



AALBORG UNIVERSITY
DENMARK

Aalborg Universitet

CLIMA 2016 - proceedings of the 12th REHVA World Congress

Heiselberg, Per Kvols

Publication date:
2016

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Heiselberg, P. K. (Ed.) (2016). CLIMA 2016 - proceedings of the 12th REHVA World Congress: volume 10. Aalborg: Aalborg University, Department of Civil Engineering.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- ? Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- ? You may not further distribute the material or use it for any profit-making activity or commercial gain
- ? You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Globally Optimised Energy Efficient Data Centres

Lara Lopez^{#1}, Enric Pages^{#2}, Dhanaraja Kasinathan^{*3}, Jacinta Townley^{*4}, Dirk Pesch^{#5},
Susan Rea^{#6}

[#]Atos Spain SA

Albarracin 25, 28037, Madrid, Spain

¹lara.lopez@atos.net

²enric.pages@atos.net

^{*}United Technologies Research Centre

4th floor, Penrose Wharf Business Centre, Penrose Wharf, Cork City, Ireland

³KasinaD@utrc.utc.com

⁴TownleJM@utrc.utc.com

[#]Cork Institute of Technology

Address Including Country Name

⁵Dirk.Pesch@cit.ie

⁶Susan.Rea@cit.ie

Abstract

Energy consumed by computation and cooling represents the greatest percentage of the average energy consumed in a data centre. As these two aspects are not always coordinated, energy consumption is not optimised. Also the penetration of energy coming from renewable sources has received limited interest due to the high investment costs needed to integrate them. Basically, data centres lack of an integrated system that jointly optimises and controls all the operations in order to reduce energy consumption and increase the usage of renewable sources. GENiC is addressing this topic through a novel scalable, integrated energy management and control platform for data centre wide optimisation. The proposed approach integrates monitoring and control of energy produced (local power generation and its sources, waste heat recovery) versus the energy consumed (computation, communication, data storage, cooling) inside a data centre. GENiC offers a distributed system architecture where the components within the platform exchange information through a message broker system. The Supervisory Intelligence component is located at the top of the hierarchical level control architecture and coordinates the primary data centre functions. Energy optimization and control algorithms are designed and implemented to achieve the high-level objectives. All these measures and considerations are represented in an integrated platform that provides a set of services to support energy optimisation and a decision support toolkit that targets a reduction of a 25% of the energy consumed plus the improvement of the operational performance (<1.3 PUE) with the target of achieving 80% renewable usage.

Keywords – Renewable Energy Sources; energy; optimization; Data Centre

1. Introduction

Data centers are a crucial part of the technology era. The rise of new technologies, such as cloud computing or high performance processing, the growth of the Internet of Everything, and the socialisation of technologies that are increasing the need of storage for the big amounts of produced data, directly relies on the capacity of data centres to cover computational and storage needs. But this also implies a direct cost related to energy consumption. On average, computing consumes 60% of total energy in data centres while cooling consumes 35%. New technologies can reduce a 40% the energy consumed on computation and cooling without taking into account coordination or optimisation. The integration of energy coming from renewable sources has received also limited interest from the data centre community.

GENiC (<http://projectgenic.eu/>), an European Commission funded project, has made a significant progress on the energy efficiency arena, at data centre level, by integrating in a single platform monitoring and control of the primary data centre energy components: computation, communication, data storage, cooling, local power generation, energy storage and waste heat recovery. The platform provides control and optimisation functions along with decision support tools that allows data centre operators to achieve a substantial reduction in energy consumption.

2. Architecture

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 and middleware platform is provided. 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. The proposed GENiC system architecture consists of six functional groups known as GENiC Component Groups (GCGs) [1] [2]:

- 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 system 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 Figure 1.




Fig. 1 Energy management use case (simplified) – from [2]

The core function of the GENiC system for continuous holistic data centre optimisation can be divided into four basic steps:

- **Monitoring** components within the management GCGs collect data about IT workload, thermal environment, cooling systems, power demand and on-site power supply.

- **Prediction** components within the management GCGs update their internal models and estimate future systems states based on the collected monitoring data.
- **Optimisation** components determine optimal policies based on the collected monitoring data and calculated prediction data. These policies are provided to the management GCGs.
- **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.

3. Data Centre Energy Efficiency Metrics

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, GreenDataNet, All4Green and CoolEmAll, 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 [3][4][5]. This approach supports the development of a common framework for monitoring and assessing the flexibility and sustainability of data centres. Each project has selected a subset of those metrics that best measure its objectives the ones selected by GENiC are listed in Table 1. GENiC is currently deploying its prototype implementation in a data centre located on the Cork Institute of Technology Bishopstown campus in Cork, Ireland.

