC, as they include temperature models for the thermal contribution of IT server workload to the server inlets, and the Supervisory Intelligence GC. Furthermore, short-term predictions, combined with equipment fault information from the Thermal Fault Detection & Diagnostics (FDD) GC, are used for fault detection and diagnostics at the supervisory level.

Power & RES Management GCG: The Power Monitoring GC integrates monitoring of the RES infrastructure for local energy generation and storage and of the data centre power consumption. This data is used by the Power Prediction GC to provide long-term predictions to support supervisory control decisions and power actuation. The Power Actuation GC determines set points for the power systems based on measured data, operational conditions, restrictions and limitations and the power profiles provided by the Supervisory Intelligence GC.

Supervision GCG: The Supervisory Intelligence GC is responsible for the overall coordination of workload, thermal, power management and heat recovery. It considers power demand and supply, grid energy price, energy storage model and determines how much power should be supplied from the electricity grid, RES and energy storage to minimize energy cost/maximize RES/minimize carbon emission accordingly over a given horizon. To this end, it provides policies for the actuation components in the Workload Management, Thermal Management, and Power & RES Management GCGs based on information from monitoring and prediction components. The Supervisory Intelligence GC provides these high-level policies to the Management GCGs for the purpose of guiding these component groups towards the Supervisory Intelligence GC strategy that has been chosen as a driver for current data centre operations; the key strategies available for selection are minimization of financial cost, minimization of carbon emissions and maximization of renewables. To detect and diagnose system anomalies, the Supervisory FDD GC compares predicted values with measurement data and collects and evaluates fault information. In appropriate situations, the Supervisory FDD GC informs the Supervisory Intelligence GC when a deviation becomes substantial enough to negatively impact system operation so that mitigation action can be taken by the platform until the fault has been corrected. The Human-Machine-Interface GC provides a framework for the user interfaces that allow data centre operators to monitor and evaluate aggregated data provided by the individual GCs.

Support Tools GCG: The GENiC platform includes a number of tools to assist data centre planners, system integrators and data centre operators:

- The Workload Profiler GC consists of a set of tools to capture application profiles that can be used by data centre operators to improve application performance.
- The Decision Support for RES Integration GC is a tool for data centre planners to determine the most cost-efficient renewable energy systems to install at a data centre facility.
- The Wireless Sensor Network (WSN) Design Tool GC is a tool to capture system and application level requirements for data centre wireless monitoring infrastructure deployments.
- The Workload Generator GC provides recorded and synthetic VM resource utilization traces for the simulation-based assessment of a GENiC based system and its implemented algorithms and policies.
- The Simulator GC supports the testing of individual and groups of GCs as well as the (virtual) commissioning of a GENiC platform before its deployment in an actual data centre.
- The Multi Data Centre (DC) Optimisation GC is a tool that exploits the differences in time-zones, energy tariff plans, outside temperatures, performances of geographically distributed data centres to allocate workload amongst them in order to minimise global energy cost and related metrics.

Integration Framework GCG: The Communication Middleware GC provides the communication infrastructure used within the GENiC platform. The Data Centre Configuration GC uses a centralized data repository to store all information related to the data centre configuration, including information on data centre layout, cooling equipment, monitoring infrastructure, IT equipment, and virtual machines running in the data centre. Finally, the External Data Acquisition GC provides access to data collected by existing components of the GENiC platform, including weather data, grid energy prices, and grid energy CO_2 indicators.

B. GENiC Communications Architecture

The GENiC platform integrates distributed software components, which are developed and maintained by individual consortium partners. A software component can implement a single GC, multiple GCs or just part of a GC to provide the required functionality to the platform. A topic-based publish-subscribe messaging architecture is a suitable mechanism to ensure a robust data exchange between individual software components. With this approach, the components do not need to be connected directly to each other. Components can publish messages to a central message broker using pre-defined topics and subscribe to the broker to topics from other components that are of interest to them. The broker forwards all incoming messages to the appropriate subscribers. The GENiC architecture defines a consistent interface specification using a common data format for all GENiC components. All interfaces are defined by hierarchically structured topics. Each of these topics has a defined message payload structure that uses the GENiC

