Project acronym: **CoolEmAll**

Project full title: **Platform for optimising the design and operation of modular configurable IT infrastructures and facilities with resource-efficient cooling**

SEVENTH FRAMEWORK PROGRAMME

Market Assessment Report (D7.4)

Author: Andrew Donoghue (451 Research)

Version: 1.0

Date: 30/04/2012

Deliverable Number:	D7.4
Contractual Date of Delivery:	30/04/2012
Actual Date of Delivery:	30/04/2012
Title of Deliverable:	**Market Assessment report (D7.4)**
Dissemination Level:	Public
WP contributing to the Deliverable:	WP 7
Author:	Andrew Donoghue (451G)
Co-Authors:	Andy Lawrence (451G) Wolfgang Christmann (Christmann) Eugen Volk (USTUTT-HLRS) Eser Ay (ATOS) Ariel Oleksiak (PSNC)

History			
Version	Date	Author	Comments
0.0	02/03/2012	Andrew Donoghue (451 Research)	Draft version
0.1	15/03/2012	Andrew Donoghue	2nd draft version
0.2	13/04/2012	Andrew Donoghue	3rd draft version
0.3	22/04/2012	Andrew Donoghue	4th draft version
1.0	30/04/2012	Andrew Donoghue	Final version

Approval		
Date	Name	Signature
30/04/2012	Ariel Oleksiak	

Abstract

This deliverable is intended to provide an overview of the main components and products that will result from project CoolEmAll and how they complement and compete with existing tools, metrics and standards already in existence.

The report is designed to reinforce the motivation for the project already detailed in the description of work and provides additional context and detail.

Each of the main components are analysed separately but analysis is also provided on how the project as a whole will add to the existing state of the art in the area of data centre energy efficiency.

A SWOT (Strengths, Weaknesses, Opportunities, Threats) analysis of each of the components is included as well as recommendations for the general product development and strategy based on an assessment of existing tools, products, metrics and standards.

The report includes input from consortium partners as well as external experts (from 451 Research and other relevant groups), and relevant documents and reports.

Keywords

Holistic, Facilities, IT, SVD Toolkit, ComputeBox, DEBB, Data Centre Infrastructure Management, metrics, modular, monitoring, sensors, low-power servers, convergence, high-performance.

Acronym	Meaning
CFD	Computational Fluid Dynamics
COVISE	COllaborative VIsualization and Simulation Environment
CPU	Central Processing Unit
CRAC	Computer Room Air Conditioner
DEBB	Datacenter Efficiency Building Block - defined in [D3.1]
DCIM	Data Centre Infrastructure Management
DVFS	Dynamic Voltage and Frequency Scaling
HPC	High Performance Computing
HTTP	Hypertext Transfer Protocol
MOP	Module Operation Platform
OS	Operating System
PSNC	Poznan Supercomputing and Networking Center
RAM	Random Access Memory
RECS	Resource Efficient Cluster Server
REST	Representational State Transfer
SVD Toolkit	Simulation, Visualisation, and Decision Support Toolkit
TIMaCS	Tools for Intelligent Management of very large Computing Systems
UC	Use Case
UI	User Interface
VM	Virtual Machine

Table of contents

Table of tables

Table of tables

1. Executive Summary

The purpose of this deliverable is to provide an initial market assessment and product analysis of the main components of the CoolEmAll project – namely the Simulation, Visualisation and Decision support Toolkit and Data centre Energy Building Blocks (ComputeBox) blueprints.

The deliverable deals with each component in turn and provides, where relevant, context on how each will complement, compete or add to the state of the art in that area.

The deliverable will also assess the potential impact of the project as a whole. The unique differentiator for the CoolEmAll project is that it tackles the issue of data centre energy efficiency from multiple angles in a holistic fashion. This kind of approach is particularly important when it comes to data centre energy efficiency as traditionally data centres are managed by two distinct groups – those concerned with the IT and those concerned with managing the mechanical and engineering components (facilities). CoolEmAll makes no real distinction between these two groups in data centre management and instead approaches the problem from a functional perspective i.e. what is the problem that needs to be solved and what tools are required to meet this objective?

The key findings from the market assessment can be summarized as follows:

- Rising energy prices, climate change, environmental legislation, and increasing demand for data services (driven by cloud computing) are putting pressure on data centre operators to improve data centre efficiency
- The adoption of energy efficient and sustainable approaches to data centre design and operation is increasing. However, lack of clear return on investment models and a split between IT and facilities in many large data centre operators is holding back this adoption.
- Data centres need to clearly understand how efficiently their facilities are being run before they can make significant efficiency improvements. This requires the deployment of monitoring tools (and related metrics) as well as new efficient IT hardware. All of these elements are included in the scope of CoolEmAll project (SVD Toolkit, DEBB blueprints, and new potential metrics)
- No other commercial platform or research project is approaching the issue of data centre energy efficiency in the same holistic way as the CoolEmAll project.
- There are existing technologies that solve some of the problems tackled by the individual components of the CoolEmAll project. Data centre Infrastructure Management tools, and some computational fluid dynamics tools, already provide some of the same functions as the proposed SVD Toolkit. Likewise some blade-server systems, modular data centres, and other integrated IT hardware platforms provide some of the same features of the DEBB blueprints (based on CoolEmAll partner Christmann hardware).
- The key differentiator and value of the CoolEmAll project is how each of its elements (SVD Toolkit, DEBB blueprints, workload management, and proposed metrics) are integrated together to provide a comprehensive suite of tools to enable data centre operators, and researchers, to better understand and manage the whole range of IT and mechanical components in their facilities.

2. Introduction

Project CoolEmAll is primarily focused on data centre energy efficiency – an issue that European authorities see as a key part of the Europe 2020 strategy to promote energy efficiency across the region over the next decade.

The CoolEmAll consortium plans to take a holistic approach to improving data centre energy efficiency that includes not just the role of IT hardware or facilities equipment but also the applications they ultimately support.

The expected results of the project include new data centre monitoring software, blueprints/designs of energy-efficient IT hardware, and potentially new energy-efficiency metrics. By coming at the problem of data centre efficiency from a number of angles, the CoolEmAll consortium hopes to develop solutions beyond the scope of other projects or industry efforts.

The main goal of the CoolEmAll project is to increase understanding of how a range of factors influences the overall efficiency of data centres. These include not only the power consumption of IT hardware and different approaches to cooling (e.g. mechanical chillers or free air cooling) but also the role of specific applications and workloads.

The output from tests and simulations carried out in the project will be used to develop the project's two main products:

- Simulation, Visualization and Decision support (SVD) toolkit – The consortium describes the SVD toolkit as an interactive data centre modelling and simulation tool. The toolkit will be developed using existing tools (such as COVISE and OpenFoam) as well as technology developed in the project. The SVD Toolkit will allow data centre planners to model the energy efficiency implications of physical placement of servers within the facility, different approaches to cooling, as well as the role played by applications and workload.
- Blueprints/designs of energy-efficient hardware – The other main outcome of the project will be a set of designs based on the server hardware used in the project. A set of detailed designs of how the equipment was used in the project will be released under an open source license. These designs, along with the SVD toolkit, should allow other projects, or potentially commercial data centre operators, to build on the research done by the CoolEmAll consortium.

3. Market analysis

3.1 Data centre demand

CoolEmAll is essentially concerned with improving the energy efficiency of data centres. The need to improve how efficiently data centre operate is increasing thanks to continued high demand for new data centre capacity combined with other factors such as the increased competition for energy resources. The financial crisis may have dampened data centre demand temporarily, but current projections indicate strong growth ahead. By 2020, it is estimated that annual investment in the construction of new data centres will rise to $50bn in the US, and $220bn worldwide[i].

According to a survey by the Uptime Institute, approximately 80% of data centre owners and operators have built or renovated data centre space in the past five years[ii]. Furthermore, 36% of survey respondents reported that they would run out of capacity by 2012, and a significant proportion of those indicated that they are planning to build a new facility in the near future.

Despite this overall trend toward expansion, large organizations are also increasingly aware of the need to improve the productivity of existing facilities and consolidate excess or inefficient capacity. The US government, for example, driven by the need to tackle its considerable deficit, is currently involved in an extensive consolidation program, which will see more than 800 of its 2,000-plus data centres shut by 2014.[iii]

3.1.1 Environmental cost of data centres

Some recent high-profile cloud-related data centre builds by the likes of Apple, Google and Facebook have increased awareness of both the financial and environmental costs of data centres. They have also raised awareness of the efficiencies that can be realized by smarter data centre design and operation (which the CoolEmAll consortium plans to add to). Environmental campaigners have begun to target carbon-intensive data centres with campaign tactics traditionally reserved for heavy industry. Environmentalists argue that data centre owners, and particularly cloud providers, could be more transparent about their energy use and efforts to improve efficiency. Greenpeace, for example, is concerned about whether providers are putting electricity costs above carbon intensity when constructing new facilities in locations with cheap, coal-based electricity generation.

Data centre operators have responded to this scrutiny with marketing initiatives, but also by developing some genuinely innovative sustainability initiatives, such as Facebook’s Open Compute strategy. Cloud providers also maintain that concentrating IT in large centralized facilities has enormous benefits in terms of how energy is consumed and managed. The emergence of more professionally run ‘information factories’ could ultimately mean fewer inefficient private facilities. Cloud service suppliers also point to their innovations around data centre energy efficiency.

Meanwhile, government regulators have also begun to take more notice of data centre energy use. For example, the UK’s CRC Energy Efficiency Scheme represented the first real ‘carbon tax’ to include data centres (See

Section 6. Legislative landscape). Other governments are expected to follow suit.

All of these factors are gradually shifting the traditional uptime-and-availability-focused view of data centre owners and operators to encompass energy efficiency and sustainability issues. This rise in energy awareness has helped drive the development of an increasing range of eco-efficient data centre strategies and technologies by traditional data centre suppliers, specialist start-ups and consultants (as well as EU projects such as CoolEmAll). Implemented in the correct way, these tools and approaches can help achieve financial savings, 'reputational savings' associated with environmental issues, and operational flexibility associated with freed-up energy.

3.1.2 Energy efficient data centre technologies

The technologies that are being adopted by data centres to tackle the growing issue of energy efficiency can be broadly categorized in two main groups: efficient IT and efficient facilities. CoolEmAll's holistic approach to data centre energy efficiency is founded on integrating both approaches but the split remains largely entrenched in industry both in terms of suppliers of data centre technologies and data centre owners and operators.

3.1.2.1 Energy efficient IT

Efficient IT refers to the optimization of the servers, storage and networking equipment within the facility. This includes approaches such as server virtualization and power management.

Technology or strategy	Impact	Maturity	Barriers to adoption	Adoption	Included in CoolEmAll
Virtualization/consolidation	High	High	Med	High	Yes
Efficient, multi-core servers	Med/High	High	Low	Med	Yes
Identify, eliminate unused equipment	Med/High	High	Med	Med	Yes
Use of cloud services to reduce energy, carbon	Med	Low	Med	Low	Yes
Storage management	Med/Low	Med	Med	Low	No

Technology or strategy	Impact	Maturity	Barriers to adoption	Adoption	Included in CoolEmAll
Power management/power capping	Med	Low	High	Low	No
Application monitoring and chargeback	Under review	Low	High	Low	Yes
Load shedding and balancing	Under review	Low	High	Low	Yes

Table 1 Overview of key energy efficient IT tools

As Table 1 illustrates, today there are really only three eco-efficient IT technologies that enjoy moderate or widespread adoption in the data centre (virtualization, multicore servers, and identifying unused equipment). A second wave of technologies – including power management and the use of cloud services – will undoubtedly make a bigger impact and enjoy wider adoption when they are more mature, and when preconceived ideas about outsourcing and server management undergo a rethink.

