
High-Performance Computing & Formula 1

by Paul Ingram - Principal Consultant



Red Oak Consulting
Expert advice, exceptional delivery



Our team of consultants have
combined experience and knowledge
spanning over 90 years.

Contents

Introduction	3
Formula 1 – Speed, Agility and Performance	4
HPC in Formula 1	5
What are the HPC Workflows in Formula 1?	6
HPC setting the direction for a season	7
What Could the Future Look Like?	8
Cloud Computing in F1 - YES	10



Red Oak Consulting
Expert advice, exceptional delivery



Introduction

Red Oak Consulting continues to provide high levels of skilled, expertise to businesses in the supply and solutions for all matters relating to HPC. Over the last 17 years we have completed over 100 HPC projects, over 30 HPC procurements and over 25 Cloud Computing projects.

Our team of consultants have combined experience and knowledge spanning over 90 years. Having worked across a diverse range of business sectors we continue to apply our unrivalled knowledge to each and every customer engagement.

Paul Ingram, Principal Consultant, has worked in HPC since 1999 and has successfully built, tested, and installed many HPC systems of varying complexities across academia, research, and government sectors in the UK. Through the course of Paul's highly successful career Paul has acquired the niche expertise for every aspect of on-premise HPC including; Linux, design, storage, benchmarking, vendor liaison, UAT and FAT's, problem solving and project management.

For the past three years Paul has worked in the highly acclaimed, prestige, corporate sector of Formula One. Here, he has managed HPC systems running aerodynamic(CFD) and Finite Element Analysis system workloads, race strategy systems, and vehicle dynamics systems, all of which led to three consecutive world championship wins.

This direct insight has shown how reliant Formula One is on HPC. With the majority of HPC being sourced via on-premise this paper looks at why HPC is so critical and what the future could possibly look like. A key question is does the future allow for more usage of HPC in the Cloud?

For the past three years Paul has worked in the highly acclaimed, prestige, corporate sector of Formula One.

Formula 1 – Speed, Agility and Performance



Formula 1, an industry many of us have a view or an opinion on, but do we really know what goes on behind the scenes. From the outside looking in we see an industry built on money, glam, drive and passion, but what really makes Formula 1 tick. Some may say that Formula 1 is an edge case organisation as it does not fit with other profit-and-loss Automotive manufacturing businesses. Why, because in mainstream automotive survival in the industry relies on selling as many vehicles as possible. Without sales there is no financial viability.

As we know, the purpose of a Formula 1 team is not to sell cars but to win the F1 world constructors and drivers' championships. Formula 1 is a very expensive sport, especially when it comes to the building and maintaining of the race car. There are a lot of stories in the press regarding the cost of an F1 car for example the front wing, or the cost of a F1 engine, but the cost of these parts compared to road-going cars is irrelevant when you consider that the infrastructure required is to build in essence two cars a year. A worthy note is there are more cars built each year, but one team can only race two cars at a time.!

F1 cars are hand-built as opposed to being built on a production line. Therefore, F1 car parts could not only be considered custom-built, but also built to the highest manufacturing standards and tolerances, whilst using the best quality materials in terms of strength and weight.

It is accepted that to design, build and race an F1 car is a very expensive business. To pay for these cars to be made and raced can be summed up in one word – sponsorship. It is sponsorship that funds the whole F1

organisation. And sponsorship can feed in to the HPC Solution. Normally, in F1 terms, success breeds success so a team that is doing well on the track is more likely to attract the big successful companies that want their name to be associated with a winning Formula 1 team.

Over recent years the governing body of F1, the FIA, have brought in rule changes to level the playing field. For previous years it has been a case of the bigger the team the bigger the budget therefore they are able to spend, spend, spend, to make their car the fastest on the track.