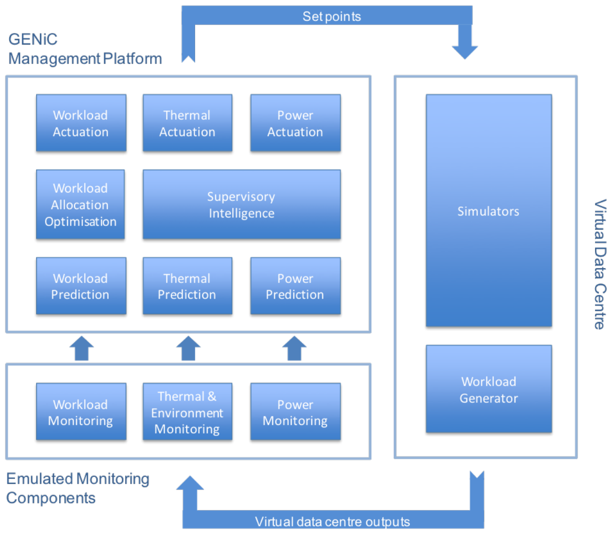

Fig. 4. Interplay between simulation and management platform (from [9])

common data exchange format which is specified based on JSON [10].

C. *Energy Management Use Case*

GENiC aims at optimally operating data centres with respect to energy related metrics. Integration of workload management, thermal management, and power management (including powering through renewable energy sources) is achieved via a hierarchical supervisory control concept. The three key optimisation criteria are (i) minimisation of total energy costs, (ii) maximisation of local RES power use, and (iii) minimisation of carbon emissions. To account for fluctuations in the IT workload demand and the availability of renewable energy supply (which includes local on-site energy production and grid power), the set points of the management sub-systems have to be adapted over time. The Supervisory Intelligence (SI) GC coordinates the individual management sub-systems, including renewable energy supply, by providing optimal policies with respect to the selected optimisation criterion. The use case scenario is illustrated in Fig. 3. Due to the limited space, we focus on the main interactions of SI with the other GENiC components. Interactions related to fault detection & diagnostics are not visualized in the figure. A more detailed description can be found in [9]. The basic operational flow is as follows:

Step 1 – The monitoring GCs, Workload Monitoring, Thermal & Environment Monitoring, and Power Monitoring, collect data from VMs, PMs, air conditioning equipment, sensor networks, power meters, and on-site energy supply systems. The relevant information is forwarded to the individual prediction and actuation GCs and SI.

Step 2 – Based on recent and historical monitoring data, the prediction GCs, Workload Prediction, Thermal Prediction, and Power Prediction, predict server power demand, thermal profile and cooling demand, RES production capacity and energy demand. The relevant information is forwarded to the individual actuation GCs and SI.

Step 3 – Additional data, i.e. weather data and grid energy prices, are obtained from external data sources and forwarded to SI by the External Data Acquisition GC.

Step 4 – Based on the inputs from the monitoring and prediction components and further interactions with the Power Prediction GC to validate the consequences of the power profiles SI is finalising on, SI provides a set of policies to the actuation GCs, Workload Allocation, Thermal Actuation and Power Actuation. The Workload Allocation Optimisation GC solves a constrained optimisation problem to determine an optimal VM allocation plan, taking the upper-bound IT power budget recommended by SI and additional inputs from other GCs (thermal and colocation and anti-colocation constraints) into consideration. The Thermal Actuation GC takes the minimum and maximum allowable data centre temperatures determined and then provided to it by SI and optimally calculates cooling equipment set points that ensure the room's thermal profile is properly regulated with minimal cooling equipment electrical power consumption.

Step 5 –Based on the inputs from SI and the Workload Allocation Optimisation GC, as well as monitoring and prediction components, the actuation GCs, Workload Actuation, Thermal Actuation, and Power Actuation, decide and apply the actual control actions. For example, the Workload Actuation GC executes the VM allocation plan and switches PMs on/off, based on the actuation requests. Faults are reported back to the optimisation GCs to be considered in the next iteration of the optimisation process.

D. *Data Centre Design Use Case*

The GENiC architecture does not only support the operational phase of data centre management but also provides components that support the design phase. A key element is the Simulators GC, which provides physics-based models for IT equipment, the DC whitespace environment, HVAC (cooling) systems, and power supply systems, and as such represents a virtual data centre. To emulate workload in the data centre, the Workload Generator GC is used to inject recorded or synthetic VM resource utilization traces into the system.