Specifically CoolEmAll will either develop or integrate with the majority of these technologies either through the SVD Toolkit (analysing impact of virtualization and consolidation, identifying underused equipment and load balancing) or the Data centre Efficiency Building Block (DEBB) blueprints (virtualization/consolidation and efficient multi-core servers) will contribute to the state of the art in these areas.

3.1.2.2 Energy efficient facilities

Efficient facilities technologies encompass the mechanical and electrical systems that control power, heating and cooling. Examples here include the use of free-air cooling and raising the operating temperature of data centres. Alternative approaches to data centre design, such as modular or containerized data centres, can also improve overall efficiency. All of these technologies and approaches have the proven or potential ability to slash energy use.

Technology or strategy	Impact	Maturity	Barriers to adoption	Adoption	Included in CoolEmAll
Enclosed hot/cold aisles	High	High	Low	High	Yes (partly)
Free or Fresh-air cooling	High	High	Med	High	Yes (partly)

Technology or strategy	Impact	Maturity	Barriers to adoption	Adoption	Included in CoolEmAll
Raise operating temperature	High	High	Low	High	Yes
Modularization	High	High	Med	High	Yes
Absorption Chillers	Low/med	Med	Med	Med	No
Containers	High	Med	High	Low	Yes
Relocate (green power, cooler)	High	Low	High	Low	No
Direct current distribution	Med/low	Low	High	Low	No
Liquid cooling	Med/Low	Med	High	Low	Yes
Thermal storage	Med/Low	Med	Med	Low	No
Heat recycling	Med	Low	High	Low	Yes
Fuel cells/local generation	Med	Low	High	Low	No
Dynamic smart cooling	Med	Low	High	Low	Yes

Table 2 Overview of energy efficient facilities tools

The high impact measures in Table 2 show what most data centre operators already know: they must allow their data centres to run warmer. They must manage the flow and mixing of warm and cold air. They should make full use of outside temperatures, usually indirectly, using economizers.

However, research shows that the majority of data centres are still slow to adopt these approaches. For example, over half of respondents to an Uptime Institute survey reported maintaining servers at or below 70°F (21°C), which is much colder than the maximum recommended server inlet temperature of 80.6°F (27°C), as stated by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE). Keeping servers colder than they need to be increases energy use.

CoolEmAll includes elements of, or will directly lead to the development of, a number of these leading facilities-side technologies. The ComputeBox blueprints and DEBB concept (applied as a main example to modular form-factor of the Christmann RECS technology) will be combined with technologies in the SVD Toolkit, which will allow data centre operators to analyse impact of approaches such as heat recycling, safely raising operating temperatures and dynamic cooling on energy efficiency of the data centre. Crucially, CoolEmAll will combine these advances in facility-side data centre energy efficiency with a number of similar advances in IT-side energy efficiency.

3.1.3 The role of CoolEmAll

The defining characteristic of the CoolEmAll project is that it will bridge this traditional gap between IT and facilities approaches to efficiency. The main outcomes of CoolEmAll - the SVD Toolkit and DEBB blueprints – will be based on a holistic rethinking of data centre efficiency that is crucially based on the interaction of all the factors involved rather just one set of technologies.

Some commercial suppliers (most notably Data centre Infrastructure Management suppliers) and consultants have recently begun to take a more all-encompassing approach to the problem by straddling both IT and facilities equipment. However, few suppliers or researchers up to now have attempted to include the crucial role of workloads and applications. That is beginning to change, and it is likely that projects such as CoolEmAll can advance the state of the art in this area.

3.2 Potential target markets

The potential target markets for the outcomes of CoolEmAll (SVD Toolkit, DEBB blueprints, new IT energy efficiency metrics) will mainly be medium to large data centres.

These data centres can be divided into four main groups: Enterprise data centres (operated and built by end-consumer of IT such as banks, pharmaceuticals, retailers), Commercial/multi-tenant data centres (operated by third-party companies to house IT owned by end-users of technology, Hosted or Cloud service providers (operated by third party companies to provide direct IT services to end-users such e.g. Cloud) and High Performance Computing facilities (Supercomputing sites owned and run largely by research organizations).

The groups described above all have different priorities when it comes to owning and operating data centres. However, the issues of energy efficiency and sustainability are common across all of these groups. Rising fuel prices, lack of capital for new data centre builds, expensive facilities technologies (mechanical chillers for example can represent up to 40 percent of the cost of building a new data centre), and emerging green legislation (such as the UK's CRC Energy Efficiency Scheme – a de facto carbon tax) mean that all classes of data centre are now focused on minimizing energy use.

But given the holistic approach, which is the basis of the CoolEmAll project, it is fair to assume that these different classes of data centre will favour certain aspects of the project over others. Other factors such as geographic location (ambient temperature, price of electricity) and vertical segment (retail, cloud services, banking) will also play a part.

The main results of CoolEmAll will consist of tools and good practises for modelling and analysis of data centres' energy efficiency. Therefore, the potential target markets will also include companies supporting the data centre classes described above, especially data centre designers, cooling equipment vendors, IT equipment providers. The CoolEmAll outcomes may be of a great importance for small and medium enterprises (SMEs) that deliver technology to data centres as they can get access to open source or affordable data centre analysis solutions. An example of such company is one of CoolEmAll project partners – Christmann, which will be able to use CoolEmAll tools to analyse energy efficiency of its solutions in a larger scale (hard to realize using real hardware).

3.2.1 Enterprise data centres

Enterprise data centres (those that provide IT services largely to one organization) make up a large proportion of the total number of facilities in existence. However, the rise of cloud, software as a service, rising energy costs, lack of capital, and green legislation are just some of the factors constraining new enterprise data centre builds. Enterprise data centre owners and operators are probably the most cost and efficiency conscious of the four groups defined above and as such may prove to be an important market for the products that result from CoolEmAll. The fact that these facilities are run by one organization may mean they are more open to experimenting with some of the emerging and disruptive technologies included in the project such as application-level monitoring and thermal and energy aware workload scheduling.

More importantly, the fact that the IT and facilities equipment is owned and managed by one organization in most enterprise facilities (as opposed to multi-tenant facilities where the IT is owned by the customer and the facilities by the third-party provider) means that these sites are more able to adopt the holistic approach to energy management which is central to the CoolEmAll project.

However, enterprise facilities' may also be capital constrained, and (like much of the data centre industry) highly conservative about implementing any new technology or strategy that could threaten the uptime of the IT services they provide. The CoolEmAll consortium will aim to understand these concerns and attempt to ensure that the products that result from the project are robust while still advancing the state of the art.

3.2.2 Co-location providers/multi-tenant data centres

Co-location and multi-tenant (MTDC) data centre operators are a class of data centre operators that essentially house IT equipment owned by a third party. These operators build large data centre facilities and lease space in them to organizations that can't afford, or don't want to go to the trouble, of building their own facilities.

The MTDC sector is very fragmented, with over 650 companies providing MTDC services globally and the top 25 global MTDC providers having less than 50% market share. Such fragmentation provides challenges to organizations looking to target this market (including project CoolEmAll).

However through 2012 there is expected to continue consolidation in the MTDC sector as providers look to add capacity and geographies and/or to solidify a market position within a particular geography.

The overall market demand for multi-tenant data centre space is expected to continue to exceed supply growth globally in each region of the world. It is predicted that growth in data centre demand will continue to outpace data centre supply through 2013. Enterprises seeking to minimize capex will continue to seek MTDC options to avoid building their own data centres. This competition for space means that MTDC operators will increasingly look to technologies that can help them improve operational efficiency.

Given that pressure, this class of data centre operator may find some interest in elements of the CoolEmAll project. For example, they can take advantage of partial project results such as application efficiency monitoring and modelling, impact of workload management policies on energy efficiency or metrics. However as the IT and facilities are essentially owned by separate companies, it is unlikely that many commercial data centre operators would be willing or able to take advantage of the holistic approach to data centre energy efficiency being explored by project CoolEmAll.

3.2.3 Hosted and cloud service providers

Data centres that provide direct computing services to other businesses have a unique set of requirements when it comes to the efficient operation of their facilities. Although they share the same control over IT and facilities infrastructure as enterprise data centres, they also have the commercial pressures associated with multi-tenant data centres (maintaining quality of service etc.). This class of data centres is also extremely important as it includes some of the largest names in IT (Google, Microsoft) and some of the largest facilities in terms of square meters and numbers of servers.

The combination of commercial pressure, scale and control means that cloud providers are amongst the most innovative adopters of new approaches to energy efficiency and sustainability. For example, Google is well known for customising servers and its wider data centre infrastructure to improve performance. The company recently disclosed that its facilities worldwide draw around 260MW of power, most of which is data centres. Google also estimated that its total carbon emissions for 2010 were just less than 1.5 million metric tons, and that an average search uses 0.3 Watt-hours of electricity.

Given the increasing demand for consumer and business cloud services, it is likely that hosted and cloud service providers will be receptive to any new technologies and research that can help them to improve how the efficiency of their operations. The CoolEmAll project plans to reach out to such organizations directly (potentially through the Project advisory board) but also via initiatives such as Facebook's Open Compute project which has some touch points (which are explored later in this report) with some aspects of the CoolEmAll – notably the DEBB and ComputeBox blueprints.

3.2.4 High Performance Computing (HPC)

High Performance Computing employs a range of tools, techniques and technologies to deliver additional performance for computationally intensive applications. Designers in aerospace, automotive and pharmaceutical markets; analysts in earth sciences, investment banking and oil exploration; and researchers in academia, defence and government are among those aiming to squeeze the last drop of performance from a system to reduce a time step or cell size – or to optimize a design – in order to bring a simulation another step closer to reality.

Historically, supercomputers that delivered HPC solutions were very different from mainstream computers, but that has all changed as the economics of deploying large clusters of low-cost commodity servers has encouraged many developers to re-architect their applications to scale effectively on clusters.

When supercomputers were expensive, proprietary architectures, the cost of powering and cooling them was not of major significance. But the three-year running costs of a commodity server-based HPC system now comfortably exceed the purchase costs. And as system components run faster, they get hotter, creating a cooling overhead that is not only expensive in terms of power consumption, but may be technically challenging in terms of heat dissipation. In one survey, 35% of supercomputer operators said heat and power had become their serious operational issue.

In terms of absolute power consumption, supercomputers may range from 20kW to nearly 13MW – larger than many medium-sized data centres.

The power consumption issues associated with high-end supercomputers present difficult technical issues associated with power availability and heat dissipation. The power is also expensive. For this reason, we expect HPC owners and operators to be the largest adopters of the technologies, which result from CoolEmAll.

The fact that two HPC research intuitions, PSNC and HLRS are involved in the project should also raise the profile of the project amongst HPC organizations and help to drive adoption.

3.3 Potential early adopters

The previous section outlines four classes of data centre operators. Trying to identify which group is most likely to adopt the results of CoolEmAll will be dealt with in future exploitation deliverables but it is useful to consider at this stage of the project the various technical and business imperatives at play in these different classes of data centre.

