ECMWF Feature article

from Newsletter Number 126 – Winter 2010/11

COMPUTING

# Green computing

www.ecmwf.int/en/about/news-centre/media-resources

doi:10.21957/cemub9f8

*This article appeared in the Computing section of ECMWF Newsletter No. 126 – Winter 2010/11, pp. 28–31.*

# Green computing

Sylvia Baylis

In recent years, environmental awareness has increased but economic pressures have also led to organisations paying closer attention to the costs, especially energy costs, of IT operations. If energy use can be reduced, carbon emissions will reduce too, bringing environmental benefit. This is particularly relevant at centres such as ECMWF where supercomputers and their associated infrastructure have large electricity needs.

Given the need to minimise energy use, a study was undertaken with the aim of having an environmentally-friendly green supercomputer installation, with particular focus on energy efficiency and low carbon footprint. This will set the scene for the future when the power and cooling requirements of supercomputers may increase significantly which will pose major engineering and budgetary challenges.

Typical data centres use twice the electrical power required by their computer equipment, whereas ECMWF uses less than an additional 50% of energy for its infrastructure, on top of that used to power the IT equipment. This reflects the importance that was attached to increasing energy efficiency during computer and infrastructure procurements in previous years.

This article describes the measures that have been taken to reduce energy consumption in the computer building at ECMWF and considers some options for further reductions.

### Where is the energy used?

The electricity used by a supercomputer of given computational power is fixed as far as the user is concerned. Manufacturers of supercomputer equipment have realised that they must limit the increase in electricity consumption of future hardware to be within electrical load limitations of major data centres and for running costs to be affordable. The infrastructure supporting a supercomputer consumes large amounts of electricity that contributes to the running costs of the data centre and to the carbon footprint.

To reduce the energy consumption in an individual data centre it is necessary to know how much power each piece of equipment uses. This leads to a requirement to measure both the total energy use of a data centre and how this is divided amongst the components. To find out how much electricity each piece of equipment uses, ECMWF installed power meters around the site. Half hourly readings are collected via a LAN and the central software logs all meter parameters: voltage, current, power factor, power and energy. The meters allow ECMWF to measure the electricity usage at power distribution unit level for computer equipment and at individual unit level for infrastructure such as chillers. Readings from the meters are recorded so that, over time, changes in the energy efficiency of the computer building can be assessed and quantified.

### Comparing supercomputer data centres

In a typical data centre the overall energy consumption is determined by the infrastructure, computer and other hardware, software, and how the system is managed. To compare the energy efficiency of one computer centre with another there is a need for standard metrics. The measure that is used at ECMWF is the Power Usage Effectiveness (PUE) that measures the efficiency of the environment around the supercomputer. PUE is defined as the ‘total power input to the data centre’ divided by ‘power input to computers’.

The PUE of a typical data centre is about 2 although there are data centres with PUEs exceeding 4. New data centres are being constructed with the aim of consuming only 10–20% additional energy for the infrastructure corresponding to a PUE in the range 1.1 to 1.2, but these centres are generally only suitable for standard IT servers and not supercomputers (see Box A). A small number are attempting to reduce their PUE further by re-using their waste heat. This is not possible for all data centres; it depends on the temperature of the waste heat, its location in relation to any potential users, and the viability of transporting the heat. Also, heat is produced all year but might not be wanted by the users in summer (e.g. for district heating systems).

 doi:10.21957/cemub9f8

## ECMWF Power Usage Effectiveness

ECMWF has been measuring its energy consumption since spring 2010 to calculate its PUE for the computer building. The power input to the computer building is calculated by subtracting the electricity consumed by the two office buildings and the conference building from the total electricity consumption for the site. This takes into account all losses within the main power distribution equipment (i.e. transformers and Uninterruptible Power Supply (UPS) systems). As the power input to the computers is known, the PUE can be determined. There is not sufficient sub-metering to enable small loads such as street lighting and ancillary buildings to be quantified.

ECMWF calculates PUE every 30 minutes from the half-hourly consumption data provided by its electricity supplier, and half-hourly data obtained from power metering equipment installed throughout the site. Averages for a day, week and month are then produced based on the half-hourly PUE figures. The intention is to build a rolling average over time to obtain the annual average for the computer building. Analysis of the data shows that the PUE of the computer building is about 1.45 (see Figure 1), with the computer building extension, where there are few non-supercomputer racks, having a lower PUE than the main computer hall. A ratio of 1.45 means that for every kilowatt of electricity used directly by the computer systems an additional 0.45 kW is attributed to the ancillary services, such as cooling and lights, and losses in the power distribution system.

