

## Performance Analysis of Applications using Singularity Container on SDSC Comet

Emily Le<sup>1</sup>, David Paz<sup>2</sup>

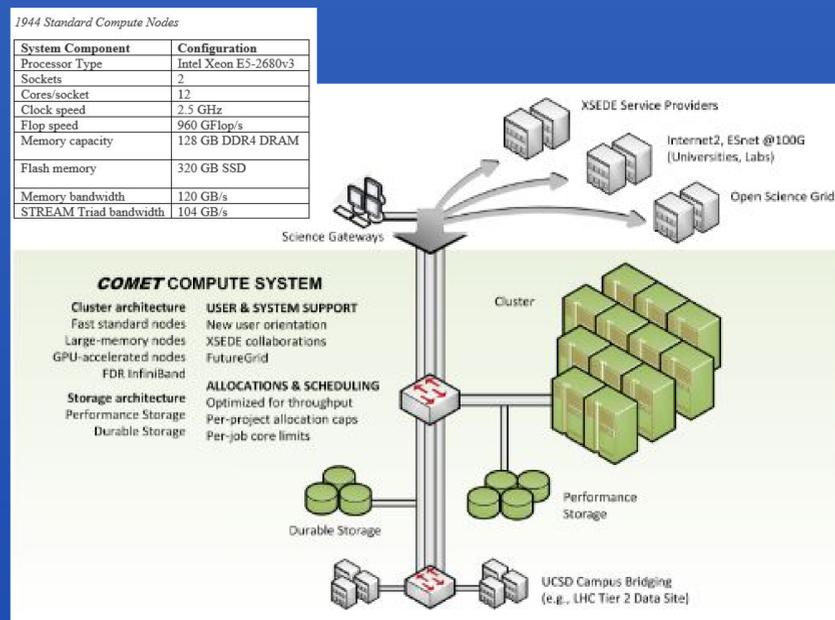
<sup>1</sup>Mathematics Department, <sup>2</sup>Electrical and Computer Engineering, University of California, San Diego

### ABSTRACT:

This study aims to analyze the properties and HPC performance implications of Singularity to determine if container benefits amortize the overhead cost. Multiple nodes on the Comet supercomputer at the San Diego Supercomputer Center were used to analyze and compare (between Singularity and non-Singularity) performance of resource intensive applications and benchmark codes such as NEURON and Intel MPI Benchmarks (IMB). NEURON software was used to simulate a complex neural network model, and the OSU and IMB benchmarks were used to calculate the latency and communication efficiency for MPI. The containerized runtimes were directly compared with the corresponding non-containerized runtime of jobs to analyze the performance of each method. For future work, we plan to explore other technologies such as Shifter and study their performance on HPC.

### INTRODUCTION

Container technologies have been widely used in recent years due to their applications and benefits in version control, predictability, and often lightweight properties. Container technologies, such as Docker, provide most--if not all--of the mentioned benefits. However, some limitations of containerization arise when high performance computing (HPC) is involved. Docker requires the user to be root which is not allowable in most HPC resources due to security concerns. However, Singularity containers inherit the user's permissions and prevent unauthorized deletion or modification of valuable data within the containers--while keeping the container lightweight and HPC supportive. Singularity is installed and made available on HPC resources such as XSEDE Comet HPC cluster.



### COMET

Comet is a HPC resource that was designed to be operated for computational research performed at modest scale. Comet also supports science gateways, which are web-based applications that simplify access to HPC resources on behalf of a diverse range of research communities and domains, typically with hundreds to thousands of users

### SINGULARITY

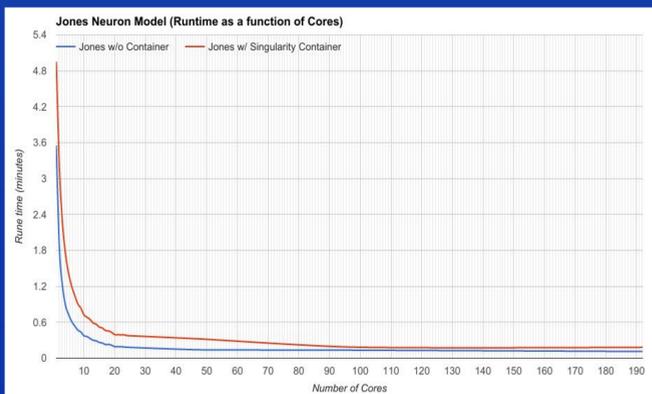
- Encapsulation of environments
- Image based containers
- No root escalation allowed (permissions are inherited)

### RESULTS

The output obtained from running the Jones Model using NEURON (a computational neuronal simulation code) suggests a noticeable overhead when running on a single core; however, as the number of cores increase, the overhead becomes constant and almost negligible.

We ran the whole Intel MPI Benchmark (IMB) suites using Singularity and non-Containerized options. Although IMB's Sendrecv suggests nearly identical data transfer rate performance, a more noticeable difference can be observed in OSU's Average Latency Benchmark. In OSU's Average Latency Benchmark, the initial latency between the containerized and non-containerized job runs is nearly negligible; however, once the input size exceeds 400,000 bytes, the non-containerized begins to deviate at a slightly higher rate. Although these outputs still differ by a small margin, this is more notable difference that can be explained by the container's overhead.

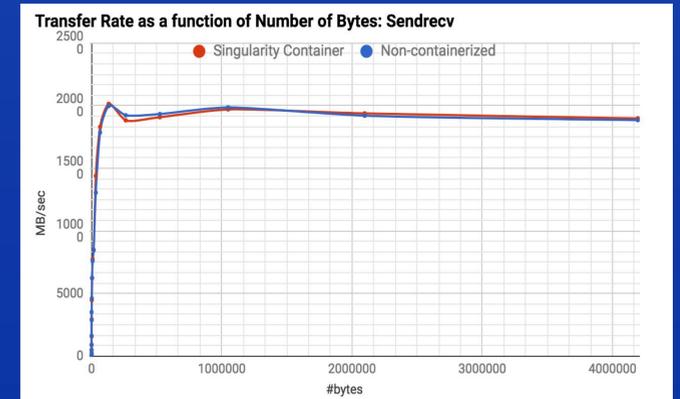
### NEURON: Jones Model



### OSU



### Intel MPI Benchmark (IMB): Sendrecv



### CONCLUSION:

Singularity features encapsulation, portability, reproducibility and prevention of root escalation privileges. This makes it ideal for use on supercomputers because once a container image is transferred over it does not permit privileged operations with root access. This prevents users from creating and executing malicious software on a supercomputer's shared resources. Singularity's performance results from high performance computing benchmarks suggest that its negligible performance overhead should not be a significant decision factor when considering its intended applications as the performance is fairly close to non-containerized jobs.

### Acknowledgements



**Mentors:**  
Amit Majumdar, Subhashini Sivagnanam  
Mahidhar Tatineni, Kenneth Yoshimoto,  
San Diego Supercomputer Center,  
University of California San Diego

We would like to acknowledge the partial support for this work provided by the National Science Foundation grant #1339856. We would like to thank XSEDE allocation on Comet.