Given the holistic approach central to the CoolEmAll project, it is most likely that only those data centres with control over both IT and facilities equipment (enterprise, HPC and cloud service providers) will be able to make best use of the results of the CoolEmAll project. The consortium will consider the different pressures and motivations of these three groups as project continues.

It is also important to consider where the very first actual deployments of CoolEmAll technology will take place. The main components of the CoolEmAll project will be developed, and used for simulation and analysis of a number of data centres operated by consortium members (primarily ATOS, PSNC and HLRS). These applications will constitute the first actual adoption of the technology but should be seen as test and development environments. Specific validation scenarios that define specific use of CoolEmAll outcomes will be described in detail in D6.1.

In addition to simulation of specific parts of selected data centres, CoolEmAll results will be also applied to model behaviour of hardware delivered by an IT equipment provider – one of project partners: Christmann. Christmann aims at delivering an efficient computing hardware following the “data centre in a rack” vision. According to this, by filling a whole rack with RECS units, Christmann plans to deliver modules containing 600 computing nodes or more. Of course, constructing such a high-density environment is expensive and difficult to realize. Using CoolEmAll tools Christmann will be able analyse impact of various hardware settings, cooling, application characteristics and workload management on energy efficiency even for large infrastructures without a need to construct real prototypes and perform complex experiments.

However, the consortium will also look to identify a small group of companies to act as early adopters of the technology. This is most likely to be organizations that are already associated with the project e.g. the project advisory board. However, there are also some other potential early adopters including existing partners, or in the case of Christmann commercial customers, of consortium members.

These organizations are summarized in the following table.

Name	Type of organization	CoolEmAll Technology already deployed
Ostfalia High school	School	RECS | Box (See section 4.2 DEBB Blueprints)
IT-School in Kassel	School	RECS | Box
HLRS	Supercomputing centre	RECS | Box, Covise visualisation environment applied for air-flow visualisation.
PSNC	Supercomputing centre	RECS | Box (CoolEmAll Test-bed).
1E	IT power management specialist (advisory board member)	None
Racktivity	Smart Power Distribution specialist	None

Name	Type of organization	CoolEmAll Technology already deployed
Christmann	Energy-efficient IT equipment provider	RECS | Box

Table 3 Potential early adopters of CoolEmAll

4. Component positioning and description

There are a number of components that will be developed during the CoolEmAll project. The key components (SVD Toolkit, DEBB Blueprints, and energy efficiency metrics) are described in the following section and then analysed versus the existing market and state of the art in Section 5.

4.1 SVD Toolkit

One of the main results of the CoolEmAll project will be the development of a simulation, visualisation, and decision support toolkit (SVD Toolkit).

The SVD Toolkit is designed to aid data centre operators in planning the layout of new infrastructures or improving existing ones. CoolEmAll will develop a flexible simulation platform integrating models of applications, workload scheduling policies, hardware characteristics, cooling as well as air and thermal flows.

As the name suggests, the SVD toolkit will be made up of simulation, visualization and decision support components. Some of these will be developed internally in the project but the consortium will also look to make use of existing tools where possible (and financially viable).

Another related component, the Module Operating Platform (MOP), will also be developed alongside the SVD Toolkit. The MOP contains energy- and thermal-aware resource management and scheduling policies that can be inserted into simulations within the SVD Toolkit. The main aim is to be able to create a thermal map of the hardware in the testbeds. This will be used to monitor the effect of changes in IT hardware performance, cooling technologies (free cooling versus mechanical chillers) and different applications and workloads on total energy use. However it is not envisioned that it will be a product per se but rather a sub-component of the SVD Toolkit a tool to help with internal development during the project. In particular, it will enable to preform experiments in a real environment (mostly CoolEmAll testbed) and compare results with those coming from simulation tools. In this way, MOP will support process of model creation and validation of SVD Toolkit results.

It is likely that the SVD Toolkit will be developed from elements of the open source CFD application OpenFoam and a visualization tool developed by HLRS – The Collaborative Visualization and Simulation Environment (COVISE). However a wide range of such tools exists and will be evaluated through the course of the project.

The project will also investigate how workload scheduling and resource management can be used to improve data centre efficiency. A set of policies will be developed on how, for example, (virtualised) applications and workloads can be moved between servers to improve utilization and avoid hotspots. The project will also investigate techniques such as dynamic voltage and frequency scaling to reduce power usage of servers during periods of inactivity.

The SVD Toolkit will essentially allow data centre planners to model the energy efficiency implications of physical placement of servers within the facility, different approaches to cooling, as well as the role played by applications and workload optimisation.

As stated, a full description of the state of the art in simulation can be found in deliverable D2.1 Report about state-of-the-art simulation and visualisation environments, however the key findings are summarised below.

Element	Tool	Description	Suitability for the SVD Toolkit
CFD	OpenFOAM	• Maintained and released as open source under the GNU General Public License. • Generally used for CFD, but can also be applied to areas like stress analysis, electromagnetics and finance	• Free of charge and publicly available. Commercial support is supplied by OpenCFD • Easy to integrate an OpenFOAM based simulation into emerging frameworks
	ANSYS CFD (CFX and FLUENT)	• Computer Aided Engineering (CAE) package • Products share a common workflow, such as CAD connectivity, geometry modelling, meshing, and post-processing.	• Well known, powerful and trusted products worldwide • BUT a commercial product, available only on licenses acquisition
	COMSOL Multiphysics	• Software suitable for the simulation of various physics phenomena	• User friendly templates for coupled heat transfer and fluid dynamics problems • BUT each module requires a paid license that may include maintenance and support, resulting in high

Element	Tool	Description	Suitability for the SVD Toolkit
			costs.
	Star-CCM+	• CFD suite developed by CD-adapco	• High quality, high performance fluid simulation code that would be very suitable for the CoolEmAll SVD. • BUT The high price tag makes it difficult to acquire, though, especially for the casual end user of the CoolEmAll SVD toolkit
	APC Data Centre Software	• Proprietary suite solution developed by Schneider Electric that allows companies to design and optimise all the aspects of their data centres	• Flexible, but the data centre owner has to buy licenses for each of the components • Costs make it most likely unsuitable for integration into the SVD toolkit.
	6SigmaDC	• Commercial simulation code developed by the company "future facilities" that is dedicated to the simulation of data centres	• The most comprehensive and feature complete of all the dedicated solutions • Recommended depending on price
	TileFlow	• Constructs a computer model of the data centre and uses CFD to assess airflow patterns and temperature distribution	• Intuitive to use and has a great user interface • BUT the absence of an external API makes it difficult to attach it to other frameworks
Element	Tool	Description	Suitability for the

Element	Tool	Description	Suitability for the SVD Toolkit
			SVD Toolkit
Simulation	GreenCloud	• Workload simulator for the evaluation of energy use in data centres	• Enables energy-aware scheduling • BUT does not allow for testing of the impact of a virtualisation-based approach on resource management
	GridSim and CloudSim	• Allow for the simulation of Grid and Cloud computing systems	• Both tools are commonly adopted and enhanced by other distributed computing simulators and are available without any cost issues • BUT may be integration issues
	GSSIM	• Highly customisable and extendable framework that offers a comprehensive simulation environment for workloads in distributed computing systems	• A standard format for the input files and the simulation results, GSSIM can be easily integrated with other SVD Toolkit components • As it is developed by PSNC, full access to the source code is available and it will be easy to extend GSSIM if needed
Visualisation	Paraview	• Open source data analysis and visualisation platform • Offers a complete set of functionality for all major visualisation tasks	• Sophisticated Open Source tool that is scalable on large clusters • As ParaView is Open Source, modifications needed for inclusion into the SVD as a visualisation platform would be possible

Element	Tool	Description	Suitability for the SVD Toolkit
	EnSight	• Visualisation package created by CEI Software	• One of the standard visualisation tools commonly used in commercial settings • The high price tag limits the suitability for CoolEmAll, especially during the development phase
	OpenDX (Open Data Explorer)	• Modular visualisation package originally created by IBM and released to the public using the open source IBM Public License (IPL).	• Although it is stable and mature, development on OpenDX is discontinued • Use of OpenDX in the CoolEmAll project is not recommended
	COVISE (Collaborative VIsualization and Simulation Environment)	• Mature tool for visualisation that also includes hooks for geometry and grid generation as well as linking directly to various simulation codes	• COVISE has proven to tackle the airflow computation of a data centre and its modular approach makes it easy to include new simulation codes in the workflow • As it is developed by HLRS, full access to the source code is available and it will be easy to extend COVISE if needed features should be amiss
	Vitrall	• New, distributed web based visualisation system designed to utilise efficiently multi-GPU	• The Vitrall core component is free and Open Source, which makes it an interesting complementary option to other tools (such as COVISE)

Element	Tool	Description	Suitability for the SVD Toolkit
		server installations	• Vitrall is developed by PSNC.

Table 4 Potential components for SVD Toolkit

4.2 DEBB Blueprints

The other main outcome of the project will be a set of hardware designs. These designs/blueprints are defined by the project as Data Centre Efficiency Building Blocks (DEBBs) and can be divided into a number of granularity levels.

Level	Full name	Description
1	Node Unit	This level reflects the finest granularity of building blocks to be modelled within CoolEmAll. This smallest unit reflects a single blade CPU module, a so-called "pizza box", or a RECS CPU module
2	Node Group	This level reflects an assembled unit of building blocks of level 1, e.g. a complete blade centre or a complete RECS unit (currently consisting of 18 node-units)
3	ComputeBox1	This level reflects a typical rack within an IT service centre
4	ComputeBox2	Building blocks of this level are assembled of units of level 3, e.g. reflecting a container filled with racks or even complete compute rooms.

Table 5 DEBB Definition

A DEBB is a concept that will allow definition of various data centre building blocks. These will be published by the project as open designs and used in simulations. A DEBB of each granularity level is described by several main elements such as composition (specification of components and sub-building blocks), physical dimensions (outer physical dimensions, and optionally arrangements of components and sub-building blocks within particular DEBB), power consumption for different load levels, thermodynamic modelling profile (describing air-flow and temperature on inlets and outlets for different load levels), and metrics assessing energy efficiency of particular DEBB. The architecture of these various classes of DEBB unit is explored in more detail in deliverable D3.1 First definition of the flexible rack-level ComputeBox with integrated cooling.

For the purpose of expediency in this market assessment it makes sense to consider DEBB units also in terms of the actual hardware they will define in the first order – namely the RECS| Box technology from Christmann. Christmann describes its RECS| Box as essentially a 'data centre in a rack.' It consists of multiple high-density server units (equivalent to the Node Groups described in Table 4) combined with storage, including de-duplication, placed into a single rack along with integrated cooling.

Since each high-density server unit contains up to 18 CPU-Modules, the Compute Box rack can contain up around 600 individual server boards, which Christmann says could make it equivalent to a medium-sized data centre.

