# CMS Job Efficiency at the RAL Tier-1

## Introduction

It has been noted that globally, CMS Job Efficiencies (CPU-Time/Wallclock-Time) are lower than some other VOs, and that since July 2020 the lowest efficiency has been at the RAL Tier-1. The RRB interest in this matter gave rise to a WLCG MB action for RAL to provide this report.

## CMS Job Efficiencies at RAL

Our analysis of CMS jobs shows that the low efficiency arises from two primary causes:

1) Poor data access speeds for jobs requiring off-site secondary data.
2) I/O contention on older generations of Worker Nodes.

## 1. CMS Off-Site Reads

The CMS computing model takes advantage of available network bandwidth to reduce site storage requirements using the AAA federation of resources. Simply put, some CMS simulation jobs pull in significant external secondary data (minimum bias overlay) that are added to simulated physics channels. The secondary data is purposefully not stored locally because of size. However, this reduces the demand on local storage at the expense of increased network activity. Our analysis of CMS workloads at RAL show that whilst less data intensive jobs have a CPU efficiency comparable with other Tier-1s, the CPU efficiency of jobs reading secondary data from off-site is significantly lower at RAL than other Tier-1s. This can be seen in Figure-1 below. In particular, the lower plot shows the efficiency over a 3-week period for (green) jobs with no secondary data; (red) jobs with all data on-site; (dark blue) jobs with secondary data off-site in Europe; and (light blue) jobs with secondary data outside of Europe.

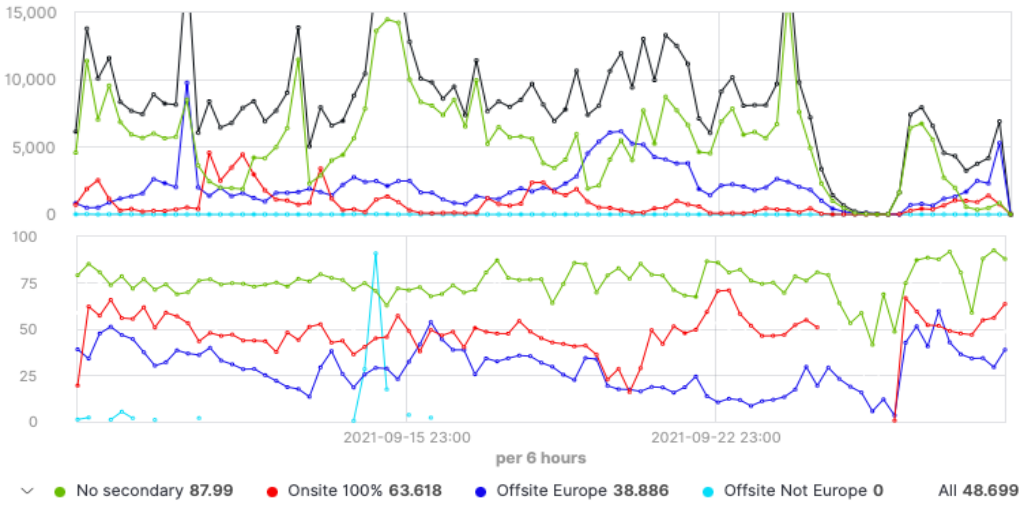


*Figure 1. The top plot shows the number of jobs completing for CMS at RAL during a 3-week period in September 2021. The bottom plot shows the CPU efficiency of those jobs in the same period. Green Line – No secondary data. Red Line – Primary and Secondary data onsite. Dark Blue Line – Primary data is at RAL, secondary in Europe. Light Blue Line – Primary data is at RAL, secondary is outside Europe.*

It is clear, therefore, that whilst the RAL Tier-1 network has been adequate, so far, to meet the needs of the other LHC VOs, it did not meet the more demanding (instantaneous) requirements of the CMS AAA model and was contributing to increased inefficiencies for some CMS workloads. A substantial network upgrade is currently almost complete; this is described later.

## 2. I/O Contention on older worker nodes

The other problem impacting CMS CPU efficiency was I/O contention on older Worker Nodes purchased before 2019 without SSD backed storage. Due to the way multi-core jobs are scheduled at RAL, it was frequently possible for entire nodes to be filled with CMS jobs. A simple test was performed where we measured the CPU efficiency as a function of the number of jobs running on the node: Initially the CPU efficiency was very high (in line with expectations) however by the time all the physical CPUs were in use the CPU efficiency was < 25%.