In simulation mode, the basic operational flow of the GENiC platform is the same as for the energy management use case presented in the previous section. This allow the calibration of algorithms within the GCs previous to the release of the GC to the main demonstrator. The main difference is that for the simulation, the monitoring components are emulated by the Simulators GC and that the set points determined by the actuation GCs are provided to the Simulators GC instead of the real system. This allows a data centre planner to test control strategies before implementing them in the real data centre. The interplay between simulation and management platform is illustrated in Fig. 4.

IV. Prototype Implementation and Evaluation

Fig. 5 illustrates a prototype implementation of the GENiC architecture presented in Section III. The GENiC distributed architecture approach with clearly defined interfaces simplifies integration of a diverse set of software components from multiple manufacturers and service providers. The architecture is scalable and flexible at the same time and is based on micro service architecture principles that offer the following benefits from a GENiC platform perspective:

- **Separation of concerns** – each service implements a single operational functionality. The architecture becomes more flexible and scalable at the same time.
- **Distributed security compliance** – each service can have different security policies allowing each service provider to maintain local security policies.
- **Freedom of service implementation** – each service provider can choose any development language without compromising the integrity of the overall platform. The only requirement is that the service needs to be able to communicate with the messaging broker.
- **Service scalability** – new instances of services can be spawned when more processing power is required.
- **Simplified API** – all modules use a common API to exchange data and trigger events used by other services.
- **Simplified testing and integration** – testing and integration is easier as testing focuses on black box testing with implementation details hidden behind APIs. Service integration hides APIs and dependencies.

A central element of the implementation of the GENiC integration framework is the use of the RabbitMQ messaging system for the GENiC exchange broker. RabbitMQ provides a range of client implementations in a wide range of programming languages, which avoids compromising the integrity of the overall platform. The individual components are implemented as individual services that communicate via

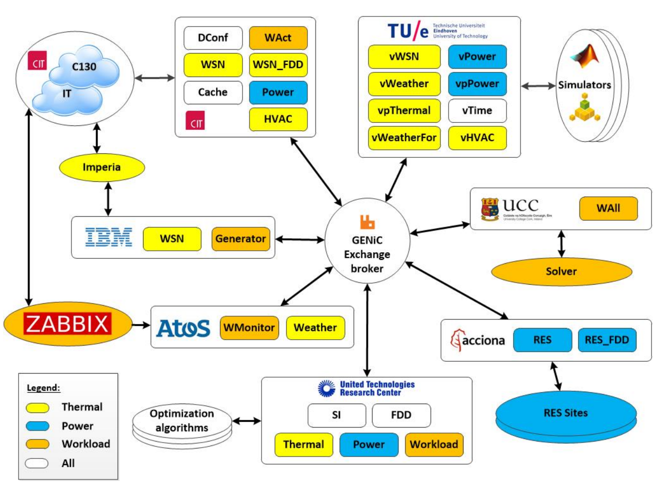

Fig. 5. GENiC architecture implementation prototype

the RabbitMQ message broker. A Generic Client architecture has been developed to allow each component provider expose their components in a distributed manner in the GENiC architecture. This client offers an easy way to integrate 3rd party (closed source) services with a minimal effort. Each of the components implemented in the GENiC prototype are shown in Fig. 5, colour coded based on the component group they belong to. Monitored data is stored in a NoSQL database backend. In the GENiC implementation we use CouchDB as a backend solution. Due to large number of stored data only short term data are available on the broker.

In order to assess the effectiveness of the GENiC architecture in terms of the energy efficiency, power management, managing increased penetration of renewable energy sources, heat reuse and data centre flexibility, the need to select appropriate metrics is of paramount importance. A cluster of FP7 data centre projects, including GENiC, DC4Cities, RenewIT, Dolfin, GEYSER and GreenDataNet funded under the FP7 ICT-2013.6.2 objective along with two previous FP7 projects, All4Green and CoolEmAll, together target data centre sustainability by increasing renewable powering, heat reuse and smart grid integration. This cluster has taken five common data centre metrics and defined twenty-one new metrics, along with measurement methodologies, to adequately capture the energy efficiency, flexibility and sustainability of modern data centres [11, 12, 13]. This approach supports the development of a common framework for monitoring and assessing the flexibility and sustainability of data centres. Each project in the cluster has selected a subset of those metrics that best measure its objectives and the metrics of specific interest to the GENiC project are listed in Table I. We are currently in the process and deploying the GENiC prototype implementation in a data centre on the Cork Institute of Technology Bishopstown campus in Cork, Ireland. We will execute the GENiC system for an extended period and gather statistics in order to evaluate our chosen metrics and to demonstrate the benefits of the GENiC system and approach.