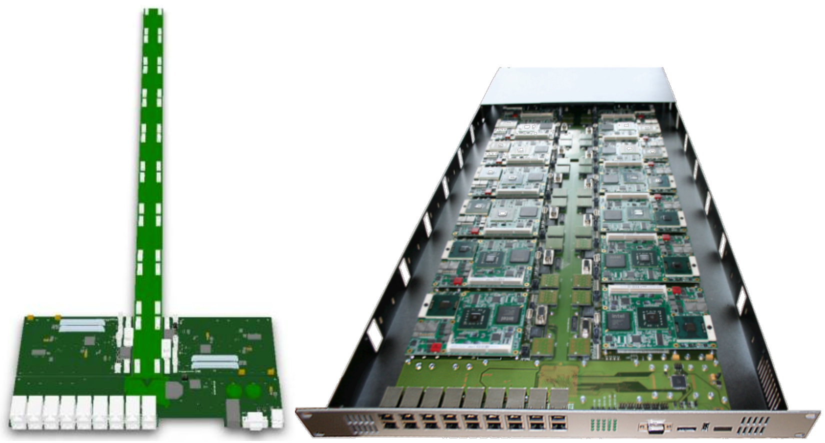

Figure 1 The RECs System (picture and technical sketch)

Along with the18 CPU boards, each server unit contains 18 thermal and 18 current sensors. Each CPU board can be plugged into a ComExpress socket and can have different performance, from a low-power ARM or Intel Atom up to an Intel i7 (Sandy Bridge) CPU with 16GB of RAM. Sensors integrated in the baseboard infrastructure enable the network status, fan speed and power usage of each of the 18 boards to be monitored. The sensor information from each board is collected by a microcontroller, which then relays the information to a master microcontroller for all 18 nodes that can be accessed via an Ethernet network.

This system of sensors and microcontrollers means each server node can be powered up and down as required, saving large amounts of power. (Christmann is developing this ability to monitor and dynamically power down nodes as part of its involvement in the EC-funded Green Active Management of Energy in IT Service Centres, or GAMES (see Section 5.4.1)

The company has also developed a rack-based power unit with high efficiency (93%) and DC power distribution, which will sit within the Compute Box. Christmann estimates a ratio of one power unit to every three to five high-density server units within the Compute Box rack, which depends on the used CPU modules.

Cooling – The extremely high density of the Compute Box makes cooling especially important. The individual server units are cooled by cold air, which is provided by a cold-water heat-exchanger system integrated into the rack. Christmann says it can adopt a variety of rack-level, cold-water systems. The final RECS| Compute Box will be delivered with an integrated, efficient cooling system that is under development with another German rack and cooling supplier.

Monitoring and management – Aside from the use of energy-efficient components and liquid cooling, Christmann also states that manageability is an important part of improving energy efficiency. It has developed a network monitoring and controlling system, which does not impair the performance of the systems under inspection. More details can be found in deliverable D3.1.

Processors – Christmann provides a choice of processors for its hardware, including the low-power Intel Atom, as well as AMD Fusion processors or even very low-power chips from ARM Ltd. Overall power use is reduced by making the form factor of the devices as compact as possible. For example, Christmann uses Mini-ITX and even smaller motherboards, which are extremely compact and require only about 15W of power (idle state), according to Christmann. The boards can also be passively cooled without the use of a CPU fan.

Since it consists of numerous skinless servers integrated into a rack with built-in cooling, the Christmann RECS| Box could be considered a form of modular or containerized data centre (again, Christmann describes it as a data centre in rack). As well as examining the thermal output at the server and rack level, the project also plans to approach the Christmann RECS| Box as a modular data centre. Versions of the system with built-in liquid cooling and external air-cooling will be examined. It is possible that the project will also investigate the thermal implications of having a group of high-density RECS| Box in close proximity.

4.3 New metrics (and standardization)

Where relevant, the consortium will look to collaborate with existing standards bodies and other organizations in the development of existing and new metrics. This will help ensure that the technology that results from the project is useful to the wider data centre industry, both in terms of research but also in commercial adoption by data centre technology suppliers, and data centre owners and operators.

The state of the art in energy efficiency metrics will be explored more fully in subsequent deliverables (notably Deliverable D5.1 White paper on Energy- and Heat-aware metrics for computing modules). However for the purposes of this market assessment it is possible to outline the main metrics currently in existence and where CoolEmAll could potentially add to the existing state of the art of look to develop new metrics, potentially in collaboration with other organizations and standards groups (see section 5.3).

5. Competitive environment

The CoolEmAll project is an attempt to develop an innovative solution to cooling and energy efficiency issues in data centres. However, this is an area that has seen high levels of innovation over the last five years driven by rising energy prices and climate change. It is important therefore that the SVD Toolkit, DEBB Blueprints and metrics components of the CoolEmAll project are developed with a thorough understanding of existing technologies and development efforts. The following section continues the description and analysis of the products outlined in the previous section but also describes the competitive environment into which these technologies will be deployed. In some instances the results of the project will compete with existing technologies but efforts will also be made to integrate with, and add to, existing commercial products, standards and research products (including complementary European projects).

5.1 Alternative approaches to SVD Toolkit

Some open source and commercial alternatives already exist to some aspects of the SVD Toolkit however currently there is no one offering which covers all the functions that the consortium plans to include. Some of the main technologies/products and the vendors that provide them are outlined below together with how CoolEmAll compliments or advances on the functions they provide.

5.1.1 Basic online energy calculation tools

A range of tools already exists which can be used to help data centre owners and operators capture energy use information (for use in calculating data centre Power Usage Effectiveness PUE for example). These include DC-Pro from the US EPA [iv] or the Data Centre Efficiency calculator from APC (Schneider Electric) [v] or the BCS Data centre Strategy Group/Carbon Trust Data centre Simulator[vi]. These tools are limited in scope however and do not address the application or workload layer contribution in the granular way in which it is intended the CoolEmAll SVD Toolkit will.

5.1.2 Computational Fluid Dynamics

Computational fluid dynamics (CFD) creates a detailed mathematical model of airflows, temperatures and other environmental variables within a space. CoolEmAll is going to apply CFD for the airflow and thermal modelling planned for the SVD Toolkit. To this end the use of OpenFoam is planned. However, project goals include integration of CFD simulations with advanced visualisation tools as well as simulation of applications and workload management policies. The important project objective is also definition of DEBBs – predefined open source models that can be re-used in simulations of data centres.

Furthermore most CFD modelling applications are not interactive tools (although some suppliers are trying to develop these). Many suppliers are still questioning how to use them with dynamic, virtualized environments in which IT heat output changes with varying workload. One of the aims of the SVD

Toolkit is to tackle this issue of interactive data centre energy efficiency analysis and how thermal characteristics change with workload.

Suppliers of commercial CFD tools include Applied Math Modelling (CoolSim), Future Facilities (6SigmaDC), Innovative Research Inc (TileFlow), and Daat Research (CoolitDC).

Suppliers such as Future Facilities are also trying to develop a more holistic approach to CFD (extra modules that allow integration with computer-aided design tools and other third-party data centre software, including that provided by nlyte Software and RF Code for example) that may result in platforms that are close to some of the aims for the SVD Toolkit. But for now most of these efforts are still relatively immature.

5.1.3 Data centre infrastructure management (DCIM)

The emerging, and important, data centre infrastructure management sector (DCIM) has some broad similarities with the holistic approach that underpins the CoolEmAll project. DCIM is difficult to define precisely. It is multifunctional, has many components, attempts to address various technical and business issues, and may consist of numerous subsystems that appear to duplicate or overlap with other systems. DCIM also goes by many names: data centre efficiency software, operational technology, data centre management (DCM) and data centre operational management (DCOM). These terms are sometimes proprietary, inconsistently used, or not widely recognized.

But despite some confusion about what constitutes a DCIM system, it is a rapidly expanding area of interest for suppliers and data centre owners and operators. By adding to the state of the art in DCIM, CoolEmAll will be contributing to an important and fast growing technology sector. The DCIM market was worth $245m in annual revenue in 2010, and will grow to $1.27bn in 2015, a growth rate of 39% per year, according to a study from 451 Research. This growth figure far outstrips that for the data centre equipment industries and for the enterprise IT segment as a whole.

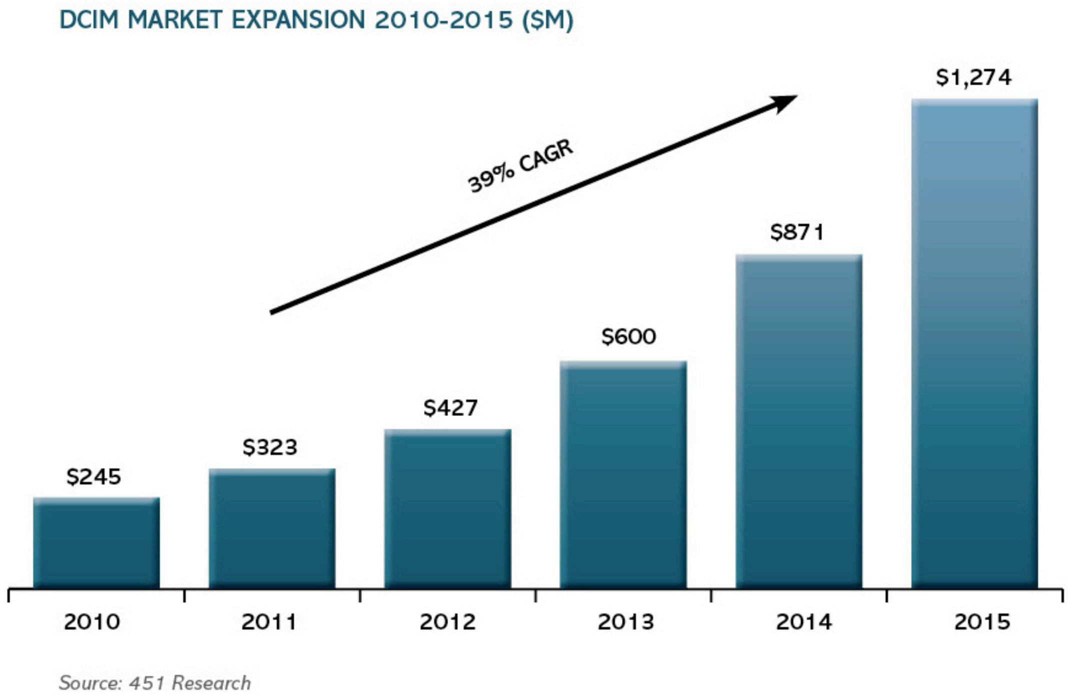

Figure 2 Projected growth in the DCIM market

At its core, DCIM is an attempt to use information from multiple sources within the data centre (IT and facilities) to improve the general management and sustainability of data centres. A data centre infrastructure management system can basically be defined as a system that collects and manages information about a data centre's assets, resource use and operational status. This information is then distributed, integrated, analysed and applied in ways that help managers meet business and service-oriented goals and optimize the data centre's performance.

A DCIM system may be defined as having several (few have all) of the following functions (the DCIM stack)

1. Data collection (meters, sensors)
2. Environmental monitoring and reporting
3. Asset, configuration and change management
4. Power/energy measurements and modelling
5. Power management, power scaling, and power capping
6. Data management, integration and reporting
7. Capacity planning, forecasting, simulation and analytics
8. Optimization, operational BI, load management

In practice, DCIM systems may vary widely in focus, and complete DCIM offerings are likely to consist of a framework or suite of products, from one or many suppliers, that are designed to interoperate with or complement each other.

Given the wide-variety of DCIM systems available, it is very difficult to assess how the relevant aspects of the CoolEmAll project (namely the SVD Toolkit and Module Operating Platform (MOP)) will complement or improve on existing DCIM technology. However, Figure 3 attempts to overlay some of the functions that the SVD Toolkit will provide against functions already included in some DCIM platforms. These improvements are two-fold. On one hand project results such as RECS and MOP will deliver technology for effective monitoring and managing hardware and software. On the other hand SVD Toolkit will enable simulation and evaluation of various settings, management policies, etc. that should be applied by DCIM tools.

