Ubiquity: An Open-Source Platform For Ubiquitous Large-Scale Computing

Chris Coates, Head of Engineering, HPC&AI - Logicalis UKI

Abstract

There are significant challenges that exist in the widespread adoption of Large-Scale Computing, many of which revolve around an increased complexity reducing the effectiveness of staffing.

Because of this need, Logicalis has developed in collaboration with various organisations a solution called Ubiquity.

Ubiquity provides a rapidly growing open-source and open-development platform designed for Large-Scale Computing, which leverages the concepts of Cloud Computing, Converged Computing, and Ubiquitous Computing.

This platform utilises container and container orchestration technologies to enable efficient and scalable solutions.

The results of this project have been impressive, and as a result Ubiquity has garnered an increasing level of adoption and has extended its support to diverse HPC infrastructures beyond the field of scientific research.

Industries such as financial services, energy, manufacturing, and life sciences have found value in leveraging Ubiquity to enhance their computational capabilities.

Introduction

Despite the immense potential and opportunities presented by Large-Scale Computing in accelerating business innovation and scientific research, several challenges exist that hinder its widespread adoption and optimisation.

At the heart of this lies a widespread lack of awareness and specialised expertise necessary to operate and leverage HPC and AI systems efficiently. This gap is further exacerbated in the academia, where the roles of software developer and researcher often converge, leading to sub-optimal software development practices, poor testing procedures, and inadequate software packaging/archiving protocols.

There exists a crucial need for an HPC solution that can address these multidimensional challenges, providing an accessible, efficient, and sustainable platform for users to harness the power of large-scale computing across hybrid environments.

Logicalis has invested in the future of these environments and is giving back to the HPC community at-large by developing a solution to this in the form of a software stack/framework called Ubiquity.

Ubiquity is a powerful, autonomous platform designed to accelerate and simplify the deployment and management of HPC resources. It is designed around the concept of Ubiquitous Computing, which aims to make the complex, easier to use and administrate.

The Ubiquity project delivers a stack that is fully open-source and platform-independent, enabling rapid deployment of containerised HPC environments across both on-premises and cloud environments.

Components/Requirements

Ubiquitous Computing, was coined by Marc Weiser of Xerox PARC in 1991:

"My colleagues and I at PARC believe that what we call ubiquitous computing will gradually emerge as the dominant mode of computer access over the next twenty years.
Like the personal computer, ubiquitous computing will enable nothing fundamentally new, but by making everything faster and easier to do, with less strain and mental gymnastics, it will transform what is apparently possible."

This concept applied to Large-Scale Computing means:

- Reduced overhead to use day-one
- Reduce deployment time (90%+ required)
- Enable mobility of research
- Be secure by-default without impeding users
- Significantly reduce overheads to administrate day-to-day
- Improve coding workflows
- Improve observability of performance
- Reduce touchpoints to make infrastructure changes

To achieve this, we have leveraged specific components – But the ultimate goal is to be autonomous wherever possible

Methodology

To deliver this functionality, we have referred to our deployment in the concept of "layers" – Key components layered, that enable other layers above to be more effective.

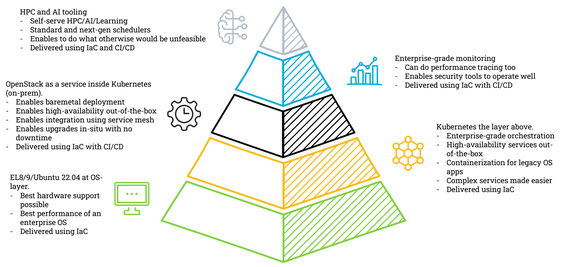

These concepts have been developed in conjunction with Intel and Lenovo, however further research was undertaken with national-scale infrastructure providers on how they address their environments in order to work on the OpenStack-Within-K8s paradigm.

This paradigm allows us to manage a huge range of bare-metal environments, alongside being able to use Kubernetes to provision underlying infrastructure at-scale, easily.

The command-set by using operators within Kubernetes means only one command-set needs to be understood – Kubernetes.

And even then? Ubiquity includes toolsets to reduce the overhead for that too.

Results

Our results speak for themselves:

- Deployed entire environment for key science research outcomes within 10 days end-to-end.
- Achieved a 46% performance uplift vs. higher-specified hardware, independently benchmarked on a hardware vendors own software stack.

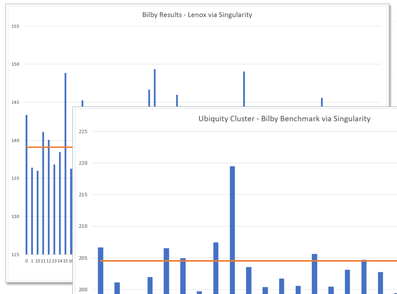

An average of 139.10, Lowest being 134.12, peak being 149.38 – A maximum worst-case scenario of ~11.3% variance (best-to-worst) across all cores on the node, of which there were 64.

An average of 204.504, Lowest being 197.91, peak being 219.75 – A maximum worst-case scenario of ~11% variance (best-to-worst) across all cores on the node, of which there were 24.

- Environment shutdown/startup time now reduced to <30 minutes with zero intervention.
- Automated code analysis via Ubiquity workflows delivers a significant performance uplift:

```
$ CC=cc RUNS_WARMUP=5 RUNS=10 ./benchmark-omp-multi.sh
...
Code       Original  Optimized  Speedup
========== ========  =========  ==============
ATMUX      0.29      0.08       72.49% (3.64x)
CANNY      6.23      4.55       26.97% (1.37x)
COULOMB    10.59     0.36       96.61% (29.51x)
HACCmk     34.97     1.60       95.43% (21.89x)
MATMUL     3.19      0.12       96.20% (26.29x)
NPB_CG     26.50     5.37       79.75% (4.94x)
PI         4.15      0.13       96.85% (31.76x)
```

- Self-serve functionality allows entire learning environments to be spun-up within minutes from a catalog of software, all controlled via GitOps
- Leveraging eBPF to monitor syscalls of software alongside full network segmentation enables security posture to be maintained whilst users are unimpeded by such technology. The performance hit of this is less than 0.5%.
- This same functionality also allows for live performance profiling of software, with minimal overhead on the software itself.

Conclusion

Logicalis and their investment into the future of Large-Scale Computing has developed the concept of applying Ubiquitous Computing Concepts towards HPC and AI environments.

As such, the performance and usability improvements alongside significant security improvements and reduced overhead to both implement and maintain make a Ubiquity platform significantly more scalable – A key concern as we move into the exascale era.

The results we achieved during the continual ongoing development of this project will continue to evolve and improve over time.

Ubiquity as a stack will be formally released before SC23 – Under a free open-source license (Apache 2.0).

Logicalis welcomes collaboration with organisations to help usher in a new era of collaborative design and technology improvement to unlock the transformation of what is apparently possible.

Recommendations/Roadmap

Going forward, this project has on its roadmap already:

- Multi-cloud refinement – Because of the ability to deploy the same underlying core K8s environment, the remainder of the infrastructure can be consistent. Axure and AWS are already onboarded, GCP is underway alongside OCI and OVH.
- Make the physical part of deploying a cluster simpler – Augmented Reality project already underway
- Leverage AI to provide root-cause analysis of system functionality/performance issues – By ingesting system logs and creating agents to do investigation and fixes and continually learn via reinforcement, the system should become more efficient at maintaining and fixing itself the longer it is in-service.
- Continue to refine and improve automation pipelines to ensure updates are continual.
- Continue to integrate storage providers into Ubiquity.