

arjuna 

Agility™ for Condor



Agility Delivers Intelligent Power Management for Condor

Introduction

Universities, like most large organisations need to make cost effective and energy efficient use of their IT resources. However their research and teaching departments have a significant and growing need for IT consumption. Attempts to resolve these potentially conflicting requirements can be disruptive and are difficult to enforce. The installation of Arjuna's Agility product at Newcastle University has allowed these conflicts to be resolved in a unique way, and paves the way for adoption of the same techniques in other similar situations.

University research

In common with many Universities, Newcastle University obtains significant funding from research grants and success in research makes the university a popular choice for quality students. Research is a core business function of the University. Much of modern research is compute-intensive and becomingly increasingly so, and it is often the case that the quality of research results can be improved by access to additional IT. Indeed there are projects for which, if economics were to be ignored, near-infinite IT capacity could be consumed.

Much of the IT infrastructure within a University is unused for significant periods of time. For example, project-funded IT is only intensively used during particular phases of a project, and much department based IT is unused outside of office hours. At Newcastle, the majority of the University's compute intensive research tasks are executed on this spare computing capacity available over the entire university spanning many departments and shared facilities. This provides an invaluable shared resource enabling researchers across campus to carry out computing tasks which would be impossible within the confines of their own project budgets.

To date at Newcastle IT resources have been shared on a quid pro quo basis with departments treating their IT resources as sunk capital costs and therefore 'free' for use by others.



Carbon and Power saving

Energy costs are escalating and are predicted to continue to do so for the foreseeable future. Those costs now exceed capital costs over the lifetime of an IT resource. (Total energy costs per annum of IT infrastructure within Newcastle University would exceed £2M per annum if all equipment were left permanently on).

Additionally, government legislation and public opinion is increasingly concerned with the Carbon footprint of large organisations. As a consequence Newcastle University has announced that it is committed to reducing both carbon footprint and energy costs, and has introduced a set of initiatives intended to make departments, projects and individuals accountable for the carbon and energy they consume.

The conflict of interests

There is unfortunately an essential conflict between those tasked with reducing energy and carbon consumption and the broader interests of a university's research community. Whilst departments are capable of effective cost/benefit analysis for work of direct concern to themselves, they tend not to take into consideration the overall benefits of other external research tasks to the University. As an example, at Newcastle increased accountability had resulted in departments powering down IT infrastructure during out of office hours, thereby reducing the spare capacity available to researchers. Unchecked this tendency would have had a significant negative impact upon University research.

This conflict can be resolved through the use intelligent power management policy for IT resources which takes into account the respective needs of all parties. This can be achieved by introducing policies which:-

1. gather information regarding both energy costs involved and the benefits accrued through the expenditure of that energy.
2. enact policy which can use that information to perform a cost/benefit analysis.
3. depending upon that analysis, either control the demand for IT capacity (by rejecting, or delaying requests for capacity) or control the supply of IT capacity (by dynamically powering up or down IT resources).

Without effective power management energy will continue to be wasted and/or research results will suffer from reduced IT capacity.

The Arjuna Solution

In order to deliver intelligent power management Newcastle University has utilised Arjuna's Agility product. Agility is a framework for federated cloud computing management, designed to improve business agility through a flexible infrastructure approach. 'Federated' infrastructures are those constructed from IT resources assigned by autonomous, cooperating, business parties within and beyond the enterprise. Although University departments all belong to a single organisation they each maintain, and jealously guard, a considerable amount of autonomy to the extent that the University may be considered to be a federation of departments, at least with regards to IT.

Agility™ is being used at Newcastle to provide intelligent power management for the University's Condor cluster whilst respecting the federated nature of the organisation.

Condor is a popular grid computing technology typically used for parallel execution of compute intensive tasks. IT resources (generally, in the case of the University, desktop PCs) register themselves with Condor when they are 'spare' i.e. not performing their regular tasks (typically at the end of office hours) and deregister themselves when they wish to return to their regular tasks (typically at the start of office hours). Users, who wish to have compute tasks executed, submit them to Condor which then distributes them over the available resources for execution. At Newcastle University there are frequently in excess of 2,000 machines within the cluster.

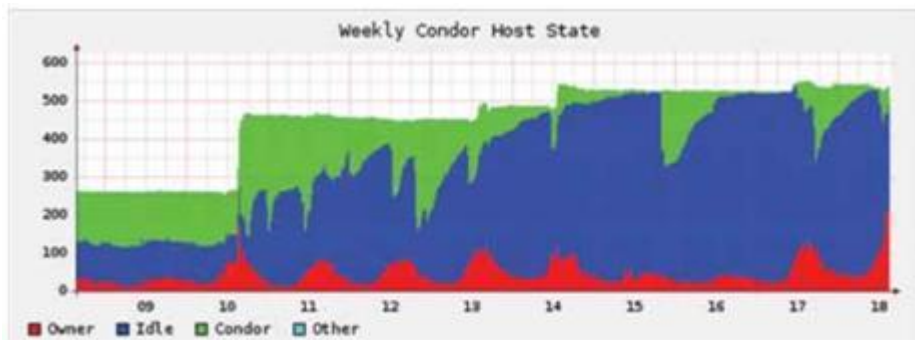


Fig. 1 Condor utilisation

Agility™ installation

Agility was deployed in order to allow the University to introduce the policies necessary to intelligently power manage the IT resources within their Condor grid. Agility enables the introduction of policies which can gather information in order to improve accountability, and to enable intelligent decisions with regard to power management. The deployment was simple to achieve and did not impact upon Condor's users for whom the interface to Condor remain unchanged. No changes to Condor were required.