Capacity planning, forecasting, simulation, analytics			Optimization, operational BI, load management		
Data management integration and reporting					
Cooling control, BMS, alarms etc.	Environmental monitoring and reporting	Asset configuration and charge management	Power, energy measuring, modelling	Power management, power capping	IT service and systems, VM mgt
Data collection, meters, sensors					
Application and workload (real-time) monitoring and management					
Thermal and energy-aware resource management					

Figure 3 DCIM stack versus CoolEmAll (Blue = Existing features of DCIM platforms, Green = Features found in some DCIM Platforms but planned for CoolEmAll, Red = Features not in DCIM platforms but planned for CoolEmAll)

5.1.4 Thermal and energy aware workload scheduling

We can broadly say that the SVD Toolkit should be seen as operating one level higher than most DCIM systems especially given its focus on the impact of workload and application management, which is not currently a feature built into most DCIM systems. SVD Toolkit will allow evaluating this impact in various settings in order to decide about proper configuration of DCIM tools.

The inclusion of thermal and energy-aware resource management, moving virtual workloads between servers or even between data centres for energy efficiency reasons, is also another area where the SVD toolkit will advance beyond the scope of most mainstream DCIM systems. Tools such as VMware vSphere or Ovirt or Platform VM Orchestrator enable VMs to be managed according to set policies but some of these suppliers have integrated their technology with existing DCIM tools. However, the level of integration between traditional DCIM functions and virtual machine management is still relatively immature and CoolEmAll should be well placed to advance the state of the art in this area.

5.1.5. Dynamic Voltage and Frequency Scaling

Another feature, which CoolEmAll plans to investigate, is the ability to enable nodes (servers) to be switched off, or put into lower power and frequency states to improve energy efficiency. Such tools already exist from suppliers such as MiserWare and EnergyWare. MiserWare for example tackles wasted power consumption by servers and desktops. It does this by putting the processors into lower power states while maintaining the machine in an active mode, a process known as dynamic voltage and frequency scaling or CPU throttling. MiserWare calls its approach Intelligent Software Power Management.

One important differentiator MiserWare claims for its technology is that it can be used to reduce server energy usage in a predictable way. Turning chips off completely or putting them into very low power states, or P-states, can disrupt applications or workloads supported by the server under power management. Granola, however, only slows down the chip, so the impact on performance is lessened and more predictable, MiserWare says. The predictability can be used to establish energy usage SLAs between MiserWare and its customers – a vital issue in the data centre where uptime is everything.

SVD Toolkit will enable performing experiments with DFVS, including it impact on thermal processes in a data centre as well as application-specific behaviour. In this way, CoolEmAll will attempt to advance DVFS methods with new thermal and energy aware workload scheduling policies taking advantage of DVFS technologies.

5.1.6 Data centre predictive modelling (DCPM)

The SVD Toolkit may also overlap with technology separate but closely aligned to DCIM called DCPM (data centre predictive modelling).

DCPM tools, such as Romonet's Prognose[vii] and Lumina Decision Systems' Analytical Data centre Capacity Planning tool,[viii] can be used for detailed data centre planning. Prognose for example allows a user to create a detailed model of a facility and then run a mathematical simulation to predict energy and cost performance. By changing the model and re-running the simulation, users can experiment with different data centre designs or operational strategies. A data centre designer might use the tool to compare the predicted performance of different cooling technologies, and to see how the answer might change in different climate zones. A data centre operator with an active site might use the model to estimate cost and energy savings from a change in temperature set point or an efficiency retrofit.

The BCS and Carbon Trust developed an open source DCPM tool, called the Data centre Simulator. (Romonet's Prognose tool was originally based on the Data centre Simulator but has been significantly updated recently).

SVD Toolkit will provide similar functionality to mentioned DCPM tools with a greater focus on detailed air-low and heat transfer processes in a data centre as well as impact of specific classes of applications and workload management policies. CoolEmAll's approach also assumes the use of open DEBB blueprints, which should facilitate exchange of models between data centre planners and hardware providers. Nevertheless, CoolEmAll needs to have a closer look at DCPM tools (especially open source such as Data centre Simulator) in order to minimize overlaps and maximize possibilities of models re-use.

5.1.7 Other data centre monitoring and management tools

Other data centre monitoring and management tools that should be considered by CoolEmAll include those developed by IT suppliers – most notable Intel but also server suppliers including HP and Dell. These tools include functions for monitoring and measuring power, temperature but also some degree of control including power capping.

Intel's Data centre Manager (DCM) tool is particularly relevant to the project goals. DCM was launched in 2009. Along with Intel Intelligent Power Node Manager (IPNM), it provides power and thermal monitoring and management for servers, racks and groups of servers in data centres. The products provide information on board-level temperatures and power use, as well as the capability to control processor power consumption by altering voltages and frequency. IPNM provides the direct monitoring and control of the processor power states (p-states) while the DCM SDK acts as an interface to IPNM that can be integrated into other consoles or management platforms. (Intel lists a number of supported servers for DCM from IBM, HP and Dell, but only those that are IPNM-enabled can be power-managed.)

The real standout benefit of DCM is that it is a cross-platform power management tool. The large server makers have also developed some monitoring tools such as IBM's Systems Director Active Energy Manager, Dell's Open Manage and HP Insight Manager, but crucially, these are proprietary. At present, power capping is only possible across one set of products.

Intel has licensed DCM to a number of data centre technology suppliers with a number of deals done in the second half of 2011. The majority of these are DCIM suppliers and include iTRACS, FieldView Solutions, Modius, Viridity Software, Rackwise and Optimum Path Systems. Most have some insight into the IT layer, but this is usually limited to monitoring rather than actual control and management. Licensing DCM provides an efficient way to add IT energy management functions such as power capping and power throttling, as well as more sophisticated monitoring (DCM is capable of real-time monitoring of actual power and inlet temp data). The reason a number of such deals were made in the last half of 2011 is not clear, but may simply be due to the fact that once one supplier licensed the technology (Modius was the first), others were forced to follow in this increasingly competitive market.

Other licensees of DCM include data centre energy management suppliers such as JouleX and Power Assure (also considered a DCIM supplier). These suppliers have more sophisticated IT energy management tools that are usually a mixture of proprietary technology and standard protocols. DCM adds to these by providing access to processor power and throttle states. Intel has also licensed DCM directly to a small number of data centre operators with in-house management platforms.

The capabilities of Intel DCM can be broadly split into two areas: monitoring and control. A number of suppliers currently offer products capable of monitoring server energy use, mostly via PDUs and power strips that can report power data. In PDUs, for example, Intel DCM competes with smart-PDU maker Racktivity as well as established suppliers such as Servertech, Schneider Electric (APC), Raritan, Emerson Network Power (including its Avocent division) and Eaton. There are many smaller suppliers, too.

Control is a more challenging area. Intel dominates the market for an x86 processor in the data centre and nearest rival AMD has no real competitive product to DCM. Virginia-based start-up MiserWare has a product called Granola Data centre that uses dynamic voltage and frequency scaling or CPU throttling to directly control server processors. Another Virginia-based start-up, EnergyWare, has similar capabilities. UK-based 1E's NightWatchman Server product also has the ability to monitor server energy use and put machines into low power states or even shut them down. VMware has some capabilities with its Distributed Power Management features. In the long term, competition to Intel's data centre business may also come from low-power servers utilizing processors based on ARM Holdings chip-set designs.

5.1.8 SVD Toolkit recommendations and SWOT analysis

As has been detailed, there are a large number of potential CFD, Simulation and Visualisation tools that could be included in the SVD Toolkit. The selection of these will depend on a range of factors including function, openness (open source versus proprietary), ease of integration and cost.

The choice of components for the SVD Toolkit should also take into account existing technologies and techniques in data centre monitoring and management. These include the growing market for DCIM tools, and other related data centre planning and operational management tools. Project CoolEmAll (through dissemination and exploitation actions) should look to cooperate with existing DCIM tool suppliers and other relevant suppliers in this area. (Efforts have already been made through the inclusion of smart PDU supplier Racktivity in the project advisory board).

However DCIM is still a relatively nascent market and there appears to be significant opportunity to add to the state of the art in this area – particularly around the area of application and workload optimization which is currently not well served by any existing supplier. The strength of the SVD Toolkit is an integration of these application and workload analysis with thermal simulations, and advanced visualisation. Additionally, the use of predefined open designs (DEBBs) should facilitate simulations and increase chances for adoption. As a result SVD Toolkit should enable interactive analysis of data centre energy efficiency supporting data centre planners, administrators, and data centre hardware and management systems providers.

Strength	Weakness
Combining simulation, visualization, CFD and workload monitoring tools into a coherent platform should make the SVD Toolkit unique in the data centre management arena	Integration between the various components of the SVD Toolkit may be challenging and costly (in terms of time and resources)
Opportunities	**Threats**
The market for coherent and effective data centre monitoring, management, and simulation tools is still emerging and there may be room for cooperation with existing suppliers and standards bodies	The ambitious proposition behind the SVD Toolkit may be too complex for the majority of mainstream data centres to use and uptake may be low as a result.

Table 6 SWOT analysis of SVD Toolkit concept

5.2 Alternatives to DEBB blueprints

The main aim of developing the Datacenter Efficiency Building Blocks (DEBB) blueprints and the individual ComputeBox components, is to attempt to model how various types of server hardware perform in the datacentre under a variety of conditions. For the purposes of the CoolEmAll project, specialised hardware from Christmann (RECS) will be the primary development and experimentation platform. The Christmann technology has excellent inbuilt sensors and monitoring tools that can be used to gather the (thermal, performance) data necessary for the project. However, the eventual blueprints delivered at the end of the project will also be applicable to the range of more standard servers (and blade servers) deployed in enterprise datacentres, as well as specialist High Performance Computing hardware.

5.2.1 Standard data centre servers

Modern server architectures that are state of the art can be divided into several main lines. Mainframes are very complex computers that have special working environments often used to provide critical business areas. Mainframes retire for the role of standard-based server that can be integrated in our environment due to their very high investment costs and often non-standard functionalities.

Common multi-purpose servers are available in various flavours and with differing configurations. Intel's x86 architecture is the most common architecture used nearly everywhere today. POWER is an architecture developed by IBM, which still has supporters but is not as popular as the x86 architecture. ARM is an architecture that gains attention in the last months due to very low power consumption at reasonable computation power. As already mentioned, HP is building an ARM based high-density server with multiple ARM processors inside.

Most of these architectures come in different power classes. Usually the low-end server class has one CPU that mostly has several cores per CPU integrated. Above this class there is a dual-CPU server class that combines the power of two CPUs on one mainboard which leads consequently makes much memory available within one server. Multiple-CPU servers are also available on the market, but are usually designed to reach a maximum memory capacity and a maximum number of CPUs on one mainboard. Besides this goal the energy-efficiency of these servers is mostly by far worse than the energy-efficiency of one- or two-CPU servers.

To reach a good trade-off between flexibility, energy-efficiency, density and sizing of the servers, the CoolEmAll consortium plan to use multiple RECS servers as the central compute servers. This will allow several CPU architectures, CPU power classes and memory capacities to be used under the same working environment and with the same monitoring and controlling infrastructure.

5.2.2 Blade servers

The high-density approach of Christmann's RECS | Compute Box technology means it will compete primarily with blade server products from suppliers such as Dell, HP and IBM.