Table 1. 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/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 |

GENiC system will be executed for an extended period to gather statistics in order to evaluate the metrics and to demonstrate the benefits of its system and approach.

4. Monitoring

In order to enable a holistic optimisation of the data centre energy consumption, the monitoring systems are key elements to guarantee that the information needed to optimize the workloads and the thermal distribution of the data centre is collected. The aforementioned monitoring components collect data with respect to IT workload (generated by both physical and virtual resources), thermal environment, cooling systems, power demand and power supply (considering renewables).




Fig. 2 Core functions of the GENiC system

The correct monitoring of each management group within the platform is mandatory to properly populate the prediction, optimisation and actuation components which operates the data centre. After collecting the monitoring information, it is published to other components through the communication middleware, which is the common interface for all the components within the platform. The proposed solution is not restricted to any monitoring system in particular, as each data centre owner can rely on their own system to capture this information, but it has to fulfil a minimum of required measurements, and provides them with the valid JSON format and periodicity in order to be able to fully populate other components of the platform through the communication middleware.

The reference implementation of the platform has been built using the following monitoring systems for each of the management groups:

- **Workload Management GCG:** The Atos Workload Monitoring system used is based on Zabbix, an open source enterprise-level software designed for real time monitoring of

millions of metrics collected from tens of thousands of servers, virtual machines and network devices.

- **Thermal Management GCG:** Temperature measurements are collected through the IBM IMPERIAL Wireless sensor network; sensors are deployed per one third of the rack across the data centre in order to provide thermal measurements. Environmental information is provided by Nimbus NICORE sensor network.
- **Power Management Group:** Data centre power consumption is provided by ResourceKraft system, deployed in CIT facilities. Additionally, a RES monitoring service provides monitoring information from the ACCIONA RES sites in Spain.

Additionally to the minimum set of metrics required to populate the GENiC software components within the platform, the monitory systems have also to provide enough intermediate measurements and metrics that enable the calculation of specific data centre energy efficiency metrics.

5. Middleware

The GENiC platform integrates distributed software components, which are developed and maintained by individual consortium partners, through a middleware platform. A software component can implement a single GC, multiple GCs or just part of a GC to provide the required functionality to the platform through the middleware. A topic-based publish-subscribe messaging architecture was implemented as a suitable mechanism to ensure a robust data exchange between individual software components through a messaging broker within the middleware. Components can publish messages to the 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. A consistent interface specification is defined 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 common data exchange format which is specified based on JSON [6].

Figure 3 illustrates a prototype implementation of the GENiC architecture presented in Section 2. The GENiC distributed architecture and middleware 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 and middleware 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 middleware 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 middleware 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 the RabbitMQ message broker.

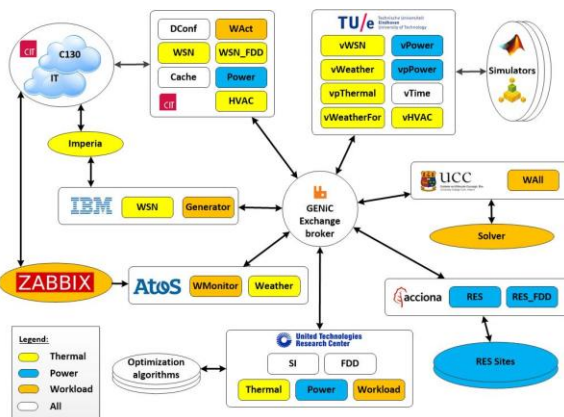


Fig. 3 GENiC architecture and middleware implementation prototype

A Generic Client architecture has been developed to allow each component provider expose their components in a distributed manner in the

GENiC system. 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 middleware prototype are shown in Figure 3, colour coded based on the component group they belong to.

6. Supervisory

As described in Section 2, the GENiC Supervisory component sits at the top-most layer of the hierarchical system-level control architecture of GENiC. This component is responsible for the optimal coordination between data centre cooling equipment, workload distribution controls, power generation and storage systems and, eventually, to maximise waste heat reutilisation. More specifically, this component makes high-level decisions that will impose optimal trade-offs between IT workload demand, safe operational DC thermal conditions and the available energy storage and power supply sources. In making such trade-offs, the Supervisory component allows for the data centre operator to also decide whether it is acceptable for IT workload to be delayed for the purpose of further improving performance. The performance of the data centre is driven by the strategy that best suits the data centre operator's needs and can be any one of the four following options:

- Minimisation of the total DC consumption,
- Minimisation of the total DC cost,
- Minimisation of the total DC carbon emissions or
- Maximisation of the usage of local renewable power supply.

