



Optimising for a changing environment

Ellexus Mistral takes I/O profiling for genomic pipelines to the next level in order to maximise performance in rapidly changing public cloud environments



“Our results show that Mistral easily identified the I/O issues underlying the genomics pipeline tested in the cloud. Creating a real HPC environment in which to profile means we can take what is already an excellent pipeline and make it even better.”

Rosemary Francis, CEO, Ellexus

Customer Profile



Company	Ellexus
Industry	I/O Profiling
Country	United Kingdom
Website	www.ellexus.com

Business Need

The team at Ellexus wished to create an environment to test the capabilities of their product, Mistral, against a typical high performance compute (HPC) application for the public cloud.

Solution

Ellexus chose a well-maintained, cloud-ready genomics pipeline and partnered with Alces Flight to quickly create a HPC cluster environment on Amazon Web Services (AWS). With Alces Flight the team at Ellexus were able to use a container holding the target genomics pipeline to be profiled with Mistral and yield performance metrics to be fed back for continued pipeline improvement.

Benefits

- Creating a predictable and repeatable HPC environment for real world testing
- Allowing containers to be used alongside ISV applications to yield beneficial I/O profiling results
- Providing specialist software providers the confidence to assist in migrating established on-premise tools to cloud environments
- Maximise the ability to enable I/O performance measurements for HPC clusters using Flight with Ellexus Mistral for software applications

Solution at a glance

- I/O Profiling
- Storage Agnostic
- Always-on Monitoring
- Container Monitoring

Ellexus is the I/O profiling company. With a passion for providing the lightest to the most detailed analysis in reporting the team are dedicated to developing solutions for all I/O profiling needs. Their latest effort? Maximising performance in cloud high performance computing (HPC).

“In cloud time is money” explains Rosemary Francis, CEO of Ellexus, “From the moment you request resources the charges begin. If you have an errant application that is running out of control you’d want to know where it is and how to stop it before the bills get too high. It’s issues like these that guided us to take on cloud HPC as our next point of focus.”

In 2016 Ellexus launched Mistral, a tool focused on monitoring application I/O and cluster performance so that jobs exceeding the expected I/O thresholds can be automatically identified and slowed down through I/O throttling. After an award-winning first year successfully profiling on hardware the team looked to expand their capabilities to cloud.

“Cloud environments can change in minutes,” Francis said, “in cloud you can theoretically have good performance one day and bad performance the next. We wanted to design an I/O profiling challenge that dealt with the difficulties involved and see how Mistral handled the pressure.”

Creating a real environment and coping with containers

In order to test Mistral’s capabilities the team at Ellexus required a cloud service provider, a platform, and an application. “We chose AWS as they are the current market-leader and through this were introduced to Alces Flight. With Alces we were able to set-up and launch our HPC cluster in minutes which meant it was all about the profiling challenge from that point on,” said Francis.

For the I/O profiling test Ellexus took the decision to take a highly maintained, publicly available genome pipeline from a well-respected cancer research organisation. “This genomic pipeline was being held in a Docker container,” Francis said. “By placing the entire environment inside a container you get as close to an operating system as possible without needing a virtual machine.

As Flight already handles containers all we needed to do was create a script for Mistral and run the Docker image. To push Mistral’s abilities to search out even the smallest of improvements we configured the tool to get the most detailed analytics possible.”

Ready for flight

The team at Alces saw the I/O profiling challenge as a means to demonstrate the strengths of both Flight and Ellexus. “We are just as concerned about optimisation and performance as Ellexus,” said Wil Mayers, Technical Director at Alces Flight. “To be able to create HPC clusters in minutes takes a huge burden off of the user who would otherwise be building from scratch. With applications, optimisation is everything. We find that, on average, Flight gets you about 80% there. You’d be surprised how much work goes into that last 20%. Adding Ellexus gets you closer to that elusive 100% by continuously looking for more ways to improve.”

The team at Ellexus completed their test in under 11 hours. “We found that changes to small writes and reads could bring about more time and cost savings in the long run,” said Francis, “From this already highly tuned pipeline Mistral could see that room for improvement that we know will always exist in cloud HPC. We’re very pleased with the results and we look forward to bringing more speed and efficiency to the cloud in the future.”

“In cloud HPC what could be good performance one day can be bad performance the next. This is the nature of the constantly changing environment and a key reason I/O profiling exists.”

Rosemary Francis, CEO,
Ellexus

Products & Services

Services

Customised set-up of HPC cluster

Shared knowledge on cloud HPC

Products

Alces Flight Solo

Cloud Service Provider

Amazon Web Services

View all Alces Flight success stories at alces-flight.com/success

The Alces Flight Connection Program

Ellexus has partnered together with Alces Flight under the Alces Flight Connections Program. This program is designed to bring ISVs for HPC into the cloud at a rapid rate by allowing them to test their tools within the Alces Flight platform on Amazon Web Services (AWS). Through continued collaboration work with ISVs it is our intent to expand the capabilities of cloud HPC by bringing cost effective solutions from trusted providers to the end user faster than ever before.