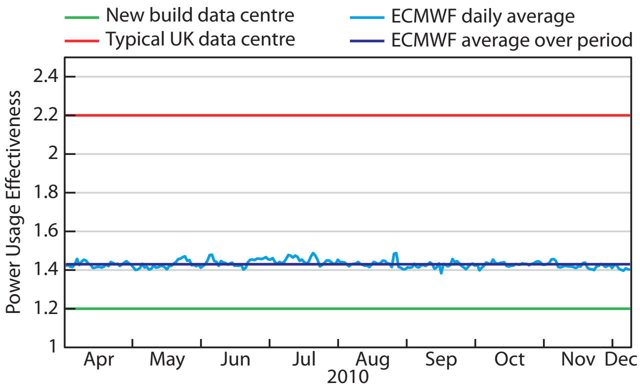


**Figure 1** Values of the Power Usage Effectiveness (PUE) for the ECMWF computer building compared with a typical UK data centre and a new build data centre.

### Current green data centres

**A**

Several UK data centres that advertise their 'green' credentials (PUEs less than 1.2) were visited as part of the study, to examine the technology in use and its applicability to meet supercomputer infrastructure requirements. These centres were a mixture of hosting sites and showcase green facilities operated by IT vendors.

Although there were interesting approaches to energy efficiency at the data centres, they were, in general, not applicable to the ECMWF data centre. Most made the assumption of equipment being in fixed, standard size racks, and that future equipment would also fit these racks. The cooling methods were designed for air cooled equipment with a maximum heat output around 20 kW per rack, and the cooling methods used could not be modified to provide chilled water directly to cool equipment, such as supercomputer racks.

ECMWF uses supercomputer equipment which, in general, is installed in custom racks and is replaced every two to three years. Successive supercomputer systems are likely to have different size racks in differing configurations. As the lifetime of the data centre building is many times that of the supercomputer systems, flexibility in the design and layout of the data centre is needed; the adaptations for each new system are likely to reduce the efficiency of the data centre as a whole.

## Power losses at ECMWF

### *Power distribution system*

The energy losses associated with power handling and distribution on-site are likely to be between 1% and about 10% of the site load. The type of Uninterruptible Power Supply (UPS) chosen will make most impact on the additional energy use. ECMWF uses a rotary UPS to protect the site from loss of power due to loss of mains electricity. It works by using the incoming electrical supply to turn a flywheel that itself is connected to an integrated generator. If the National Grid supply is lost, the flywheel continues spinning and feeding energy into the electrical distribution system long enough for a diesel engine to cut-in and drive the generator. When power is flowing normally through the UPS from the National Grid, the UPS conditions the power supply, e.g. by removing transient spikes. These UPS machines are about 96% efficient. Static UPS machines, which support the computer equipment power load on batteries between the occurrence of a mains failure and the start of the backup generators, are typically 90% efficient. However, improvements in static UPS technology in recent years has increased efficiency to around 95% for the newest such machines.

There are also losses of 2–3% in the transformers which convert the incoming 11 kV supply from the National Grid into 400 V supplies for the site. Further, small, losses occur in the electrical distribution system comprising switchboards in the basement of the computer building, power distribution units in the computer halls and all the associated cables.

### *Cooling system*

The centralised chilled water systems at ECMWF use packaged air-cooled chillers incorporating Turbocor compressors (see Figure 2) that are designed to be more efficient at part load and at low ambient air temperatures. Traditional screw or reciprocating compressor based chillers are most efficient at full load, and the efficiency variation with ambient temperature is minimal. The average annual efficiency for a Turbocor based chiller is between 5.5 and 6.5, as opposed to being in the range 3 to 3.5 for a more traditional style chiller. This means that a Turbocor chiller uses half as much electricity as a conventional chiller to provide the same cooling. All chillers at ECMWF operate at less than full load because they are run in a redundant mode with spare capacity available, and they run continuously throughout the year as in most data centres. Consequently the saving to be made using Turbocor compressors is significant.


**Figure 2** One of ECMWF's Turbocor based chillers.


doi:10.21957/cemub9f8

## Reducing energy usage

### *Electricity distribution system*

Power losses in the on-site electricity distribution system are more or less fixed. The only way to reduce the losses would be to run all or part of the computer equipment from raw mains to remove the losses in the UPS. This would have an impact on the resilience of the computers and is not suitable at an operational centre where the computers are required 24 hours per day. These losses are small when compared to the power used by the cooling systems for the computer equipment.

### *Cooling*

Over the last few years ECMWF has had a programme of replacing old, relatively inefficient chillers with new chillers based on Turbocor compressors. As described above, these have reduced the electricity used by the chilled water systems. The most efficient cooling option can only be determined from a detailed analysis of a specific site, operating conditions, loads etc. The solution may also be influenced by the level of resilience required.