It is not possible to always meet all of these objectives simultaneously. Hence, the data centre operator has the flexibility to choose the most important strategy from their operational standpoint. However, in solving one strategy, it is not unlikely that other one or two strategies are also solved.

After the data centre operator has selected one of the above options, the SI component takes in the necessary inputs from the GENiC platform to produce a 24-hour recommendation for how the system should be coordinated. This high-level recommendation is composed of three parts. Each part is as follows:

- **IT Power Demand Policy:** This is the optimum workload in kW per hour that the Workload Management group should aim to not exceed on a per hour basis.
- **Thermal Policy:** This is the minimum and maximum temperature range for the DC over each hour that the Thermal Management group must adhere to. It is important to note here that there is also an absolute minimum and maximum temperature range specified by the DC operator that must never be exceeded, e.g. a DC operator may wish to apply the ASHRAE [7] temperature recommendation of 17°C to 27°C for DCs in certain locations. The optimal temperature range

provided by the Supervisory component differs in that it is the ideal thermal envelope of the DC whitespace, based on the predicted IT power budget for that hour, and this envelope will always still fall within the absolute minimum and maximum thermal bounds specified by the DC operator; such an adherence would be required in any operational facility.

- **Energy Supply Policy:** This is the recommended power supply strategy that the Power & RES Management group should follow for powering the entire DC on an hourly basis – the demand of the entire DC is the power required to run the servers, the cooling system and the auxiliary equipment -. At the Supervisory level, power supply usage is grouped into two categories: Uncontrollable RES power supply usage profile, that are those renewables whose power generation is dependent on environmental conditions, e.g. wind and solar, and so cannot be turned on and off at will; and Controllable RES supply usage profile, that are renewables whose power generation capacity can be turned on and off at will, e.g. Organic Rankine Cycle (ORC), energy storage and grid energy. In addition to recommending power supply profiles, SI also considers when it is best to charge the battery for later usage and when to sell surplus energy back to the grid.

The Supervisory component jointly optimises all of these three policies simultaneously every hour for the next 24 hours, driven by the performance option selected by the data centre operator, as already explained above.

The three supervisory policies are then fed into the corresponding management groups and these groups bear the responsibility of reacting accordingly in order to meet the high-level policy goal.

The policies described above are generated using an optimisation algorithm which relies on prediction, monitoring and historical information from several components in workload, thermal and power management groups. For those components that rely on supervisory policy information for their next hour of operation only, the supervisory algorithm recalculates each of the three 24-hour policies every hour to produce the most accurate and up-to-date policies for a dynamically changing environment based on the latest prediction information available.

The rationality behind why the policies of supervisory component are designed this way is to keep the design scalable to any data centre with respect to:

- any choice of data centre size (small, medium, large-scale) in terms of the number and capacity of servers,
- any configuration of the cooling infrastructure (e.g. multiple chillers),

- any realistic number and choice of local uncontrollable RES power sources (e.g. solar panels, wind turbines, geo-thermal generator, or any combination) and
- any realistic number and choice of local controllable RES power sources (e.g. natural gas based CHP, biomass fueled ORC, or any combination).

7. Conclusions

The paper presents the reference architecture, implemented by GENiC project, that enables a holistic and integrated energy management of data centres. This implementation provides the optimisation of the energy consumed by a data centre integrating monitoring and control of computation, data storage, cooling, local power generation, energy storage and waste heat recovery. The architecture includes open interfaces, common data formats, control and optimisation functions, and a decision support toolkit to achieve a substantial reduction in energy consumption. It also allows the integration of renewable energy generation and energy storage equipment and operates it as a complete system to achieve an optimal energy and emissions outcome.

Acknowledgment

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] Project GENiC Deliverable D1.2 “GENiC Requirements & Architecture”, http://www.projectgenic.eu/?page_id=57
- [2] Project GENiC Deliverable D1.4 “Refined GENiC Architecture”, to be published, <http://www.projectgenic.eu/>
- [3] 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.
- [4] M. Rey Porto, S. Sanjoaquin Vives and A. Fonellosa Caro, Cluster Activities Task 3, 2014, European Commission under the 7th Framework Programme
- [5] Cluster Activities Task 4, “Data Centre Integration – Energy, Environmental and Economic Efficiency Metrics: Measurement and Verification Methodology”, 2015, European Commission under the 7th Framework Programme
- [6] Internet Engineering Task Force (IETF), “RFC7159: The JavaScript Object Notation (JSON) Data Interchange Format”, March 2014, <https://tools.ietf.org/html/rfc7159>
- [7] ASHRAE TC9.9, “Thermal Guidelines for Data Processing Environment - Expanded Data Center Classes and Usage Guidance”, 2011.