The most compact ones (e.g. IBM Blade Center E) allow up to 84 server blades in a 42 U 19" rack cabinet. Blade server solutions have switches and network devices integrated, as well as integrated monitoring and controlling solutions. This integrated infrastructure makes blade chassis quite expensive. The mechanical efforts and the needed material for building up the chassis and the blades are quite high.

Some vendors offer high-density systems. They mostly use a simpler kind of blade centre, which is not as high integrated as the classical blade centres. Supermicro offers a system with 16 low power CPUs (Intel Atom) in 2U, see Tyan (2012). This allows a density of 4 systems per rack unit. A little more density is delivered by Tyan: 18 nodes (with low power Intel Xeon quadcore CPUs) in a 4U enclosure. These systems have only a low integration level, usually each single node is completely separated with no common interconnect or monitoring and controlling.

5.2.3 Low-power servers

Chip manufacturers have been making huge investments in low-powered processors for mobile devices and tablet computers. These have been built from the ground up to run as efficiently as possible, in order to lengthen battery life. Coupled with new developments in software and a new generation of workloads (typically associated with Internet datacenters and increasingly Cloud Computing), these low-power, low-cost CPUs are starting to be used within servers, with potentially dramatic cost savings and at much greater densities. For vendors and users building out cloud infrastructure, both internal and external, the differences can be too big to ignore.

Over the next few years, datacenters looking to scale out their system architectures to support the growing demand for cloud infrastructure will require significantly more efficient alternatives to the current x86 standard server racks and blades. And those alternatives are already starting to appear on the marketplace: racks of servers that rely on low-power CPUs, mostly unmodified from their original mobile and embedded computing designs, and used where there are energy-supply limitations or where maximum rack density is a requirement. Given the right workloads, such systems can turn out to be much more efficiently than larger, monolithic systems. Utilization rates go up, capital costs go down, high levels of redundancy are achieved, and running costs are lowered, from both lower power consumption and less supporting cooling infrastructure.

Hewlett-Packard become the first major server vendor to endorse the emerging concept of extreme low-energy servers. The company announced Project Moonshot in early 2012[ix], a multiyear strategy to develop a new line of hyper-scale servers using CPU cores originally developed for mobile and embedded systems devices by ARM Ltd. The first server in the family, the ProLiant SL Redstone server, is due out next year and will use a system-on-a-chip developed by Austin, Texas-based startup Calxeda, a company ARM is a strategic investor in.

However, there are significant barriers to adoption at this stage for the low-power server concept. The most obvious is that mobile CPUs don't yet support many of the 'must have' features of full-blown x86 CPUs. ARM, for instance, hasn't yet implemented 64-bit support – though it has just announced the new ARM8 architecture, its first with a 64-bit instruction set, and says it will disclose the first ARM8 processors next year, meaning they will likely reach the general market in 2014.

Meanwhile Intel's primary competitor to ARM, the Atom processor, faces different hurdles. Atom is 64-bit and supports the x86 instruction set. It is far less efficient than ARM, doesn't benefit from volume deployments in other market sectors (at least not when compared to ARM), and it doesn't yet support ECC memory (that's planned for next year). However, server startup SeaMicro has been shipping real systems for over a year, and has some significant customer case studies to prove out its story.

The RECS Box technology from Christmann is capable of using processors from ARM (and Intel’s ATOM) and as such the CoolEmAll project is capable of modeling the performance of low-power servers in the datacenter and exploring this important and emerging area. However, when large numbers of these low-power chips (ARM and ATOM) are combined in a single platform (as with the RECS Compute Box) it is also possible to use the combined processing power for High Performance Computing applications.

5.2.4 HPC devices

It is also possible that HPC suppliers will also adopt the ComputeBox blueprints. Existing HPC-specific providers include Bull, Cray, SGI and T-Platforms, Convey Computer and ClusterVision.

Today HPC is built up mostly with standard components: There are many existing HPC-systems that are based on blade servers (e.g. the “Mare Nostrum” supercomputer of the Barcelona Supercomputing Centre, see BSC (2012)). Some still use systems with even lower density, for achieving lower capital expenditure or to realize configurations that are not doable with blade servers.

Some vendor-specific systems are non-standard solutions with new form factors and special components that are allowing a higher density. A well known example of this is the IBM Blue Gene with 1024 nodes in one rack.

5.2.5 Modular data centres

The project plans to develop several classes of ComputeBox prototypes and blueprints (defined as Data centre Efficiency Building Blocks DEBB). The most complex of these DEBBsDEBBS will be the ComputeBox2 that is designed to reflect a containerized data centre or a collection of racks. However, the ComputeBox1 could also be considered as equivalent to a data centre in one standard rack (this is how Christmann has characterised its RECS| Box technology).

Container data centres are essentially large-scale ‘data centres in a box,’ with power, cooling and rack equipment all within the same ISO container.

Existing data centre container products include IBM's Portable Modular Data Center, Microsoft (container data centre in Chicago), Dock IT (purpose-built container park in Virginia) and Pelio and Associates (container colocation data centre in California) – all are big supporters of the container approach. Rackable Systems (renamed SGI), HP (POD and EcoPOD) and Dell also play in the container market.

However, the term modular data centre refers to more than just containerised data centres. The data centre industry is currently undergoing a general shift from data centres being build to order (bricks and mortar) to a more industrialised approach where individual modules (of IT, power, cooling) are constructed in factories and then assembled on site. This approach has been characterised as Data centre 2.0 by industry analysts 451 Research[x].

	Pre-fab IT	Prefab Power	Prefab Cooling	ISO Container	Power/Cooling/IT individually scalable
Traditional data centre					X
Containerised data centre	X	X	X	X	
Traditional data centre with modular components		O	O	O	X
Data centre 2.0	X	X	X	O	X

Table 7 Data centre architecture comparison

X = has characteristic, Null = does not have characteristic O = may have characteristic

The CoolEmAll project should be cognisant of this shift towards modularity and its implications for data centre design and operation – particularly energy efficiency.

While modular design does not have any specific impact on the energy efficiency of data centre operation, the shift towards modules of IT, power and cooling being pre-packaged, as opposed to constructed on-site by IT and facilities administrators may eventually limit the amount of customisation which data centre operators are prepared to do themselves on-site. Fortunately, the Christmann technology, lends itself to this idea of pre-packaged IT, cooling and power being delivered to the data centre with little on-site adjustment required. However this kind of approach has yet to gain wide-spread acceptance from the majority of data centre operators so it is important that the project also considers how the DEBB blueprints, and other components of the project can be applicable to mainstream facilities.

5.2.6 Facebook OpenCompute project

Facebook's Open Compute Project[xi], launched in April 2011, is an initiative to share the company's custom designs for servers, power supplies and UPS units developed for its first dedicated data centre in Prineville, Oregon. The initiative coincided with the negative publicity Facebook received from Greenpeace and others in the wake of its decision to use a coal-intensive power provider in Prineville.

Open Compute could have been written off as a Facebook marketing exercise, but now appears to be gaining genuine traction in the industry. For example, Digital Realty Trust joined the project in Q3 2011, and has announced plans to adapt its data centre designs to include the option to adopt elements of Open Compute designs. It has also been reported that elements of Open Compute will also be deployed in Facebook's colocation space in Virginia. DuPont Fabros Technology (DFT) has agreed to adopt elements of Open Compute in its Ashburn, Virginia, facility.

The second Open Compute Project Summit took place in Q3 2011, attended by the likes of Intel, Rackspace and Goldman Sachs, and included the formation of a number of working groups to help develop standards around elements of the initiative. More events are scheduled for 2012, and the momentum around Open Compute is only likely to increase.

Project CoolEmAll should try to work with, or become a member of the Open Compute project, and exchange information particularly around the potential crossover between DEBB blueprints and designs developed by OpenCompute.

5.2.7 DEBB recommendations and SWOT analysis

The DEBB is an abstraction for computing and storage hardware and describes a data centre building block on different granularity-levels including its energy efficiency information. Precise specification of a DEBB format as well as development of specific DEBBs will facilitate data centre simulations and analysis. It will also help to disseminate project results. However, CoolEmAll should monitor other similar initiatives such as Open Compute

Project described above in order to collect requirements and enable compatibility if possible.

Main DEBBs defined within CoolEmAll will reflect Christmann RECS. Christmann's technology has benefits in terms of the built-in granular monitoring and management which will allow large amounts of data on performance to be extracted from the test-bed machines (more so than with the standard servers). The ability to include different processors is also an important feature. In this way, detailed definitions of DEBBs and their validation will be possible.

However it is also important that the Christmann test-bed technology being complemented by tests on standard servers (this is still being evaluated) as it is important that whatever results are produced, including the eventual DEBB blueprints are applicable to standard data centre environments.

Strength Open DEBB designs will facilitate data centre simulations and exploitation of project results. Definition of DEBBs for Christmann RECS will enable use of its in-built sensors, which means that the project will be able to extract large amounts of useful data.	**Weakness** The Christmann technology is non-standard in terms of form-factor and as such the blueprints that result from the project may have limited uptake
Opportunities The DEBB blueprints will be unique in terms of the information they provide on energy efficiency in the industry, especially when coupled with the SVD Toolkit	**Threats** New server, and data centre, form-factors and designs are emerging and the project will need to adjust the project goals around the DEBB blueprints to ensure that they remain relevant.

5.3 Alternative and existing energy efficiency metrics

The majority of data centres today strive for energy efficiency, to make the most of their existing facilities, reduce utility bills and meet environmental goals. Currently, most metrics for measuring data centre efficiency break it into two pieces: the efficiency of the data centre facility, whose job is to support the IT equipment; and the efficiency of the IT equipment itself, whose job is to transform electricity into useful IT work.

The Green Grid's PUE (Power Usage Effectiveness) metric has been widely adopted for measuring facility efficiency. To judge IT efficiency, the industry would like to have a measure of 'useful IT work per watt of energy consumed.' However, creating a tolerable definition of 'useful IT work' has been an enormous challenge.

Some metrics propose a definition of useful IT work and attempt to measure work per watt directly, while other metrics sidestep the issue and look for other markers of server efficiency, such as an efficient power supply. Several

IT energy efficiency metrics exist today, each with their own pros and cons. This report provides an overview of them.

(It should also be noted that IT efficiency is only *one* of many important aspects of effective data centre operations. Other aspects include facility efficiency, effective utilization of IT and facilities assets, reliability, cost and flexibility.)

The table below provides a summary of existing data centre energy efficiency metrics

Organisation	Metric	Focus	Status
Gartner	PPE	Useful work / total facility power	Presumed active, but little public information available
IT industry	DCD	Overall data center power per data center surface	Active
Green Grid	PUE	Overall data center power per IT power	Active
Green Grid	DCiE	Fraction of data center power used by IT equipment	Active
Green Grid	DCeP	Useful work per Watt	Completed as a placeholder concept
Green Grid	Productivity proxies	Useful work per watt	In development
Green Grid	ScE / DCcE	Fraction of a server's time spent on useful work	Completed as a concept
McKinsey & Co., Uptime Institute	CADE	IT efficiency portion still not formally defined	Used in McKinsey engagements
Green500	MFLOPS/Watt	MFLOPS/watt	Active

Organisation	Metric	Focus	Status
SPEC	SPECpower_ssj2008	Operations/sec per watt	Active
TPC	SPECpower_ssj2008	Transactions/sec per watt	Active

Table 8 Summary of existing DC energy metrics

5.3.1 IT Industry – DCD

The DCD (Data Center Density) indicates the ratio between the total power of the data center and the data centre's surface. This metric is largely adopted by the IT industry and is used for data center comparison as well, even it is not clear if it should be considered as data centre's surface the dedicated technical areas where could be lonely dedicated to facilities equipments as UPS, transformers or motor generators.