There can be little doubt that the most efficient cooling method is the use of ‘free cooling’ by using the ambient (external) air to cool the air or water that provides the direct cooling to equipment in the data centre (see Figure 3). Raising the internal design temperatures and consideration of the location may allow ‘free cooling’ to operate for the majority if not all operating hours. Where lower temperatures are required for water-cooled systems such as ECMWF’s supercomputers, free cooling can still be an effective option but will be available for less hours.

The main disadvantage of free cooling systems is that plant and equipment tends to be larger to take the maximum advantage of small temperature differences. This increases capital costs and space requirements but these costs will often be justified by life cycle cost analysis. The most important factor however is that free cooling systems will significantly reduce carbon emissions that would otherwise arise from the generation of electricity to drive mechanical cooling systems.

### *Data centre design and layout*

Other measures that can be used to increase the energy efficiency of a data centre involve changes to the layout of equipment. These are mainly relevant for air-cooled (non-supercomputer) equipment and are only likely to be practical to implement when a major change of IT hardware is taking place.

Air-cooled computer equipment works by drawing air from the computer hall into the equipment cabinets, passing it across the equipment, and venting it into the computer hall. Air typically enters the data centre under the floor, passes through vents in front of the cabinets containing the equipment and then is extracted through vents mounted high in the walls. There are two ways of increasing the effectiveness of the cold air supply.

- *Cold aisle containment*. This uses a physical partition so that all the cold air goes directly into the equipment racks and there is little recirculation or mixing with hot air. Cold aisle containment increases the external temperature at which free air becomes ineffective.
- *Hot aisle containment*. With this approach the air passed across the equipment is vented directly into return ducts.

**Figure 3** External view of a typical free cooling module. Photograph provided courtesy of Keysource (www.keysource.co.uk).

doi:10.21957/cemub9f8 

Modelling the airflow within the computer hall allows the system to be optimised to use the minimum amount of cold air. This system works best with equipment that fits into standard racks so that computer equipment can be changed without having to modify the cold/hot aisle containment.

It may also be possible to increase the temperature in the data centre. General purpose IT equipment is manufactured to operate within standard temperature and humidity ranges – known as the ASHRAE standards. These standards have evolved with time, and equipment being manufactured now is capable of being used reliably at higher temperatures than older equipment. Raising the temperature of computer halls can make a large difference to the cost of cooling. However, ECMWF has so little non-supercomputer IT equipment that raising the air temperature would have minimal impact on the cost of cooling.

## Carbon neutrality

As well as reducing the energy consumption on the site ECMWF has investigated ways of working towards carbon neutrality. The strategies available are:

- Purchase 'green' electricity through a conventional supplier;
- Invest in a supplier of renewables, 'earmarking' some of the renewable energy for use by ECMWF – the benefits being that the price can be negotiated explicitly (and thus is less susceptible to market variations) and that the source can be identified;
- Site a data centre where it is possible to generate renewable energy. In practice, as ECMWF does not wish to be a power generator, this may be a variant on the second option.

Of these options, the second is least affected by the location of the data centre and has the potential for electricity price stability. However, this is a long-term option which requires further study.

In the short term, the first option is being pursued. 'Green' tariffs are available from most large energy providers based on the premise that a proportion of the electricity supplied is generated from sustainable sources. At present the sustainable proportion is only in the order of 10%, although this is likely to increase over the next few years if government policy objectives are to be achieved. However, it is unlikely that a significantly higher proportion of grid-supplied energy will be produced from sustainable sources in the near future. The demand for green electricity exceeds the available supply. Much of the green electricity in UK is reserved for local government which has the most stringent targets for sustainable operations and reduction of carbon footprint. The shortage has produced a rise in the cost of renewable energy and, so called, Good Quality Combined Heat and Power (GQCHP). For the contract starting from April 2011, ECMWF's electricity provider will supply 25% of its electricity from renewable and GQCHP sources on a cost neutral basis. Use of green energy in the long term is being pursued.

## What next?

ECMWF operates a very large computing system, by commercial standards, and of necessity uses infrastructure equipment that can deliver low unit costs without having to share its use. The trend for supercomputers is to move to watercooling (for which there are no standard configurations), higher density (resulting in greater floor loading than for conventional IT), and non-standard racks. Together these mean that each new supercomputer installation will need changes to the supporting infrastructure. As the infrastructure is changed the opportunity will be taken to install more energy efficient equipment, where possible. Alongside this, ECMWF is looking at options for the use of free cooling to provide at least part of the required data centre cooling capacity.



 doi:10.21957/cemub9f8