Newer hardware is more than capable of handling the more data intensive workloads and there are plenty of less data intensive workflows run by both CMS and the other LHC VOs that can run fine on older nodes. The old generations of hardware have now been configured to run fewer jobs and we have configured the batch farm to try to avoid scheduling multiple data intensive workflows on the same node.

## RAL Site Network

The current RAL Tier-1 network can be thought of as 3 separate subnets, each with a 40Gb/s link to the outside world. Up until April 2021, the bandwidth to the subnet where the worker nodes are located was limited to around 15Gb/s due to the site firewall (which is external to the Tier-1). The firewall was upgraded in April 2021 and since then there are spikes of much higher bandwidth, as shown in Figure-2.

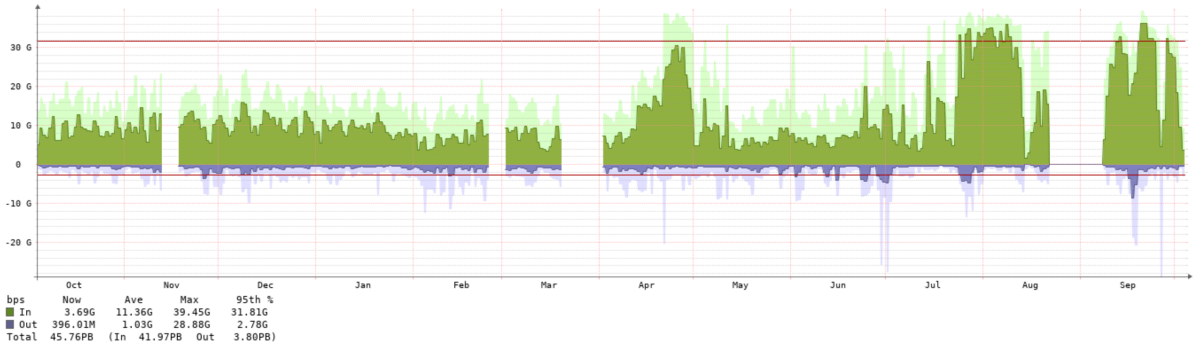


*Figure 2 showing the network traffic to the RAL Tier-1 Worker Nodes over the last year.*

In addition, as new worker nodes (with SSD storage) were brought online in July 2021 the network link started frequently to saturate. As well as these network saturations, it has been observed that, when the number of connections reaches a certain level, many transfers start to fail.

In the second half of 2019 a relatively modest project to upgrade the RAL Tier-1 LHCOPN link and the relevant internal routers was started. This was interrupted by the start of the pandemic and with the subsequent CMS CPU efficiency issues and developing predictions about the HL-LHC data rates, it was clear this upgrade would not be sufficient. Therefore,

since September 2020, the RAL Tier-1 has been building a completely new network. This replacement network will have significantly higher bandwidth; avoid the site firewall; and will greatly simplify the logical layout of the network. The details of the design are presented in a vCHEP21 paper[1].

As of October 2021, the new RAL Tier-1 network is close to completion. Most of the physical cabling was completed by May 2021, and work on commissioning the links is underway. The new tape service Antares has been built entirely on the new network and all disk and CPU purchased since 2019 is scheduled to be connected. The remaining links should be commissioned in October and then the new network will be advertised to the LHCOPN/ONE. The new network will remove the current bottleneck and should result in significantly improved off-site reads for CMS. Of course, it is always possible that the removal of one bottleneck reveals another, but we are confident of resolving any residual issue[2] before the start of Run-3.

## Conclusions

The lower CMS efficiency at RAL is understood to be related to the demands on the network of the AAA model for off-site reading of secondary data, and due to I/O contention when multiple CMS jobs were scheduled on older hardware. A completely new network has been provisioned and is currently being commissioned. All new hardware since 2019 has been backed with SSD storage and will be able to meet the I/O requirements of CMS. We are confident that these actions will bring the RAL CMS efficiency in line with other Tier-1 sites.

[1] https://doi.org/10.1051/epjconf/202125102074

[2] We may try turning off CMS Lazy-Download, which in principle could cause increased network traffic. However, it is currently needed internally to optimise I/O from ECHO; it is also used by KIT, where it does not seem to cause any significant problems. It might, for example, be possible to use it internally but not externally at RAL.