Additionally, and being power an instant value, it is not clear if it must be taken maximum, minimum or average value within a non-defined time period. Therefore, in case this metric will be considered useful for CoolEmAll's purposes it should be deeply reviewed.

5.3.2 The Green Grid – PUE, DCiE

PUE (Power Usage Effectiveness) and DCiE (Data Center infrastructure Efficiency) were introduced by Malone in 2006 as indicators of how much power of the data center is finally dedicated to the IT equipment. Since 2008, The Green Grid, an IT consortium, has been supporting and redefining them. Since then, these two metrics and moreover PUE, have been widely adopted by IT industry as data center performance metrics and for comparison purposes.

Despite of that, some facts made them to be indicators for data center comparison. The different criteria of measurement as they could be monitored at the PDU (Power Distribution Unit), at the IT equipment directly or at the UPS output and furthermore, being defined as power based metrics (an instant value), they can be calculated with average, maximum or minimum values within different time periods, make them unreliable for data center comparison. Because of that, some efforts have done to solve this inconsistency by the Global Metrics Harmonization Task Force in 2011 defining different PUE levels.

5.3.3 The Green Grid – DCeP, productivity proxies, ScE, DCcE

The Green Grid, is pursuing a definition of 'useful IT work,' but it remains elusive. In 2008, the organization defined an abstract efficiency metric called Data centre Energy Productivity (DCeP),[xii] which measures useful work per unit of energy consumed by a data centre (note that this means it takes into account both IT and facility efficiency). The metric remained abstract, however, since no procedure for measuring useful work was defined. The Green Grid is currently investigating a set of eight 'proxy' metrics that can be measured and that will (in theory) track 'true' productivity. These proxies have been under discussion since early 2009.

The Green Grid has also proposed a metric called Server Compute Efficiency (ScE), which expresses the fraction of the time that a server is doing useful work. The key feature of the metric is that it distinguishes between mere utilization and work that is actually 'useful' in terms of a server's primary function. For example, even a server running at 100% CPU utilization may not be doing much useful work, if most of its processor cycles are spent scanning itself for viruses or backing up nonessential data. However, the ScE metric does not define useful work or quantify how much of it a server does. A companion metric, Data centre Compute Efficiency (DCcE), expresses the average ScE across all servers in a facility.

5.3.4 McKinsey/Uptime Institute – CADE

Management consultancy McKinsey & Co[xiii] and The Uptime Institute[xiv] (a division of The 451 Group) jointly developed another metric for overall data centre performance called Corporate Average Data centre Efficiency, or CADE. CADE consists of four pieces: examining facility efficiency, facility utilization, IT efficiency and IT utilization. For the IT efficiency piece, a true definition of 'useful work per watt' has not been defined. The Uptime Institute has indicated it may do further research in this area.

5.3.5 Green500

Green500 is a list of the most energy -efficient super computers in the world, maintained by a group of professors at Virginia Polytechnic Institute and State University. Essentially, the Green500 seeks to add energy awareness to as a complement to the TOP500, a supercomputing list that focuses on performance alone. Green500's efficiency metric is millions of floating point operations per second per watt (MFLOPS/watt).

5.3.6 SPEC – SPECpower

One metric that does directly address IT work per watt is SPECpower_ssj2008, promulgated by the Standard Performance Evaluation Corporation. SPEC is a non-profit corporation that creates and maintains computing performance benchmarks. SPECpower_ssj2008 measures the performance of a server-side Java application (in operations per second) and power consumption for the tested server at a variety of throughput levels, to create a set of 'ssj_ops/watt' measurements. As of May 2010, server suppliers had submitted SPECpower benchmarking data for over 250 systems.

Critics of the benchmark point out that SPEC's artificial Java workload may not be representative of the workloads a server would see in actual deployment. Critics also note that server vendors have some discretion over certain variables during test runs, and thus may be able to 'game' the metric. (Vendors can control the amount of memory, whether power management features are enabled, etc.) Two other SPEC metrics, SPECweb2009 and SPECvirt_sc2010, can also include power measurements, although few suppliers have chosen to report power data thus far.

5.3.7 TPC – TPC-Energy

Another non-profit corporation that promulgates benchmarks is the Transaction Processing Performance Council. TPC's performance benchmarks focus on database transactions, and they have been extended via the TPC-Energy specification to provide a measure of watts/performance, where performance is defined in terms of (transactional) work per second. However, energy measurements are optional when reporting TPC results, and to date only about 10 systems have submitted energy data.

5.3.8 Gartner – PPE

In 2009, IT research company Gartner proposed an IT efficiency metric called Power to Performance Effectiveness (PPE), which examines rack density, server utilization and energy consumption. PPE allows users to define their own optimal performance levels, and then measure actual performance relative to those. Little information on the exact definition of PPE is publicly available, but one public document indicated that it is essentially a measure of 'useful work' divided by total (datacentre) facility power.

5.3.5 Metrics recommendations and SWOT analysis

There are many active metrics that address IT efficiency, with several more in development. Over time, there has been a clear shift toward more performance-oriented metrics that try to capture useful work per watt, rather than focusing only on power supply efficiency or other indicators. However, there is still no dominant definition of what constitutes useful work. It may be possible for CoolEmAll to contribute to work in this area of useful work via organizations such as the Green Grid.

However the CoolEmAll project should also strive to make the best use of the efficiency metrics that exist, while acknowledging that current metrics are imperfect. End-user organisations will likely make different choices about what metrics to consider and precisely how to use them, based on the organization's particular evaluation of the trade-offs among accuracy, ease of use, availability of data and any costs to implement. Server suppliers should consider measuring themselves against competitors using multiple efficiency metrics, if they are not already. They should resist the temptation to acknowledge only metrics where their products excel.

Wide adoption will be crucial to any metrics that the CoolEmAll project looks to contribute to or develop itself. A defensible metric with up-to-date data on a wide variety of servers will be more valuable to the industry than a metric that is technically superior but seldom used. A successful metric will have a central repository of easily accessible scores for many, many server types. Ideally, it would also have data on multiple configurations of each server (ideally defined within DEBBs), since variations in processor model, RAM and other components can matter a lot.

Strength	Weakness
Developing new metrics or contributing to existing ones would help to raise the profile of the project and an add an extra dimension to the work done around the SVD Toolkit and DEBB Blueprints	Although the CoolEmAll project will run for 30 months, metrics take time to develop and a significant investment of resources.
Opportunities	**Threats**
The consortium is well placed in terms of expertise and contacts (both academic and commercial) to contribute to the development of new or existing metrics	Even if the project manages to develop to contribute to or develop new energy efficiency metrics there is no guarantee that these will be widely adopted in the industry.

Table 9 SWOT Analysis of energy efficiency metrics

5.4 Relevant and complimentary EU projects

There are number of on-going EU projects that are attempting to address some of the same issues as CoolEmAll and the consortium will cooperate with these projects where possible. However while there is some crossover in terms of aims, CoolEmAll has a distinct remit compared to other projects.

Project aim	GAMES	Fit4Green	OPTIMIS	CoolEmAll
Modular design				✓
Simulation-aided Design and Management				✓
New data centre energy metrics	✓	✓	✓	✓
Thermal and energy visualization		✓		✓
Application characteristic driven workload policies			✓	✓
Impact of cooling and heat-reuse on				✓

Project aim	GAMES	Fit4Green	OPTIMIS	CoolEmAll
energy efficiency				
Thermal and energy driven workload management	✓	✓	✓	✓

Table 10 CoolEmAll positioning versus Euro projects

5.4.1 Games

The results of another EC project on data centre efficiency will also feed into CoolEmAll. Partners HLRS and Christmann are also part of that project, called Green Active Management of Energy in IT Service centres (GAMES).

The stated goal of GAMES (Green Active Management of Energy in IT Service centres) is to develop more sophisticated tools for data centre energy monitoring and control. The project organizers argue that current tools work in isolation and do not consider the interaction between applications, computing hardware and the physical facility (such as cooling and power). The GAMES project aims to produce energy-monitoring and -control tools that factor in these interactions. These tools will help managers to design and operate more energy-efficient facilities. Consortium members argue that data centres that adopt the tools will achieve a 25% increase in efficiency. The GAMES consortium is made up of business and research organizations including IBM and the University of Stuttgart. The project is set to run over 30 months from 2010.

5.4.2 Fit4Green

Fit4Green (Federated IT for a sustainable environmental impact) is focused on creating a series of software plug-ins for existing data centre management tools. The plug-ins will be designed to facilitate the movement of virtual machines or virtual workloads between servers within a data centre, but also between federated data centres. The aim is to allow virtual workloads to be moved to the optimal environment from an energy-efficiency perspective. This will include the ability to switch off servers that are no longer being used because virtual machines have been moved to another, better-utilized device. Fit4Green is a 30-month project which begun in 2010 and includes organizations such as Imperial College London, Hewlett-Packard and the University of Mannheim. Viviane Reding, former member of the European Commission responsible for Information Society and Media (replaced by Neelie Kroes), outlined the importance of cloud computing in improving social engagement and the sustainability of government IT. Several cloud projects are being funded under the FP7 program in line with this strategy. However, there has been little attempt to investigate the energy use of cloud computing beyond the nominal assumption that sharing of IT resources are more efficient than locally installed systems.

5.4.3 Optimis

Launched in June 2010, Optimis is a three-year, €10.5m (approximately $14.7m) research and development project that aims to create a framework to improve interactions between cloud providers and customers. The main goal is to enable organizations to automatically place services and applications in the cloud using trustworthy and auditable cloud providers. The initiative involves companies such as BT and cloud specialist Flexiant cooperating with academics and researchers (including Tier1 Research/The 451 Group) to develop an open source platform for procuring cloud services. Part of the project is establishing how a consumer of a cloud service can automatically assess the eco-efficiency of the provider – for example, by collecting energy-efficiency data. If possible, the toolkit that the project produces will include mechanisms to select or even apply eco-efficiency strategies.

5.5 Other relevant EU groups and bodies

The EC also provides support to groups and consortia made up of representatives from industry, academia and the public sector, which may prove useful to CoolEmAll and vice versa.

5.5.1 The ICT4EE Forum

Established by the EC and elements of the IT industry on February 23, 2010, the forum focuses on the two key aspects of eco-efficient IT: how the technology industry can curb its energy use, and how it can help industry generally do likewise. Four industry associations have signed up so far, representing the European, Japanese and American IT industries: Digital Europe, the Global e-Sustainability Initiative, the Japanese Business Council Europe and TechAmerica Europe.

5.5.2 Euro Green IT Innovation Centre

Launched in January 2010, the Euro Green IT Innovation Centre is a public-private partnership and a non-profit association based in Brussels. The centre is focused not only on developing more sustainable IT equipment and services, but also improving the energy efficiency of other sectors through better use of technology. The organization will hire a number of staff directly and also intends to fund five to 10 start-ups a year in the field of energy-efficient IT. Funding is provided from a €750,000 grant from Belgium taxpayers, but it is also being backed by IBM, Cisco, Microsoft and the Climate Savers Computing Initiative.