The cost cap rule that has been recently brought in has had the biggest impact on how each team is funded. The top four F1 teams are having to make significant changes in their operational activities to bring their spending down within the limits of the cost cap. The bigger teams have had a choice to make in this regard; they can either trim their business down to bring themselves within the cost cap or diversify and use their resources in other ways outside of F1. In order that teams do not lose access to high value specialist machinery, and more importantly a highly skilled workforce with many years of experience working in an F1 environment, most teams have chosen the diversification option.



Red Oak Consulting
Expert advice, exceptional delivery





HPC in Formula 1

HPC is a necessity within F1. It is a critical tool, one which has been used for over 20 years. Without High-Performance Computing the world of F1 would be very different, possibly even non-existent. The first use of HPC was for CFD (Computational Fluid Dynamics). CFD in basic terms is used to simulate the flow of air and liquid through mathematical modelling and the effect this flow has on the surroundings (i.e., car bodywork). This process can be applied to a range of uses such as a car exhaust system, a lubrication system, or the aerodynamic shape of a car. It is this area where the most powerful HPC systems are used.

In the early days, the problems that are now solved using CFD used to be solved in other ways. The most well-known and one that is now used alongside CFD is the wind tunnel. Where the wind tunnel was first used in aviation to measure the lift of a wing shape, it is now used in F1 to measure downforce, the method of keeping the car firmly on the track to maintain good grip from the tyres. And as we all know in F1 the holy grail is to have a car with high downforce but low drag.

Although the advent of CFD has brought down the reliance of the wind tunnel to develop an aerodynamic shape, it is still used to prove the simulated results

that CFD produces. In terms of Aerodynamics, the CFD side (computer simulated) and the Wind Tunnel side (actual Aero using a 60% scale car in a wind tunnel environment) are governed by FIA regulations limiting the amount of work that can be done in the wind tunnel (limited on the number of hours the Wind tunnel is operational) and CFD is limited on the number of FLOPS used in a particular reporting period. There are also limitations on the type of compute you can use for CFD work. GPU's are not permitted for CFD and turbo mode must be disabled. These are the two main limitations. There are two other aspects of Aero which are not governed by FIA regulations which are meshing and pre/post processing. Meshing is used to divide a geometry into many elements which is then used by the solver to represent flow regions. Pre- and post- processing are basically data conversion exercises used to prepare data sets prior to solving and then converting results into graphical representations of the model in terms of air flow and pressures. Aero engineers then use these representations to validate that a particular change has been successful.

What are the HPC Workflows in Formula 1?

To calculate the effects that airflow has over a car bodywork, such as drag or downforce, can take a lot of compute power. For just a single car aero CFD simulation run, using the most powerful home computer on the market today would take up to 14 days or longer depending on the complexity of the run. For a fast-paced formula 1 environment, this is far too long and where HPC takes over. These days, most people are aware that a HPC system comprises several compute elements or nodes that are not dissimilar in basic design to a normal, reasonably powerful, home computer system. For a HPC system to be effective, the same single car aero CFD simulation that took 14 days on a single system, would need to run in a fraction of that time across several systems, or nodes. This is known as scaling. If the same code that ran in 14 days could be run in parallel across 14 computers and take 1 day this would be 1:1 scaling. In practical terms this is difficult to achieve as several factors would need to be considered, but to get near to this for most workloads would be acceptable. (see Amdahl's law)

Other HPC areas of F1 are run restriction-free, such as Finite Element Analysis (FEA). FEA is like CFD in that it is based on mathematical modelling. With FEA the emphasis is on structural analysis rather than fluid dynamics.

FEA codes can run on a similar design HPC to those that CFD runs on. However, since CFD is closely regulated there will be separate HPC systems for CFD and FEA codes.

In F1, FEA is used for designing car components where strength and weight are critical. For example, a suspension component will need to have certain physical qualities to perform its function reliably and safely, whilst at the same time not causing a burden to the car in terms of excessive weight.