V. CONCLUSIONS

The paper has presented an architecture specification to enable holistic, integrated energy management of data centres. The aim of the architecture is to provide data centre wide optimisation of energy consumption by integrating monitoring and control of computation, data storage, cooling, local power generation, energy storage, and waste heat recovery. The architecture has been design with interfaces and common data formats for effective component communication, control and optimisation functions and decision support to achieve a substantial reduction in energy consumption. The interfaces and distributed communications architecture enable a range of technology providers (manufacturers) to integrate their technology into an integrated system. This will enable data centres to have such an integrated solution in the absence of a single provider of such a system. Future development plans include improving data processing and automation support in the GENiC middleware, so data centre stakeholders will have the capability to get quick access to telemetry data and services. We are currently finalizing the integration phase of GENiC components and are at the first stage of system demonstration and evaluation of the GENiC approach at a real-world data centre.

Table I. GENiC evaluation metrics

Metric	Goal
PUE - Power Usage Effectiveness	Energy/Power Consumption
CER - Cooling Effectiveness Rate	
CUE Carbon Usage Effectiveness	
Energy Effectiveness of Cooling Mode in a Season	
ERE Energy Reuse Effectiveness	Energy Recovered/Heat Recovered
APCren - Adaptation of Data Centre to Available Renewable Energy	Data Centre Flexibility-Energy Shifting
DCA - Change in Data Centre Energy Profile from Baseline	
RenPercent - Share of Renewables in Data Centre Electricity Consumption	Renewables Integration
Renewable Energy Factor	
CO_2 Savings Change in Data Centre CO_2 Emissions From Baseline	Primary Energy Savings and CO_2 avoided emissions

ACKNOWLEDGMENTS

The authors acknowledge the contribution of the whole GENiC consortium We also acknowledge the financial contribution of the EC under framework programme contract no. 608826.

REFERENCES

[1] How Dirty is Your Data? A Look at the Energy Choices that Power Cloud Computing. Greenpeace Report, May 24th 2011. http://www.greenpeace.org/international/en/publications/reports/How-dirty-is-your-data/

[2] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's Not Easy Being Green," *Proc. ACM SIGCOMM*, Helsinki, Finland, 2012, pp. 211-222.

[3] A. Qurush, "Power-demand routing in massive geo-distributed systems", Ph.D. thesis, MIT, 2010

[4] W. Deng, F. Liu, H. Jin, B. Li, and D. Li, "Harnessing renewable energy in cloud datacenters: opportunities and challenges", IEEE Network, vol. 28, no. 1, Jan-Feb 2014

[5] T. Cioara, I. Anghel, M. Antal, S. Crisan, and I. Salomie, "Data center optimization methodology to maximize the usage of locally produced renewable energy," *Sustainable Internet and ICT for Sustainability (SustainIT*), Madrid, Spain, April 2015

[6] R. Das, S. Yarlanki, H. Hamann, J. O. Kephart, and V. Lopez, "A unified approach to coordinated energy-management in data centers", *Proceedings of the 7th International Conference on Network and Services Management* (CNSM '11), pp. 504-508, 2011

[7] T. Jiang, L. Yu, and Y. Cao, "Energy Management of Internet Data Centers in Smart Grid", Springer Verlag, 2015

[8] Project GENiC Deliverable D1.2 "GENiC Requirements & Architecture", http://www.projectgenic.eu/?page_id=57

[9] Project GENiC Deliverable D1.4 "Refined GENiC Architecture", to be published, http://www.projectgenic.eu/

[10] Internet Engineering Task Force (IETF), "RFC7159: The JavaScript Object Notation (JSON) Data Interchange Format," March 2014, https://tools.ietf.org/html/rfc7159

[11] A. I. Aravanis et al., "Metrics for Assessing Flexibility and Sustainability of Next Generation Data Centers", accepted for publication at IEEE International Workshop on Green Standardization and Industry Issues for ICT and Relevant Technologies in GlobeCom 2015, 6-10 December 2015.

[12] M. Rey Porto, S. Sanjoaquín Vives and A Fonellosa Caro, Cluster Activities Task 3, 2014, European Commission under the 7th Framework Programme

[13] Cluster Activities Task 4, "Data Centre Integration – Energy, Environmental, and Economic Efficency Metrics: Measurement and Verification Methodology", 2015, European Commission under the 7th Framework Program