Name	Format	Focus	Time Frame
ICT4EE Forum	Industry group	Industry/public initiative focused on green IT	Ongoing
GAMES	Project	Data centre energy Efficiency metrics and	2010 - 2012

Name	Format	Focus	Time Frame
		management	
Fit4Green	Project	Data centre energy efficiency and load management	2010 -2012
The European Green IT Innovation Centre	Industry group	Sustainable IT equipment and and services	Ongoing
Optimis	Project	Cloud computing (efficient)	2010 - 2013

Table 11 Summary of relevant industry groups and projects

6. Legislative landscape

As well as assessing the potential of the CoolEmAll project from a commercial and research perspective, it is also necessary to understand how policy and legislation is shaping the market for data centre energy efficiency tools.

There a wide-range of international laws and policies that could shape some aspects of project CoolEmAll (e.g. US Energy Star and Energy Star for Data centres). However given the European focus of the project, it makes sense to focus on relevant local legislation and policy initially. This may be broadened as the project progresses.

6.1 Code of Conduct on Data Centre Energy Efficiency

The EU Code of Conduct on Data Centre Energy Efficiency sets out guidelines for data centre energy efficiency levels, and sets up a metrics and monitoring system for tracking the progress of participants (see below). Participation is voluntary, but the code is seen by many as a framework document and a pilot data-collection methodology for a future European directive (although we believe active legislation is many years away).

The EU's goal is to ensure that data centres are demonstrably improving. To that end, it will collect significant amounts of data from data centres, including energy use, adoption of technologies and best practices. It will also develop, adopt and publicize metrics, so that data centres can be compared and (eventually) given targets. As a start, it will use The Green Grid's DCIE (Datacenter Infrastructure Efficiency) ratio (the reciprocal of PUE) and at least two others that, when ready, will attempt to get a view on the overall effectiveness of the IT operation.

To qualify for code status, participating data centres must file a detailed report, as well as monthly IT and total-facility energy-use reports, at least twice a year. In this way, the EU is creating a framework for data collection for the future. Initially, at least, the EU will collect the data for auditing and for anonymized analysis. If the code works well, it could be made mandatory under European law to encourage energy efficiency among nonparticipants; conversely, if it doesn't produce results, the EU will seek a tougher approach.

The code was bolstered in October 2010 by the decision of six major companies, including France Telecom and Telecom Italia, to join the 26 existing signatories. Aside from the participants, data centre owners and operators that comply with the code fully, there are more than a hundred equipment and service providers endorsing the initiative, including IBM, EMC and Microsoft.

Project CoolEmAll should be familiar with the Code of Conduct and how it could help shape the future direction of the project. The project should look to foster good relations with the individuals and organisations overseeing the Code as well as the companies and organisations adhering to it.

6.2 EU Emission Trading Scheme (ETS)

Outside the US, carbon legislation is more established. The European Union Emissions Trading System (EU ETS) is the cap-and-trade system for GHG emissions in the EU. The system operates in 30 countries (the 27 EU member states, plus three additional countries), and it currently covers around 11,000 installations, which account for roughly 40% of the EU's total GHG emissions. The EU ETS began operating in 2005 and is now in Phase 2 of its implementation. Phase 3 will begin on January 1, 2013.

Phase 3 will be more stringent than previous phases, which failed to make a significant impact because large numbers of free carbon permits were distributed to emitting industries, mostly power companies. With permits free or very cheap, industries had little incentive to reduce emissions. Phase 3 will include more types of GHGs, cover more industries, reduce the number of emissions allowances issued, and gradually shift toward auctioning allowances rather than giving them away.

The ETS impacts data centres mainly via its impact on electricity prices so it has minimal relevance to CoolEmAll at the moment. However, carbon prices within the ETS may go up starting in 2013, with the advent of Phase 3 and its tighter rules. However, carbon prices may also be more stable if the ETS implements some proposed changes to reduce volatility. Project CoolEmAll should continue to track developments in the EU ETS and consider how it may focus attention on energy and carbon efficiency approaches to data centre operation and design.

6.4 UK CRC Energy efficiency scheme

The UK's CRC Energy Efficiency Scheme is one of the few carbon taxes in the world that regulates the carbon emissions of business consumers of energy, rather than of energy producers. The CRC came into force in April 2010. The data centre industry in the UK has lobbied furiously against the initiative from the outset, mostly on the grounds that it will push up costs for energy-efficient operators in the UK, and force business offshore, where regulations are weaker. Now the program looks as though it will be scrapped or replaced (as of early 2012).

But until the future of the CRC is confirmed, data centre operators will have little choice but to act as if the rules still stand. This may be onerous for some, but others will already have reporting systems in place to make that process easier. For those that don't, investment in accurate energy and carbon measurement and reporting systems is worthwhile even if the CRC is eventually eliminated.

Under CRC rules, all organizations with a usage of 6,000 MWh of electricity per year (measured by half-hourly meters), including UK data centre operators and owners, were required to become full participants in the initiative and register by September 30, 2010. The CRC has undergone a number of important changes over the last 23 months. The decision to possibly axe the scheme altogether is the latest chapter in a history of rewritten legislation and ambiguity.

In Q2 2012 (as part of the annual budget statement), UK chancellor George Osbourne publicly described the CRC as cumbersome, and said it imposed unnecessary costs on businesses. If it cannot be improved, it should be replaced, he argued. The UK Department for Energy and Climate Change (DECC), the government agency charged with overseeing the implementation of the CRC, plans to launch another consultation (the CRC has been through at least five consultation rounds in to date) slated to begin on March 27. This round will focus on simplifying the CRC by adjusting the reporting process and clarifying how long participants will be required to store information.

If the latest consultation is unsuccessful in resolving all the many outstanding issues – which seems likely given previous attempts – the CRC will be replaced by an as-yet undefined environmental business tax. (DECC has not provided further details on what form this tax would take.)

Project CoolEmAll should continue to monitor developments in the CRC as it will have implications not only for the adoption of energy and carbon efficient technology in the UK but also provides a model for the adoption (or not) of similar laws in other countries.

7. Conclusions and next steps

CoolEmall is undoubtedly an ambitious project with a wide scope. However such ambition is necessary to deal with the myriad of issues around data centre energy (and carbon) efficiency. The following section summarises the recommendations and conclusions from the report and suggests future ideas that may be incorporated into the project as it develops.

7.1 SVD Toolkit

- The choice of components for the SVD Toolkit should also take into account existing technologies and techniques in data centre monitoring and management domain. (However, the SVD Toolkit is built on top of monitoring and management tools – provided by MOP (module operating platform). The monitoring part of MOP is flexible enough to be adapted to any source of information).
- The SVD Toolkit should be seen as operating one level higher than most DCIM systems especially given its focus on the impact of application characteristics and workload and management, which is not currently a feature built into most DCIM systems.
- The strength of the SVD Toolkit is an integration of application and workload analysis with thermal simulations, advanced visualisation, and predefined open designs (DEBBs).
- DCIM is still a relatively nascent market and there appears to be significant opportunity to add to the state of the art in this area.
- Project CoolEmAll (through dissemination and exploitation actions) should look to cooperate with existing DCIM tool suppliers and other relevant suppliers in this area.

7.2 DEBB Blueprints

- Open, predefined DEBB designs should facilitate data centre modelling, simulation and visualisation as well as exploitation of project findings.
- The ComputeBox1 DEBB could be considered as equivalent to a data centre in one standard rack (this is how Christmann has characterised its RECS| Box technology).
- The data centre industry is currently undergoing a general shift from data centres being build to order (bricks and mortar) to a more industrialised approach where individual modules (of IT, power, cooling) are constructed in factories and then assembled on site.
- The CoolEmAll project should be cognisant of this shift towards modularity and its implications for data centre design and operation – particularly energy efficiency.
- The Christmann technology, lends itself to this idea of pre-packaged IT, cooling and power being delivered to the data centre with little on-site adjustment required.
- Project CoolEmAll should try to work with, or become a member of the Open Compute project, and exchange information particularly around the potential crossover between DEBB blueprints and designs developed by OpenCompute.

7.3 Metrics

- There are many active metrics that address IT efficiency, with several more in development.
- There is still no dominant definition of what constitutes useful work. It may be possible for CoolEmAll to contribute to work in this area via organizations such as the Green Grid.
- The CoolEmAll project should strive to make the best use of the efficiency metrics that exist, while acknowledging that current metrics are imperfect.
- Wide adoption will be crucial to any metrics that the CoolEmAll project looks to contribute to or develop itself.
- A successful metric will have a central repository of easily accessible scores for many, many server types. Ideally, it would also have data on multiple configurations of each server, since variations in processor model, RAM and other components can matter a lot (should be defined within DEBBs).

7.4 Overall project

- There is growing interest in improving data centre efficiency driven by a range of factors including facility operating and capital expenditure costs, and environmental legislation.
- Project CoolEmAll should continue to track developments in the EU ETS and consider how it may focus attention on energy and carbon efficiency approaches to data centre operation and design.
- Project CoolEmAll should be familiar with the DC Code of Conduct and how it could help shape the future direction of the project. The project should look to foster good relations with the individuals and organisations overseeing the Code as well as the companies and organisations adhering to it.
- Project CoolEmAll should continue to monitor developments in the UK CRC Energy Efficiency scheme.
- There are existing technologies that solve some of the problems tackled by the individual components of the CoolEmAll project.
- The key differentiator and value of the CoolEmAll project is how each of its elements (SVD Toolkit, DEBB blueprints, workload management, and proposed metrics) are integrated together to provide a comprehensive suite of tools.
- Reaching out to end-users, standards bodies, and supplier groups will be vital to ensuring the results of the project are exploited successful. CoolEmAll should start defining details of exploitation by potential early adopters (Christmann, ATOS, HLRS, PSNC) as soon as possible.

8. References

[i] Projecting Annual New Data Center Construction Market Size, Microsoft Global Foundation Services (March 2011)

[ii] 2011 Data Center Industry Survey, Uptime Institute, June 2011,

http://uptimeinstitute.com/publications

[iii] Memorandum for Chief Information Officers, Vivek Kundra (Federal CIO) February26, 2010, *http://www.cio.gov/documents/Federal-Data-Center-Consolidation- Initiative-02-26-2010.pdf*

[iv] US Department of Energy, Data Center Profiler Software Tool Suite, http://www1.eere.energy.gov/manufacturing/datacenters/software.html

[v] Data Centre Efficiency calculator from APC (Schneider Electric), http://www.apc.com/prod_docs/results.cfm?DocType=Trade-Off%20Tool&Query_Type=10

[vi] BCS Data centre Strategy Group/Carbon Trust Data centre Simulator, http://dcsg.bcs.org/

[vii] Romonet Prognose, (recently renamed Romonet Software Suite http://www.romonet.com/content/product-overview)

[viii] Lumina Decision Systems' Analytical Data centre Capacity Planning tool, http://www.lumina.com/products/capacity-planning-tools/data-center-capacity-planning-tool/

[ix] HP ProjecMoonshot, http://www.hp.com/hpinfo/newsroom/press/2011/111101xa.html

[x] Datacenter 2.0 – The Industrial Evolution, 451 Research, September 2011, https://451research.com/report-long?icid=1663

[xi] Facebook Open Compute Project, http://opencompute.org/

[xii] Green Grid, Data Center Energy Productivity Metric, http://www.thegreengrid.org/Global/Content/white-papers/Framework-for-Data-Center-Energy-Productivity

[xiii] McKinsey & Co, http://www.mckinsey.com/.

[xiv] Uptime Institute, http://uptimeinstitute.com/