At the start of an F1 season, each team's car must undergo a series of crash tests to ensure the car construction is safe to compete. The way the car construction behaves when subjected to the stresses that sudden impacts can inflict will be critical to ensure the car passes these tests. The behaviour of the car construction when undergoing these tests can be simulated using FEA. This can save considerable time and money in developing a safe car structure without having to destroy a lot of car parts in the search for a safe design.

Creating a championship winning car is not just about having a fast car with good aerodynamics. As F1 has evolved, teams have found that attention to the smallest detail can make the difference between finishing first or second, or scoring points, and not scoring points. No two races are the same, even if they take place on the same track; there will always be differences in weather conditions, wind speed and direction, and also which tyres to use and when to change them during a race. There will also be events such as engine failures, crashes etc. How all of these events and variables are handled by each team falls in to two main categories, each having their own compute requirements; car setup and race strategy.

Car setup generally is preparing the car for a particular track. Each track on the racing calendar has its own characteristics. For example, Silverstone and Monaco are two very different circuits and you would prepare the car accordingly. Silverstone being one of the longest circuits in the F1 calendar has a lot of straights so top speed will be an important factor. To make the car faster, the downforce will be reduced. This will have the effect of increasing the top speed, but cornering efficiency will be less. The Monaco circuit, however, is the total opposite, being one of the shortest circuits; Not to mention being a street circuit which is narrower with more corners, impacting the need to be taken at a much lower speed. For this circuit, cornering performance is critical so the car will be set up with more downforce.

HPC setting the direction for a season

So how is this achieved? Answer, simply through HPC. Increasingly, the car setup is simulated in HPC. There are so many different permutations in a car setup such as: suspension settings, ride height, how much downforce/ wing level to apply, tyre pressures, brake, and brake cooling settings. These are just a few examples of the number of variables involved that all have an impact on car performance and reliability. The effect of these settings both on the car and with each other is simulated using a HPC based system.

Race strategy is closely related to car setup as quite often changes are made to the car setup for strategic reasons. Generally, a race strategy comprises with a set of decisions made before and during a race to maximise the potential of the car. It is possible to simulate an entire race in a fraction of a second using a specifically written application across HPC cores. This is known as a “Monte Carlo Simulation”. An entire race is simulated thousands of times, each time with a slightly different set of parameters such as weather conditions, tyre changes, any safety car deployments, and other variables during the race. All this data is then compiled and analysed to provide a probability figure of a given outcome for a particular series of decisions. During a race, an event can occur and the response of the team to that event can make the difference between winning the race or finishing down the field, irrespective of if you have the fastest car or not. This is how important race strategy is.

Over the years, High-Performance Computing has predominately been facilitated on-site at each team HQ. As you would expect to find in most organisations where HPC is used, the normal life expectancy of a system will be between 3 and 5 years. After this time Moore’s Law would normally dictate that by then there will be faster and more efficient hardware available which will necessitate replacement of the system. Depending on the team, and if it has any partnership agreements in place with a particular supplier, systems will be supplied by HPE, IBM and Intel, among others.



What Could the Future Look Like?

So, questions are being asked, should HPC remain an On-Premise solution for F1? Could there be a better way of providing the Solution? Can Cloud cope in the world of Formula One? Does Cloud have the right level of security for Formula One?


It is fair to say, Cloud computing has a role to play in F1. Many of the current teams within F1 have already started using Cloud based services to some degree. The FIA used a cloud based HPC system to simulate the effect of the 2022 Aero changes. The 2022 changes are being brought in to reduce the effect of backwash that the current car aero design produces on following cars. This makes it difficult to overtake. The aero changes deflect the backwash up and over a following car, rather than right into its path. The idea of this is to make overtaking easier.

