

## **InfiniCortex:**

A path to reach Exascale concurrent supercomputing across the globe utilising  
trans-continental InfiniBand  
and  
Galaxy of Supercomputers

Marek Michalewicz, Yves Poppe, Tan Tin Wee, Yuefan Deng  
A\*STAR Computational Resource Centre  
Singapore

A position paper for the 3<sup>rd</sup> Big Data and Exascale Computing,  
Barcelona, January 2015

We have developed the idea, the method and the prescription; and have demonstrated the first implementation of a concept called **InfiniCortex** where a set of geographically distributed High Performance Computing and Storage resources can provide concurrent supercomputing and pave the way to Exascale in a concurrent and distributed fashion, alleviating the pressures of colossal power and infrastructure requirements, not to mention associated data replication and disaster recovery issues associated with a centralized approach.

The approach *is not* a grid or cloud based, but utilises extremely efficient, lossless and encrypted InfiniBand transport technology over global distances allowing RDMA and straightforward implementation of both concurrent supercomputing over global distances and implementation of very efficient workflows – and here it serves as an ideal vehicle to serve both Big Data and Exascale computing requirements.

### **InfiniCortex:**

SC14 in New Orleans has offered the ideal platform to demonstrate the merging of four key concepts, integrated for the first time together to realise the InfiniCortex demonstration:

- 1) Very high bandwidth (100 Gbps) intercontinental connectivity between Asia and the USA;
- 2) InfiniBand technology over trans-continental distances using Obsidian Strategics range extenders including routing and BGFC based subnetting;
- 3) Connecting separate InfiniBand sub-nets with different net topologies to create a single and topologically optimised computational resource: Galaxy of Supercomputers;
- 4) Running six different workflows and applications on such a distributed, concurrent computational infrastructure, especially using ADIOS I/O framework.

In a world premiere, 100Gbps trans-Pacific connectivity was demonstrated with InfiniBand transport layer, connecting three continents: Asia, Australia and North America; four countries: Singapore, Australia, Japan and the USA; seven Universities and two large research organisations: Agency for Science, Technology and Research (A\*STAR) Singapore and Oak

Ridge National Laboratory (ORNL) USA. The demonstration took place at the Supercomputing 2014 conference in New Orleans during the period of 17th -20th November 2014.

The demo set up of InfiniCortex comprised computing resources residing in Singapore (A\*CRC), Japan (Tokyo Tech), Australia (ANU) and in the United States (Georgia Tech) as well as on the floor of SC14 Exhibition hall, at the A\*CRC booth. Several applications ranging from real-time confinement analysis in nuclear fusion experiments, medical image analysis in cancer diagnostics, deformations of nanostructures in materials science, genome sequence alignment in genomics, were successfully run either as workflows or concurrent distributed applications.

A relevant point to our BDEC discussions is the fact that long distance InfiniBand has demonstrated spectacular gains in effective throughput between remote HPC facilities as compared to the use of traditional TCP/IP based file transfer (FTP) techniques. Transmission of a 1.143Tbyte file of genomics data over a dedicated a 26,000km, 10gbps connection between ANU in Canberra and A\*STAR CRC in Singapore via Seattle, was reduced from 12 hours 33 minutes to 24 minutes.

In order to actively and constructively participate and contribute to big science as well as the paving of the road to Exascale, we invite all participants and members of BDEC to participate in our trials and experiments to familiarize themselves and evaluate the merits of InfiniCortex and long distance infiniband. We also propose to form a group of likeminded HPC practitioners to work with our respective NREN's to plan the implementation of global long distance InfiniBand transmission capabilities, both nationally and internationally.

As a result of SC14, projects are currently ongoing or are in planning stages between Singapore and a number of American and at least five European universities and institutions as well as demonstrations of applications, using long distance InfiniBand and InfiniCortex, between Singapore and Europe (both eastward, and via trans-Pacific + trans-continental USA + trans-Atlantic links) at the GEANT Association annual conference in Porto, Portugal in June 2015 and at ISC15 in Frankfurt, Germany in July 2015. We place particular importance to demonstrating new and relevant application with each of our new InfiniCortex partners. For example, we will be demonstrating GPGPU applications with Reims University in France, asynchronous linear solvers with University of Lille, and globally distributed weather and climate modelling together with real-time visualisation of the workflow progress with ICM Warsaw, Poland, and several other applications with others.

Funding for the Singapore NSCC (National SuperComputing Centre) has been approved and tender will be open soon for a 1-to-3 PetaFLOP machine to be operational in the third quarter of 2015. Connectivity to Europe, Japan and the USA for the specific purpose of HPC collaboration is also important part of this initiative.

On the Big Data side, major content providers as well as some top tier data centres and cloud providers have demonstrated interest in long distance InfiniBand as a means to improve their data replication and disaster recovery requirements. Several entities involved in Smart City projects in Singapore and China have also expressed interest.