With Agility in place initial Condor Gateway policies were introduced to enable Condor tasks to be recorded as they pass through the system. An administrator may then access a console which displays, for each task, the identity of the person submitting the task, the details of the task, when the task was submitted, when it completed, the outcome of the task, the resources utilised, and the power (wattage) consumed. The data may be presented so as to illustrate the power consumption over periods of time for individual users, projects or departments. By itself this feature goes a long way to improving accountability.

Management with Flexible Policies

In addition, Agility has allowed the University to introduce Power Management policy designed to significantly reduce power consumption in two ways. Firstly, the policy modifies each Condor task with an attribute that makes Condor take into consideration the Power Efficiency Rating for the machines present in the cluster. As a consequence Condor will assign work to the available machines with the lowest power consumption (with a power saving compared to older, power-hungry machines of anything up to 80%). Secondly, the policy looks at the identity of the submitter and if they are identified as a student the task is marked as requiring an already powered-on machine. This prevents Condor from powering up a machine for this class of lower priority user.

Both of these modifications represent relatively simple policies but Agility allows other more complex policies to be added to (or removed from) the system at run-time. For example, policy could be introduced to set (or check) priorities on tasks based upon the identity of the user, could make estimates of the wattage which will be consumed by a task and then decide to postpone or cancel a task if the costs are judged too high, could abort low priority tasks if particularly high priority tasks were submitted. Policies can also be introduced which take into account changes to carbon footprint legislation as and when these apply.

The strength of Agility is the ease with which new policies may be added to the system, each hopefully improving the overall efficiency of the system, one step at a time, without disrupting existing use. These policies can be concerned with any aspect of managing the IT infrastructure and are in no way limited to issues around power management alone.

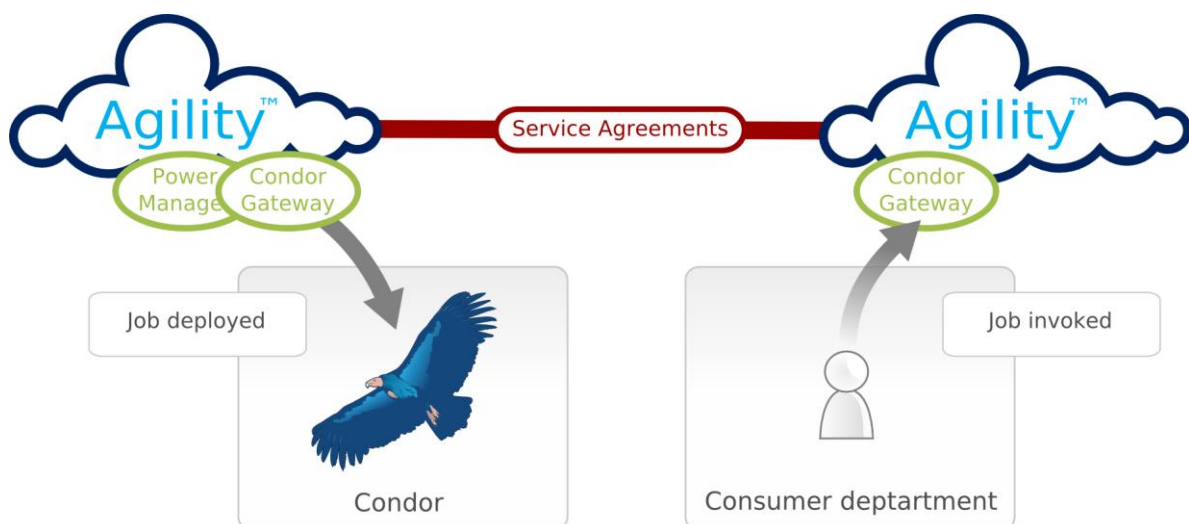


Fig. 2 Policies for Condor

Future Proofing

By deploying Agility the University has future-proofed itself against changes in business and IT Infrastructure. Agility supports flexible and dynamic policy supporting business and operational changes, and can support any IT infrastructure providing it can be programmatically managed. Newcastle University is therefore in an excellent position to deliver access to other clusters/grids or clouds, as and when these become available, all managed through a common interface.

Agility allows multiple domains to be created within the University, for example one per department. Each domain can deploy its own Policy to control its own IT infrastructure and to enable sharing with other departments in a controlled and accountable way. Sharing between Universities can also be facilitated in a similar way. As an example a Condor task submitted at Newcastle could be redirected through a Brokering policy to another instance of Agility running at a different University and submitted to their Condor cluster. This could only happen if both Universities have implemented policy which allows this form of sharing. Through the use of Service Agreements, Agility would retain a full record of the task which could be used to ensure quid quo pro sharing, or even for cross-charging.

Obtaining additional capacity from a Public Cloud also becomes feasible through Agility. Once again this would only be allowed if University policy permitted it. Ultimately a cost/benefit decision can be made on a job-by-job basis as to from where the most appropriate capacity can be obtained.