In terms of HPC, apart from the areas of CFD, for obvious reasons, many HPC workloads can be migrated to the cloud, to some degree. Below are areas that could be looked at with a view to producing some POC systems for analysis:

Agile workloads

Microsoft's Cycle Cloud product enables the building of a Cloud based HPC cluster with built-in scheduler and with high-speed interconnect (if required) within a matter of minutes. This would be particularly useful with an unexpected workload cropping up where all the current on-premises HPC systems are fully utilised. Where now some F1 teams are diversifying into non F1 projects, and where these projects require HPC work, they can be up and running almost immediately and without the long lead time procurement cycle usually associated with HPC systems. In addition, where a team has several customers, some who might be in competition with each other, and wish their IP and data to be kept away from any other customer, Cloud compute services offers this facility by default. It is much easier to separate and segregate Cloud compute resources than it is on-premise compute resources.

It should be noted that with the restrictions on the type and quantity of CFD solve work that can be done, it may not be possible to run the solve part of CFD in the Cloud. Turbo mode and any GPU compute work is not allowed for solving; however, they can be used for other aspects of CFD. This restriction may mean that the solve element of CFD may not be possible to be run in the cloud.



What Could the Future Look Like? Continued...

Hybrid systems

If a team has a compute cluster that is not performing to expectations or its use-case has grown and the system is having difficulty coping with the increased workload, it is possible to add virtual Cloud-based nodes to the cluster and “burst” into the Cloud when the on-premises system is running at maximum load and needs more compute resource. This makes the best out of both worlds, by still using the on-premises machine and only using the Cloud resource when it is really necessary.

Periodic HPC usage

Where the traditional on-premises HPC runs to full-load wherever possible to make use of the capital invested in it, where you have a use-case that requires HPC level compute power but is only run periodically, say for example 2 or three times a week for short periods, this would ideally fit into a cloud environment.

Project-based workloads

Where a team has a non-F1 customer paying for the benefit of the technical expertise that an F1 team possesses, where there is requirement for some HPC work, it would be far better to spin up a temporary Cloud-based HPC system than piggy-back on to an existing system. Once the work is completed, the cluster can be spun down and deleted.

You can see from this list that there are many areas where HPC Cloud-based solutions can be offered and trialled. The beauty of cloud computing is you do not have to buy hardware, run some tests, find that the tests are not as performant as you hoped they would be, and then realise you have

not bought enough, or you have bought the wrong hardware. With Cloud you can try a VM type, run some tests and compare those results with a different type and see which one works best for you. You can do this with interconnect options too. Cloud offers you the flexibility to grow and shrink your HPC resource and move to the latest compute technologies as and when they become available, and not have to wait for next year's budget allocation.


As mentioned at the beginning of this piece, I spent over three years working in an F1 environment so understand the challenges and problems that are faced by the team IT departments on a day-to-day basis, particularly now in the cost cap era. To survive without making deep and severe cuts to infrastructure and more importantly, a highly skilled workforce, changes to ways of working are needing to be made.

Cloud offers you the flexibility to grow and shrink your HPC resource and move to the latest compute technologies...



Red Oak Consulting
Expert advice, exceptional delivery





Cloud Computing in F1 - YES

Cloud computing can offer quick deployments of an HPC resource when it is required. Where a team has diversified and taken on non-F1 project work, cloud can be used to provide IT resources for the life of the project and then decommissioned. This negates the need to make costly hardware acquisitions that are only needed for a short period. Where a team does not have the spare budget to replace their HPC system, it's life can be prolonged by adding cloud VM nodes to the compute capacity hence converting it into a hybrid system. The relative cost to use a Cloud based HPC for a short period of time compared to the capital expenditure required for an on-premises HPC and all the related and required paraphernalia to keep it running can be substantial saving.

Cloud computing certainly has a place in F1. Whilst some companies are using Cloud to a small degree, today, it is fair to say there is more scope to use more. Cloud usage will grow, undoubtedly, but it is fair to say it is early days. If you would like any further information or are interested in knowing how Red Oak Consulting can assist you then please get in touch.

To find out more about our services and how we help our clients harness the full potential of HPC, contact an expert today.

Speak to an expert

Visit our website: redoakconsulting.co.uk



Red Oak Consulting
Expert advice, exceptional delivery

